



Scientific Workflows

Methods, Applications, Research

Ulf Leser, Humboldt-Universität zu Berlin



Knowledge Management in Bioinformatics

- Statistical Bioinformatics
- Data Integration, Databases
- Biomedical Text Mining
- Scientific Workflows



TRR 54



MapTorNet



T-SYS

OncoPath



BIOBANKCLOUD
Your PaaS for Biobanking



Big Data Processing

Big Data

? TBs of
data every day




12+ TBs
of tweet data
every day



25+ TBs
of
log data every
day

30 billion RFID
tags today
(1.3B in 2005)



76 million smart
meters in 2009...
200M by 2014



4.6 billion
camera
phones
world
wide



**100s of
million
s of
GPS
enable
d devices
sold
annually**
24 billion
people on
the Web
by end
2011



<http://www.>

© K. Kannan, IBM Research Labs, 2013

Big Scientific Data



Large Hydron Collider at CERN:
15 PB a year

Australian Square kilometer
array: 7,5TB/sec =
230 exabyte/year



Supercomputers to the Rescue



Also Big Scientific Data



NovaSeq Genome Sequencer: 3TB/day

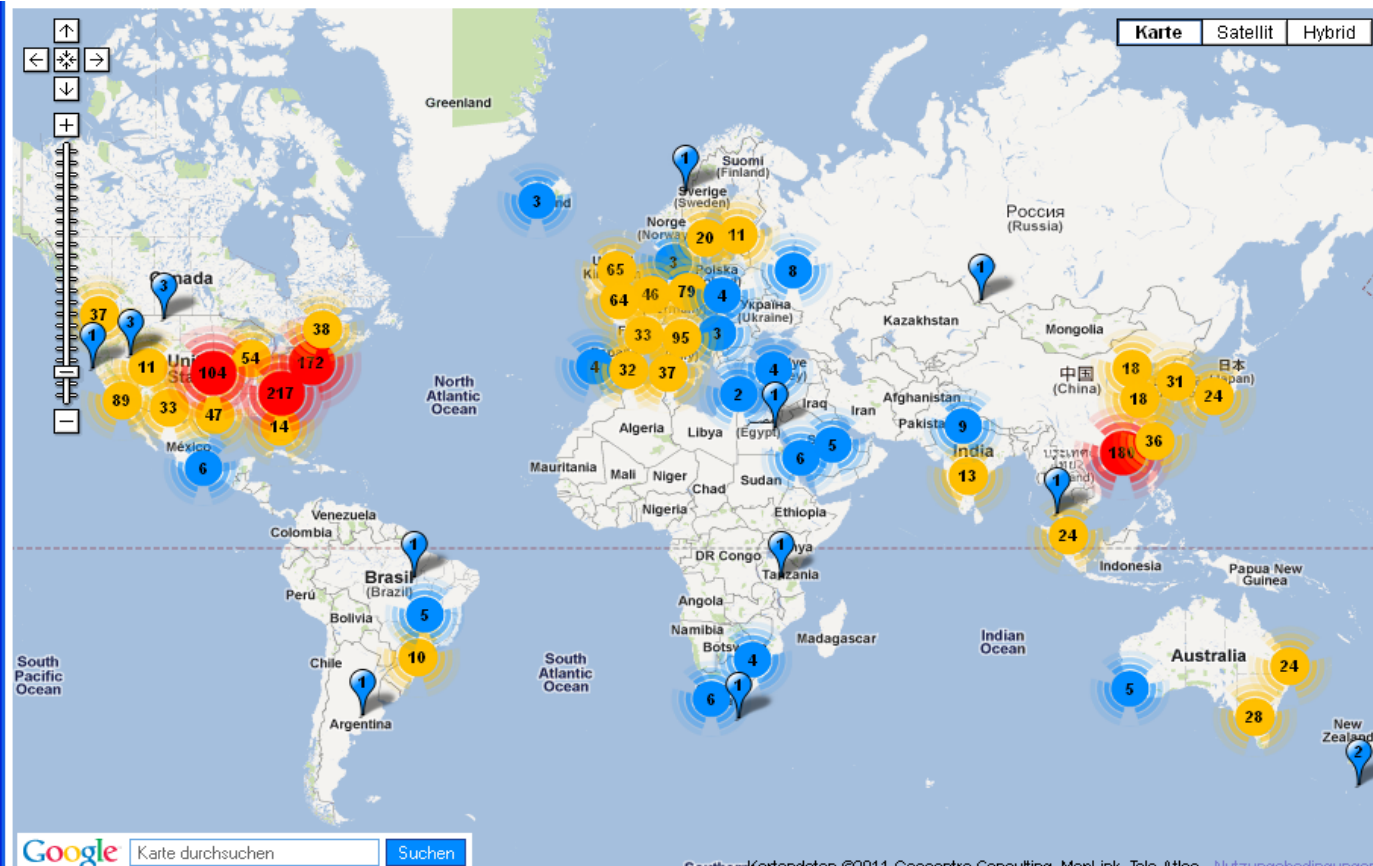
Entire human genome in one day

First human genome: 12 years

Everyday Big Scientific Data

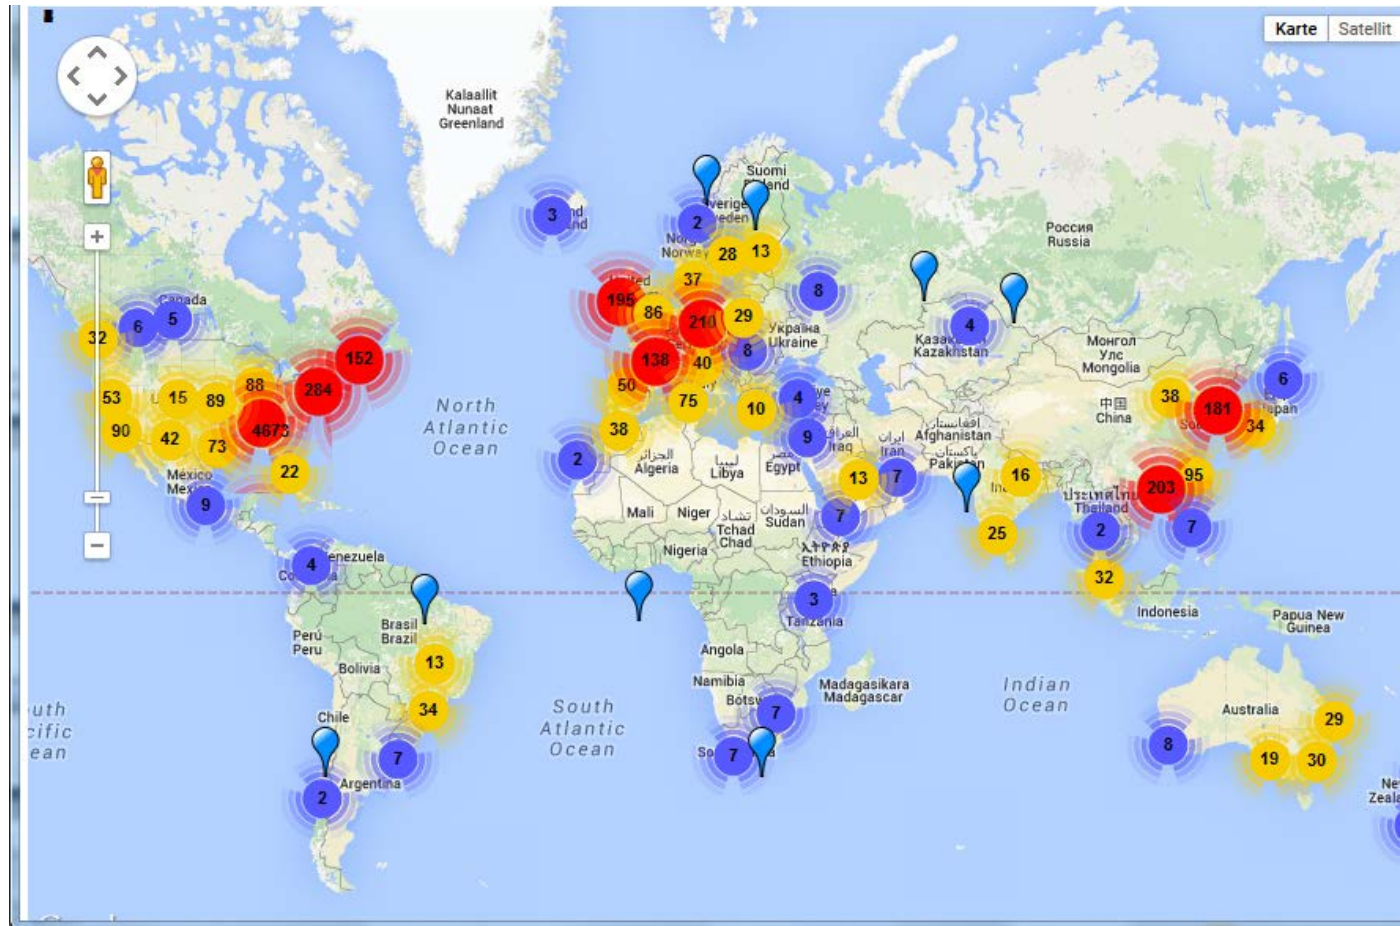


<http://pathogenomics.bham.ac.uk/hts/>



<http://pathogenomics.bham.ac.uk/hts/>


Everyday Big Scientific Data



<http://pathogenomics.bham.ac.uk/hts/>


Genome Sequencing World-Wide

All cells of all humans?



Stephens ZD, Lee SY, Faghri F, Campbell RH, Zhai C, et al. (2015) Big Data: Astronomical or Genomical?. PLOS Biology 13(7): e1002195. <https://doi.org/10.1371/journal.pbio.1002195>

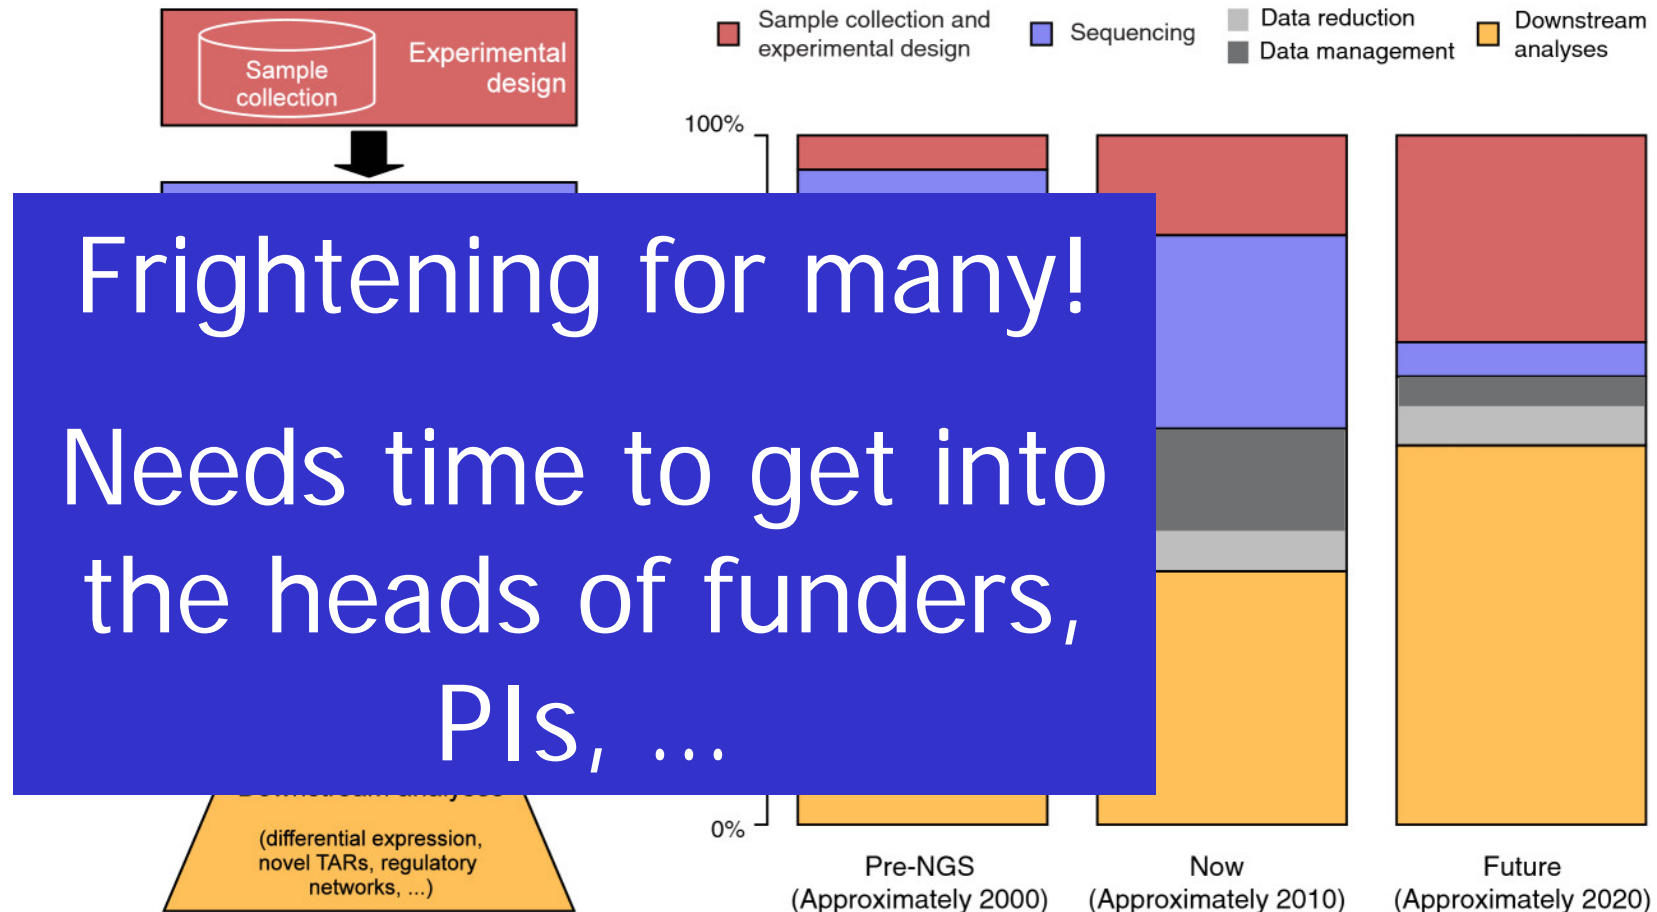
Possible Through Cost Reduction



What
does
this
mean
?

<http://www.genome.gov>

The „real“ Cost of Genomic Sequencing



Sboner, A. (2011). The real cost of sequencing: higher than you think! Genome Biology 2011



Will Computers Crash Genomics?

you-go service, accessible from one's own desktop, that provides rented time on a large cluster of machines that work together in par-

Pennisi, E. (2011). *Science*

sciencemag.org on February 10, 2011

Supercomputers to the Rescue?



Consequences

- Not every hospital will buy a supercomputer
- Actually: Google, Facebook, Twitter etc. don't buy supercomputers
- Do we always need a supercomputer for Big Data?

Data Parallel Problems

Example: SNV Detection

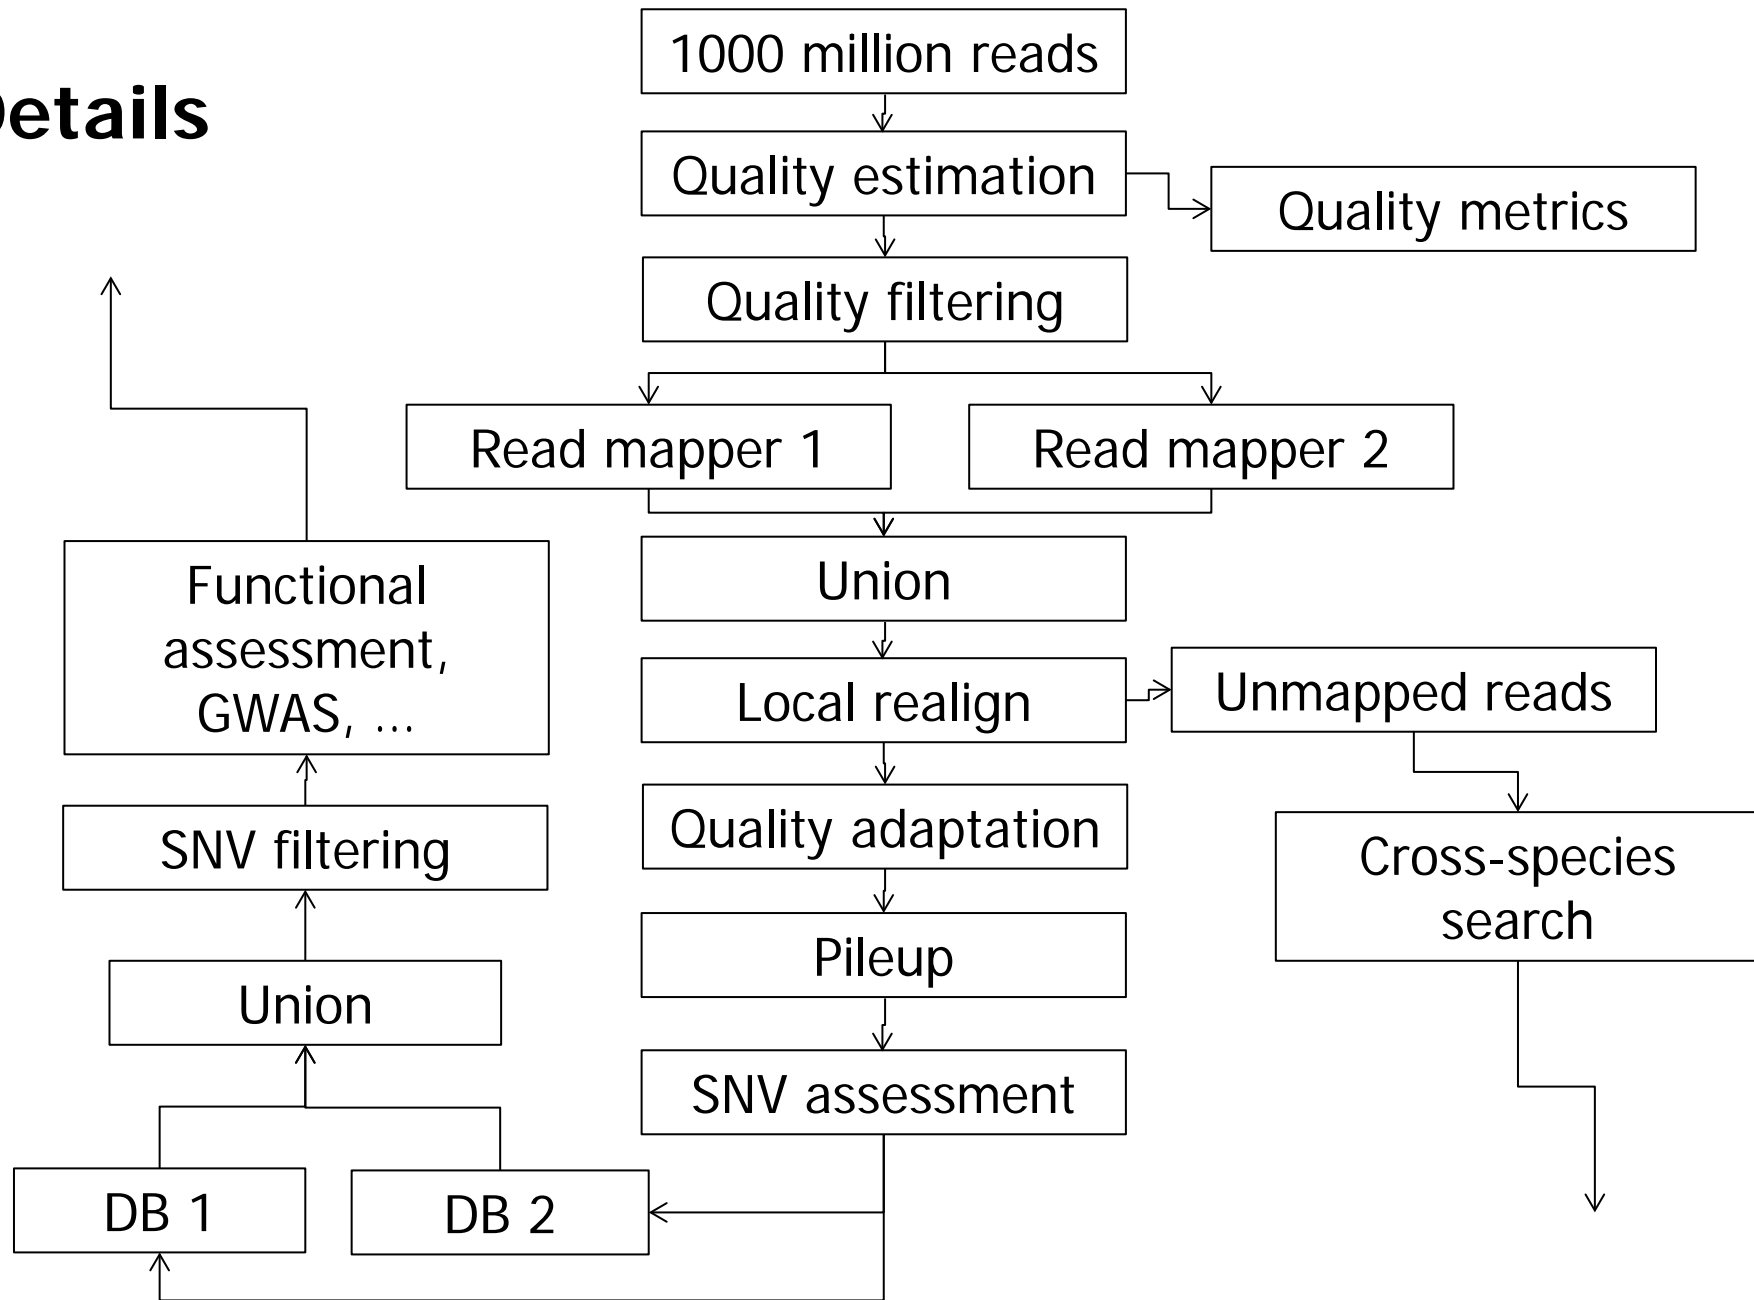


DNA/RNA

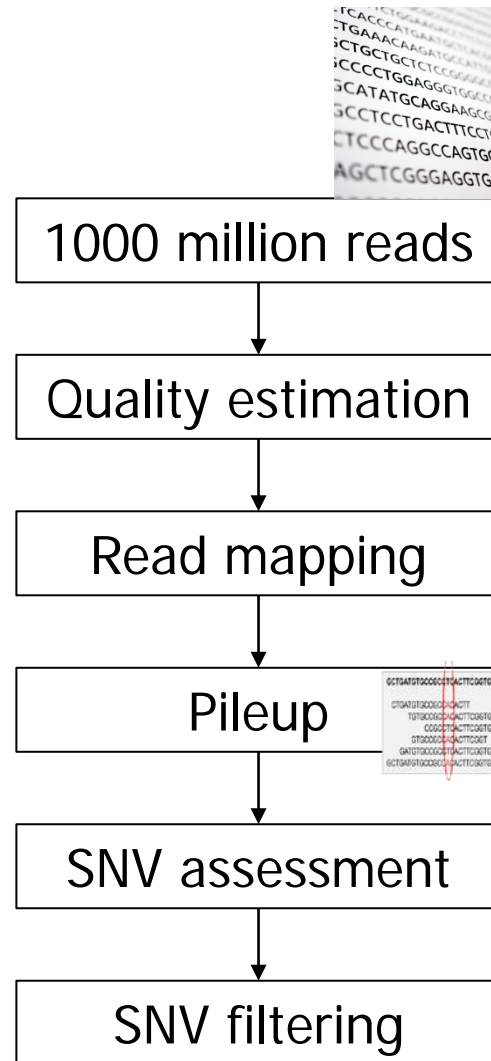



GCTGATGTGCCGCTCACTTCGGTGAGGTG	Reference sequence
CTGATGTGCCGCCCACTT	read 1
TGTGCCGCCCACTTCGGTGGT	read 2
CCGCCTCACTTCGGTGAGGT	read 3
GTGCCGCCCACTTCGGT	read 4
GATGTGCCGCTCACTTCGGTGGA	read 5
GCTGATGTGCCGCCCACTTCGGTGAGGT	read 6


Details



Size of Data




Degree of Parallelism



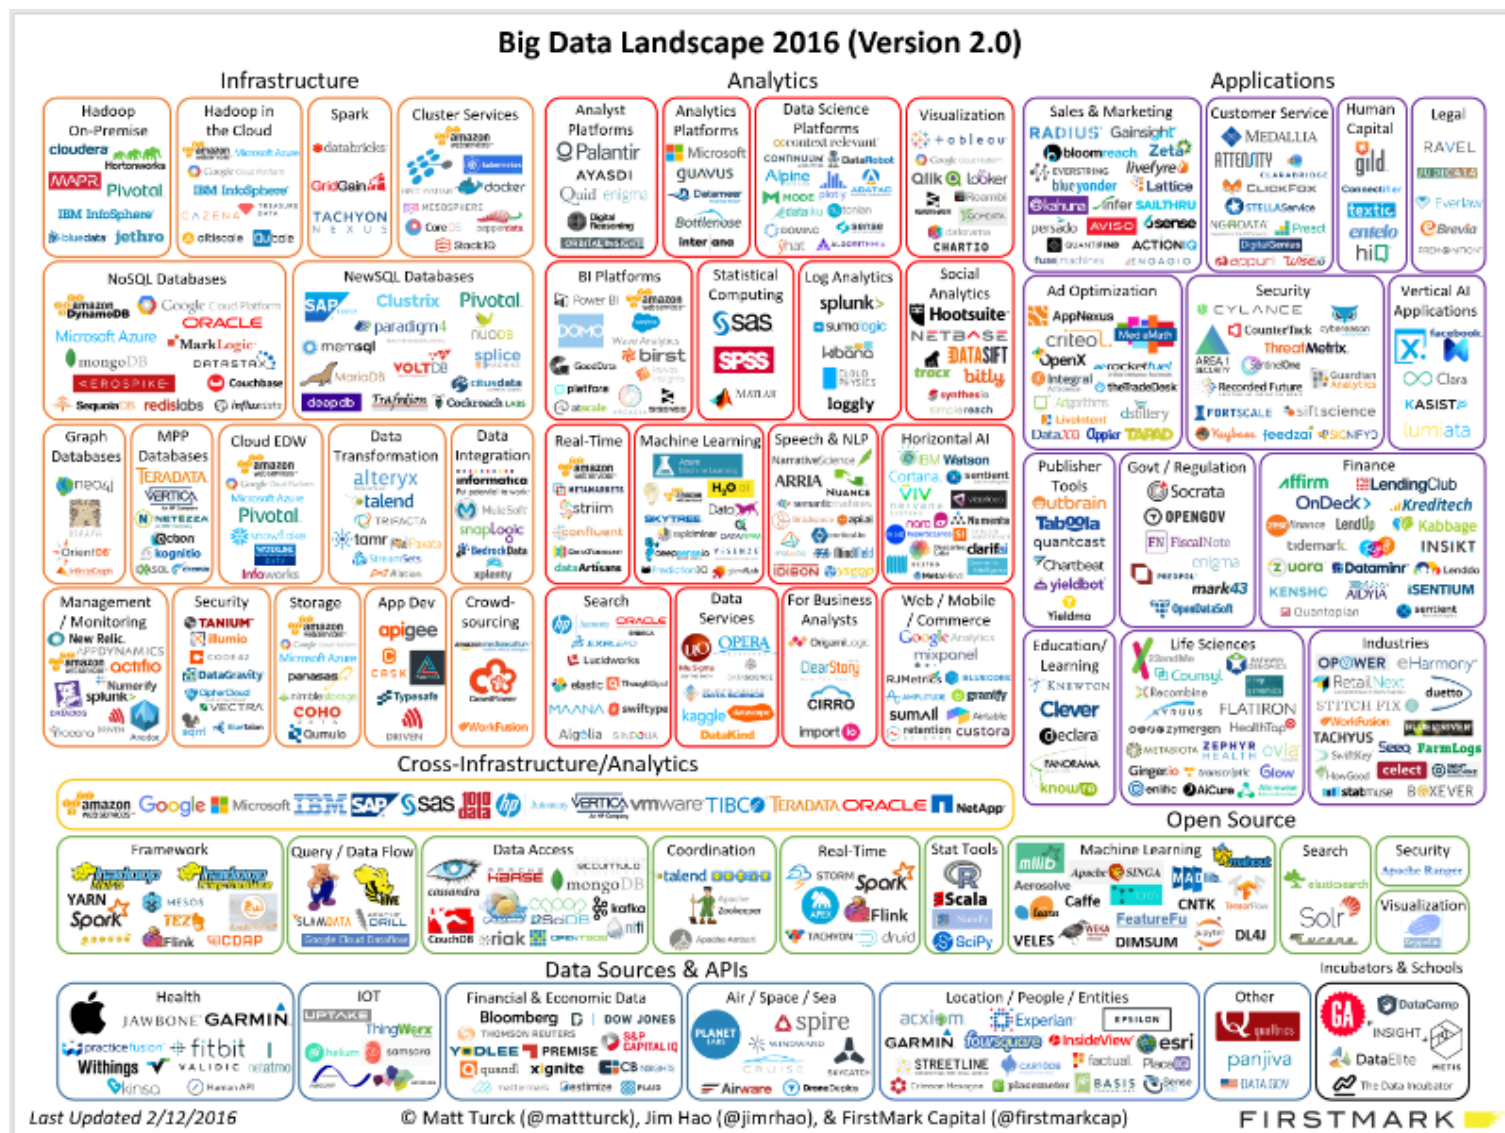
Embarrassingly Parallel

- Many scientific data analysis problems are mostly **embarrassingly parallel**
 - Exceptions: Simulation
- Provided the right software, these can be run perfectly on clusters of **commodity hardware**
 - Many cheap machines
 - Many examples

SNV W



Ulf Leser: Scientific Workflows, DESY, 03/2018



Embarrassingly Parallel

- Many scientific data analysis problems are mostly embarrassingly parallel
 - Exceptions: Simulation
- Provided the right software, these can be run perfectly on clusters of commodity hardware
 - Many cheap machines
 - Examples: Hadoop, Mesos, Flink, Spark, ...
- Will it be fast “enough”?

	sonic	dbis	rubix
type	monolithic	cluster	cluster
SWfMS	SaaSFee	Saasfee	Snakemake
nodes	1	24	111
total threads	80	576	3784
memory per node	512 GB	24 / 32 GB	128 / 188 / 500 / 1000 GB
cost	11,535 €	94,719 €	~ 2,000,000 €
exec. time	24 h	7.62 h	1.14 h
throughput (exec. per year)	365	1149.61	7684.21
throughput / €	0.0316	0.0121	0.0038

	sonic	dbis	rubix	EC2 (estimated)
type	monolithic	cluster	cluster	cloud
SWfMS	SaaSFee	Saasfee	Snakemake	Saasfee
nodes	1	24	111	16 (r3.4xlarge)
total threads	80	576	3784	256
memory per node	512 GB	24 / 32 GB	128 / 188 / 500 / 1000 GB	122 GB
cost	11,535 €	94,719 €	~ 2,000,000 €	~ 500 € (per run)
exec. time	24 h	7.62 h	1.14 h	~ 14 h
throughput (exec. per year)	365	1149.61	7684.21	arbitrary
throughput / €	0.0316	0.0121	0.0038	


Downsides

- You don't need **expensive hardware**, but ..
- You must care about failures
- You must care about IO / network
- You need to know how to program distributed systems
- You need **expensive developers!**
 - "... **human productivity** [in large-scale scientific data analysis] arguably still is the **most expensive resource**, trumping power, performance, and other factors ..."
[NSF/DOE Workshop 2015]

Cost of Implementing Pipelines

Published algorithm, all
tools available as open
source:


~3 PM implementation



	sonic	dbis	rubix	EC2 (estimated)
type	monolithic	cluster	cluster	cloud
SWfMS	SaaSfee	Saasfee	Snakemake	Saasfee
nodes	1	24	111	16 (r3.4xlarge)
total threads	80	576	3784	256
memory per node	512 GB	24 / 32 GB	128 / 188 / 500 / 1000	122 GB
cost	11,535 €	94,719 €	~ 2,000,000 €	~ 500 € (per run)
exec. time	24 h	7.62 h	1.14 h	~ 14 h
throughput	365	1149.61	7684.21	arbitrary
throughput / €	0.0316	0.0121	0.0038	

Cost of Implementing Pipelines


Experts available on
both systems:
~2 PM port



	sonic	dbis	rubix	EC2 (estimated)
type	monolithic	cluster	cluster	cloud
SWfMS	SaaSfee	Saasfee	Snakemake	Saasfee
nodes	1	24	111	16 (r3.4xlarge)
total threads	80	576	3784	256
memory per node	512 GB	24 / 32 GB	128 / 188 / 500 / 1000	122 GB
cost	11,535 €	94,719 €	~ 2,000,000 €	~ 500 € (per run)
exec. time	24 h	7.62 h	1.14 h	~ 14 h
throughput	365	1149.61	7684.21	arbitrary
throughput / €	0.0316	0.0121	0.0038	

Cost of Implementing Pipelines

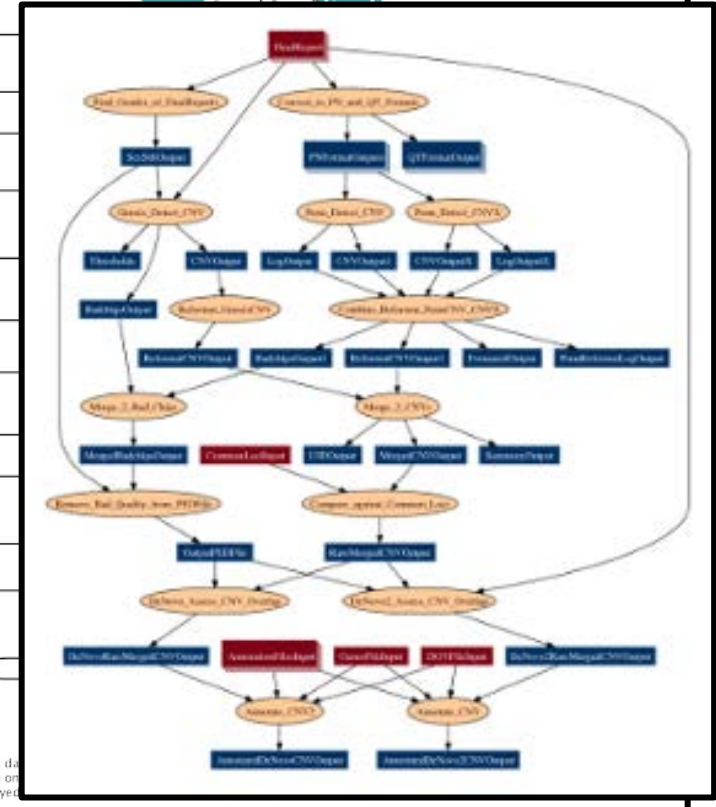
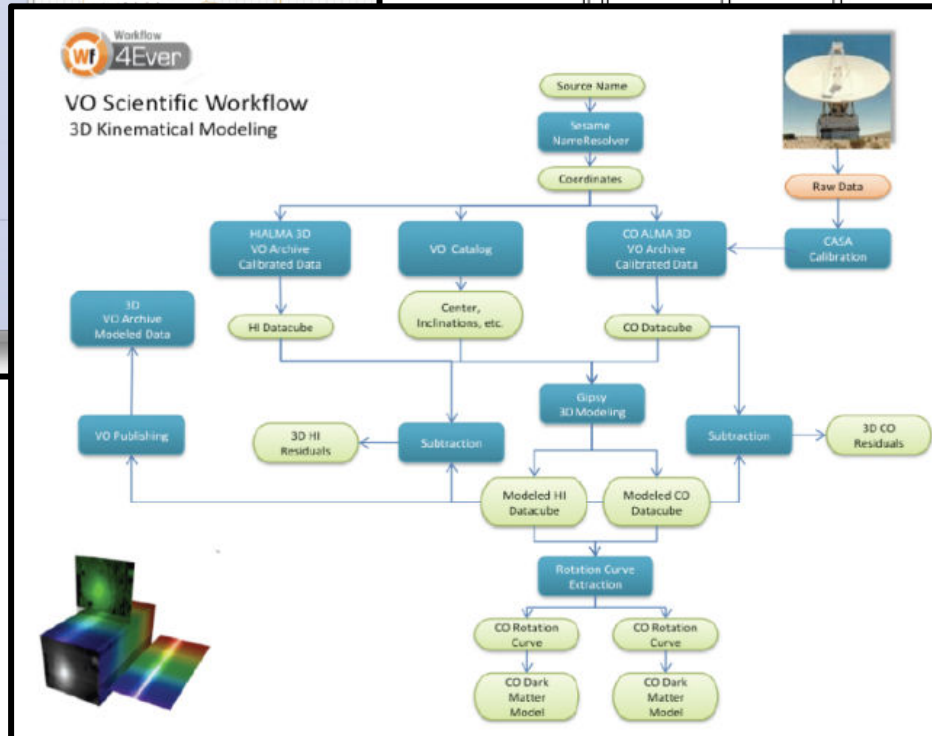
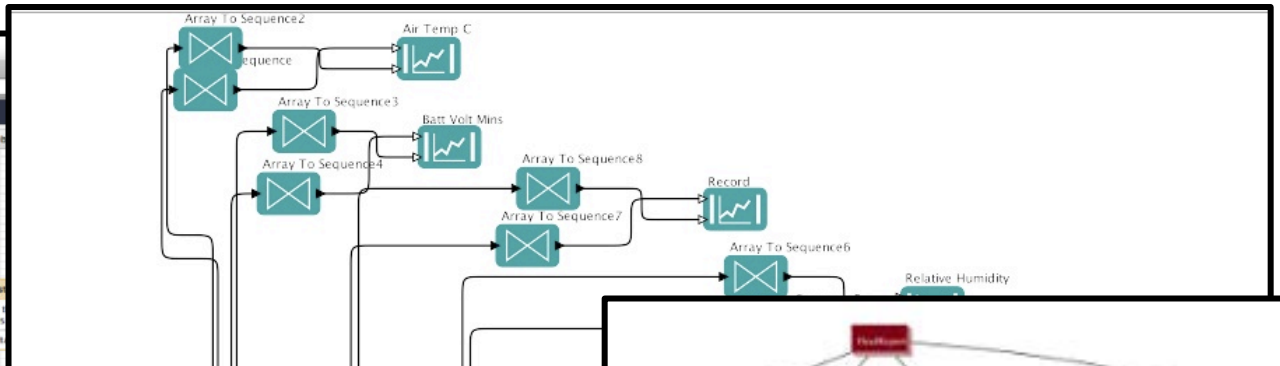
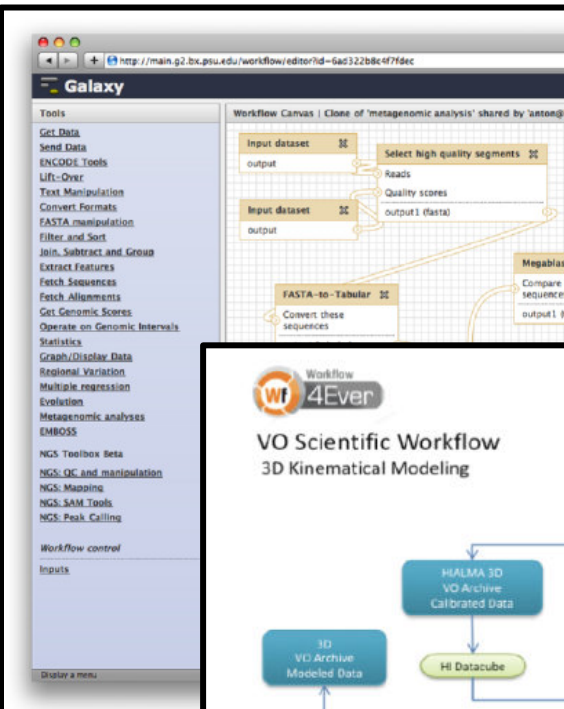
System supports
different engines:
~0 PM port



	sonic	dbis	rubix	EC2 (estimated)
type	monolithic	cluster	cluster	cloud
SWfMS	SaaSfee	Saasfee	Snakemake	Saasfee
nodes	1	24	111	16 (r3.4xlarge)
total threads	80	576	3784	256
memory per node	512 GB	24 / 32 GB	128 / 188 / 500 / 1000	122 GB
cost	11,535 €	94,719 €	~ 2,000,000 €	~ 500 € (per run)
exec. time	24 h	7.62 h	1.14 h	~ 14 h
throughput	365	1149.61	7684.21	arbitrary
throughput / €	0.0316	0.0121	0.0038	

Scientific Workflow Management Systems

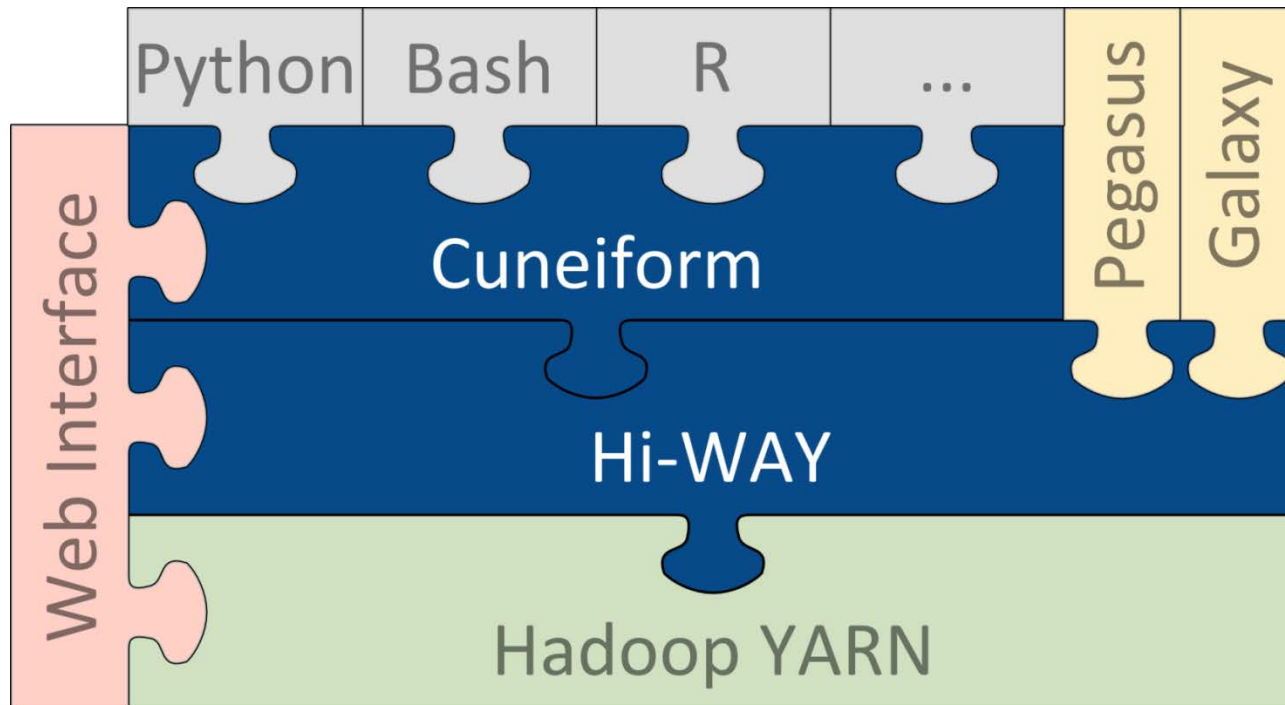
Many!



Scientific Workflows for Data Analysis

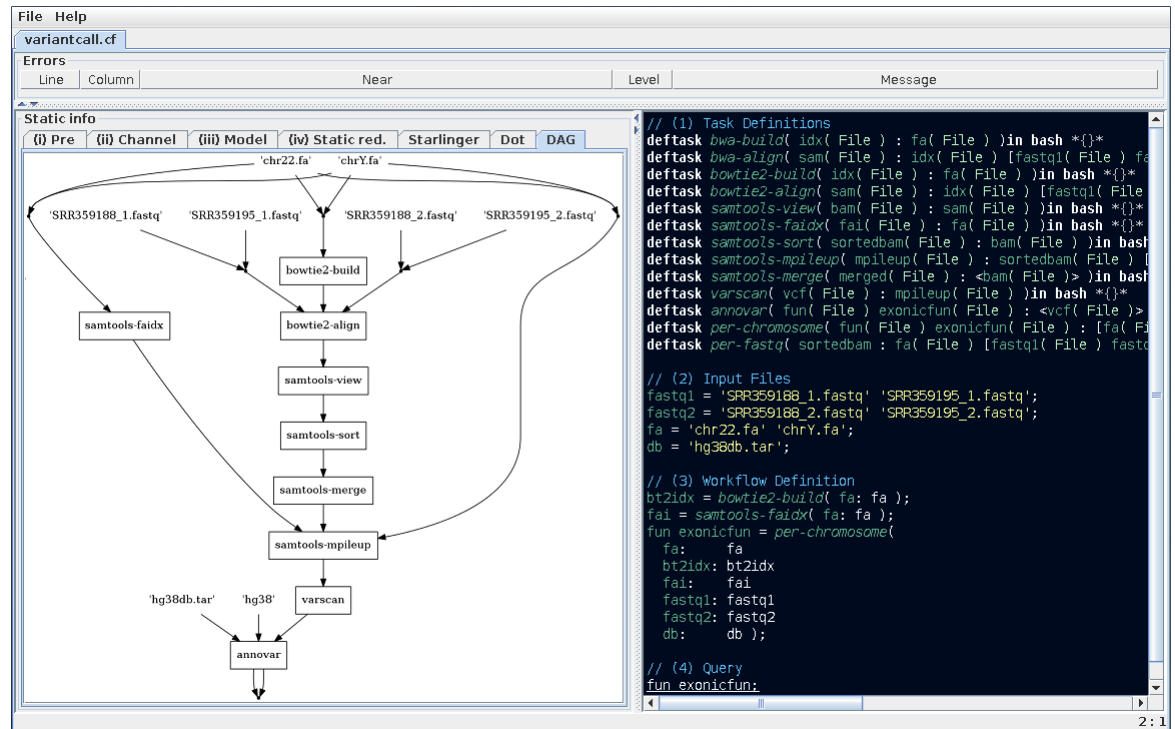
- Data parallel problems
 - Scheduling on distributed execution engines
- One-shot batch processing
 - No streaming, no data management, no indexing, ...
- DAG-structured workflows of black-box tasks
 - Easy to create, hard to optimize
- Commodity hardware: Fault Tolerance & slow IO
 - Replicate data, control activities of worker nodes
- Scientific principles: Reproducibility
 - Provenance management and analysis
- Scientific workflow management system

SaaSFee: Merging MapReduce and SWF



www.saasfee.io

Cuneiform



- Light-weight statically typed functional dataflow language
- Compiles into dynamic pipelines of black-box tools
- Make **foreign code integration** as easy as possible
- Allow complex, **iterative workflows**

Foreign Code Interface

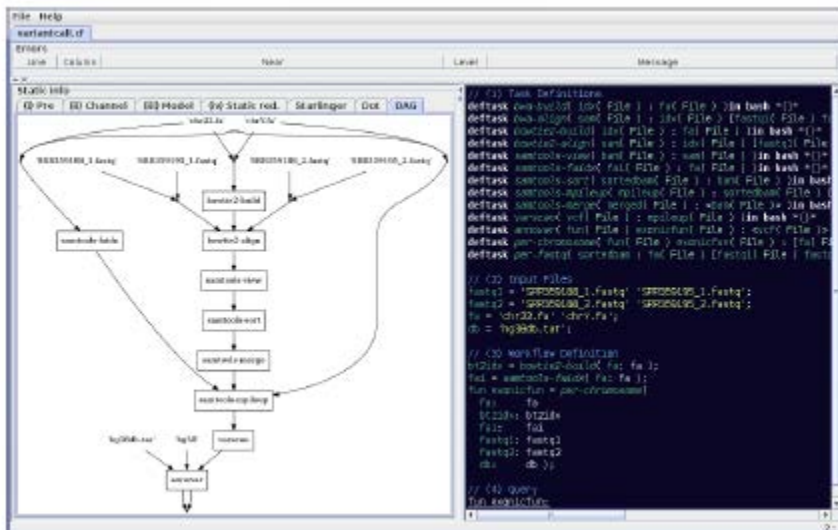
- **Directly integrates** BASH, LISP, R, MatLab, Python ...
- No wrapping, no data (un)marshalling, no API
- Communication via variables or files
- **Mixing of several languages**
- Snippets are shipped and executed by Hi-Way

```
deftask greet( out : person )in bash *{  
    out="Hello $person"  
}*  
deftask greet( out : person )in r *{  
    out = paste( "Hello", person )  
}*
```

Achieving Parallelism in Cuneiform

Task Parallelism

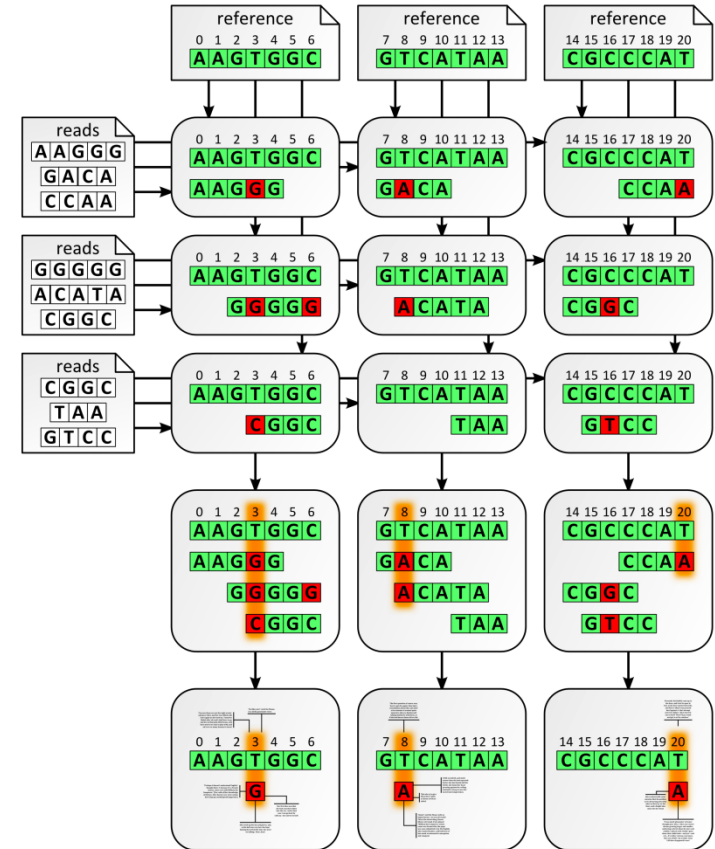
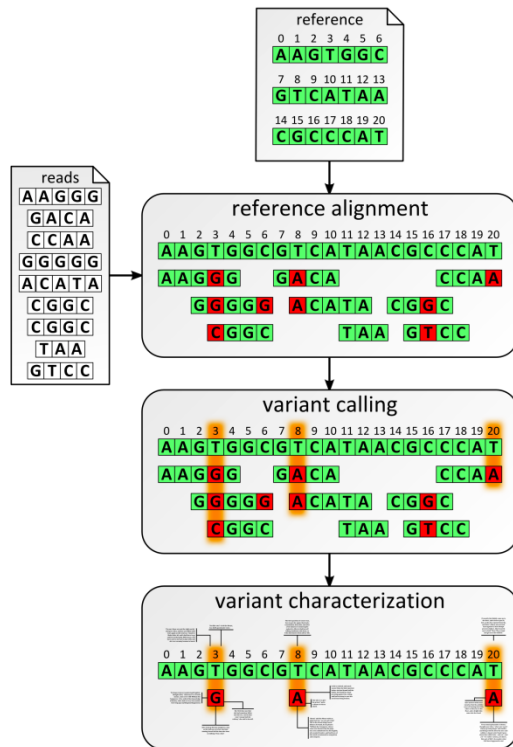
- Data dependencies



Data Parallelism

- Custom partitioning
 - Default for record-oriented files
- Algorithmic skeletons
 - Map, cross-product, dot-product, aggregation, ...
- Subsumes Map&Reduce

Parallelism



Dynamic Workflows

- No static execution plan but graph reduction
- Hi-Way polls ready-to-exec tasks, reports completed tasks

```
deftask classify( labeled( File )
    : dataset( File ) meanset( File ) )in python
deftask refine( meanset1( File )
    : labeled( File ) )in python
deftask hasconverged( <converged> q1
    : dataset( File ) meanset( File ) q )in python
deftask kmeans( result( File )
    : dataset( File ) meanset( File ) q ) {
    labeled = classify( dataset: dataset
                        meanset: meanset );
    meanset1 = refine( labeled: labeled );
    converged q1 = hasconverged( dataset: dataset
                                meanset: meanset1
                                q:      q );
    result = if converged ( result : ) then {
        result = meanset1;
    } else {
        result = kmeans( dataset: dataset
                        meanset: meanset1
                        q:      q1 );
    }
}
```

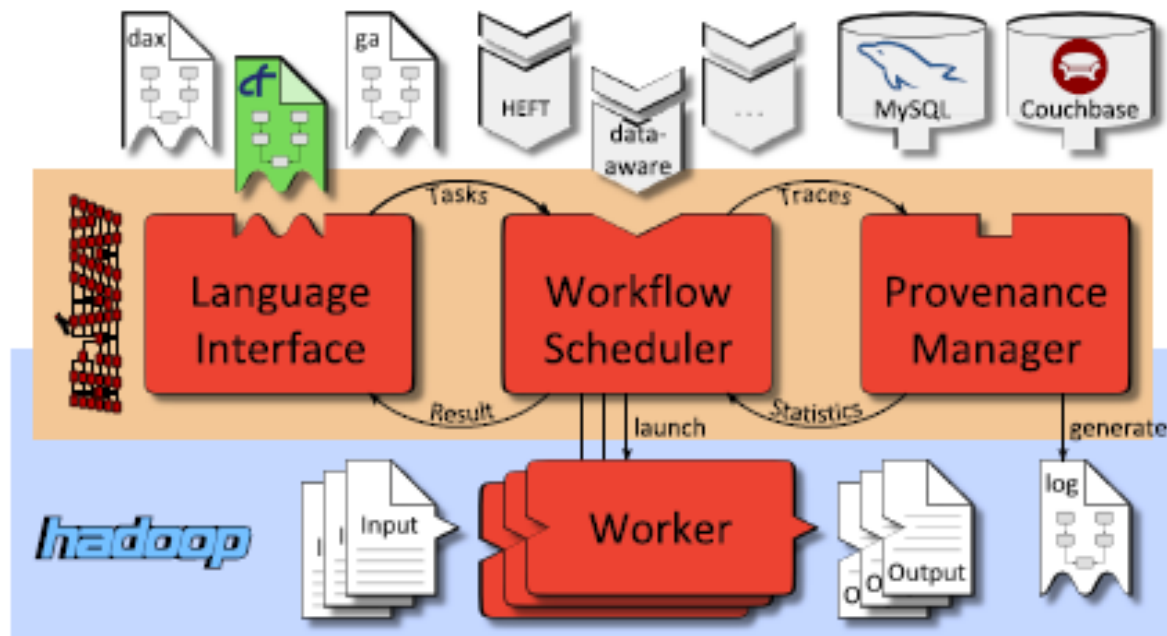
classify
refine
hasconverged

classify
refine
hasconverged

classify
refine
hasconverged

Hi-Way

- Hi-Way Workflow Application Master for YARN
- Executes workflows on Hadoop YARN
 - Scalability, maintenance, fault tolerance, ...
- Full provenance tracing, **executable provenance**
- Runs Cuneiform, Galaxy, Pegasus (DAX)
- Various (adaptive) schedulers
- Dynamic workflow interface




Research Topics (at WBI)

- Workflow Similarity Search
- Workflow Optimization (by reordering)
- Workflow Scheduling
- Workflow Scheduling using Online Learning

All Analytics is the Same?

ROUTINELY UNIQUE

Over 18 months, 46 data-analysis projects undertaken at the bioinformatics core of the University of Texas Health Science Center at Houston required 34 different types of analysis — most were used infrequently. Each project demanded unique combinations of analyses, demonstrating how bioinformaticians must be versatile, creative and collaborative.




Chang, Core services: Reward bioinformaticians, Nature 2015

myExperiment


The screenshot shows the myExperiment website interface. The top navigation bar includes links for Home, Users, Groups, Workflows, Files, and Packs. A search bar is present with a 'GO' button. The main content area displays a list of workflows under the 'Workflows' tab. A search filter bar at the top of the workflow list shows '2388' workflows, which is circled in blue. Below this, there are two workflow cards. The first card is for 'Taverna 2 Pathways and Gene annotations for QTL region (7)' by Paul Fisher, created on 19/11/09. The second card is for 'BioAID_DiseaseDiscovery_RatHumanMouseUnipr... (4)' by Marco Roos, created on 15/12/08. Both cards include a description, a rating, and a 'Download' button. The left sidebar contains filters by type (e.g., Taverna 2, Taverna 1, RapidMiner) and by tag (e.g., example, mygrid, bioinformatics).

Workflow Similarity Search



- Naïve
- Task sets
- Topological



Base, embl_id,
fetch_emb,
save_to_file, file-
location, ...



Database, program,
maskedSeq,
blast_ddbj,
text_blast_out, ...








Named Entity Recognition

Z-100 is an *arabinomannan* extracted from *Mycobacterium tuberculosis* that has various immunomodulatory activities, such as the induction of *interleukin 12*, *interferon gamma* (*IFN-gamma*) and beta-chemokines. The effects of *Z-100* on *human immunodeficiency virus type 1* (*HIV-1*) replication in *human monocyte-derived macrophages* (*MDMs*) are investigated in this paper. In *MDMs*, *Z-100* markedly suppressed the replication of not only macrophage-tropic (M-tropic) *HIV-1* strain (*HIV-1JR-CSF*), but also *HIV-1* pseudotypes that possessed amphotropic *Moloney murine leukemia virus* or *vesicular stomatitis virus G* envelopes. *Z-100* was found to inhibit *HIV-1* expression, even when added 24 h after infection. In addition, it substantially inhibited the expression of the pNL43lucDeltaenv vector (in which the *env* gene is defective and the *nef* gene is replaced with the *firefly luciferase* gene) when this vector was transfected directly into *MDMs*. These findings suggest that *Z-100* inhibits virus replication, mainly at *HIV-1* transcription. However, *Z-100* also downregulated expression of the cell surface receptors *CD4* and *CCR5* in *MDMs*, suggesting some inhibitory effect on *HIV-1* entry. Further experiments revealed that *Z-100* induced *IFN-beta* production in these cells, resulting in induction of the 16-kDa *CCAAT/enhancer binding protein* (*C/EBP*) *beta transcription factor* that represses *HIV-1* long terminal repeat transcription. These effects were alleviated by SB 203580, a specific inhibitor of *p38 mitogen-activated protein kinases* (*MAPK*), indicating that the *p38 MAPK* signalling pathway was involved in *Z-100*-induced repression of *HIV-1* replication in *MDMs*. These findings suggest that *Z-100* might be a useful immunomodulator for control of *HIV-1* infection.

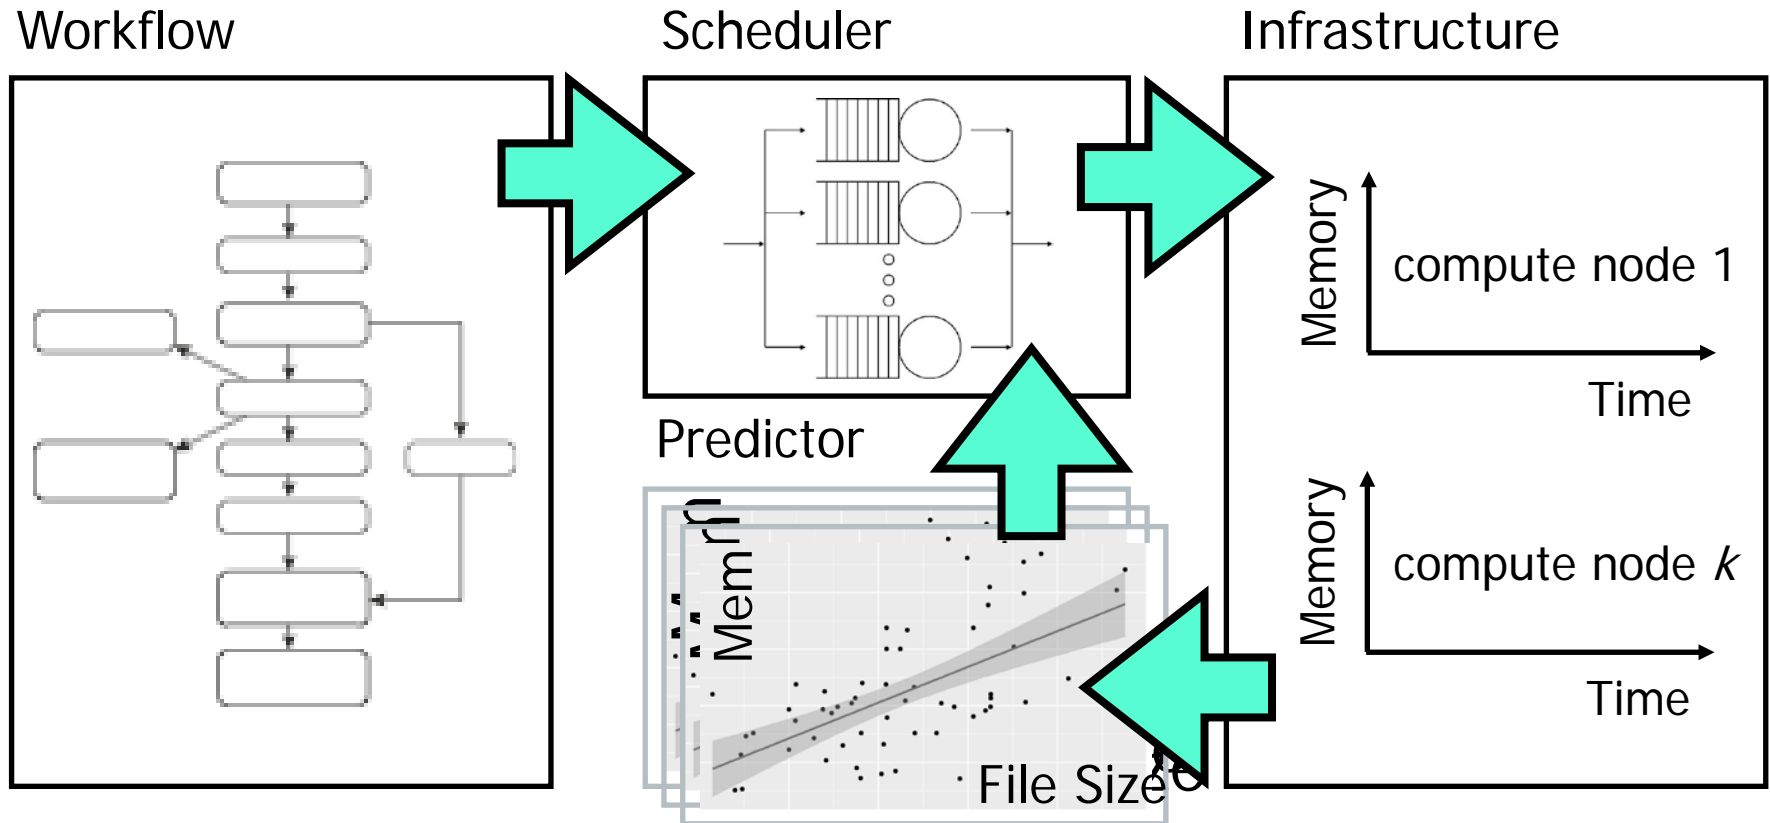
Find Relationships

Z-100 is an **arabinomannan** extracted from **Mycobacterium tuberculosis** that has various immunomodulatory activities, such as the induction of **interleukin 12**, **interferon gamma** (**IFN-gamma**) and beta-chemokines. The effects of **Z-100** on **human immunodeficiency virus type 1** (**HIV-1**) replication in **human monocyte-derived macrophages** (**MDMs**) are investigated in this paper. In **MDMs**, **Z-100** markedly suppressed the replication of not only macrophage-tropic (M-tropic) **HIV-1** strain (**HIV-1JR-CSF**), but also **HIV-1** pseudotypes that possessed amphotropic **Moloney murine leukemia virus** or **vesicular stomatitis virus G** envelopes. **Z-100** was found to inhibit **HIV-1** expression, even when added 24 h after infection. In addition, it substantially inhibited the expression of the pNL43lucDeltaenv vector (in which the **env** gene is defective and the **nef** gene is replaced with the **firefly luciferase** gene) when this vector was transfected directly into **MDMs**. These findings suggest that **Z-100** inhibits virus replication, mainly at **HIV-1** transcription. However, **Z-100** also downregulated expression of the cell surface receptors **CD4** and **CCR5** in **MDMs**, suggesting some inhibitory effect on **HIV-1** entry. Further experiments revealed that **Z-100** induced **IFN-beta** production in these cells, resulting in induction of the 16-kDa **CCAAT/enhancer binding protein** (**C/EBP**) **beta transcription factor** that represses **HIV-1** long terminal repeat transcription. These effects were alleviated by SB 203580, a specific inhibitor of **p38 mitogen-activated protein kinases** (**MAPK**), indicating that the **p38 MAPK** signalling pathway was involved in **Z-100**-induced repression of **HIV-1** replication in **MDMs**. These findings suggest that **Z-100** might be a useful immunomodulator for control of **HIV-1** infection.

Workflow Optimization




Online Scheduling



Runtimes

- Median runtime relative to lower bound



Other Research Topics

- Specification Languages: Powerful, intuitive, extensible
- Generic execution engines: De-coupling specification from execution
- Cross-layer optimization: Considering data locality and task placement
- Stream-based systems: Extreme pipelining
- Portability: Containerization / Virtualization
- Workflows for data integration
- Workflows for data visualization
- ...

Acknowledgements



Bayer HealthCare
Bayer Schering Pharma

DAAD

Deutscher Akademischer Austausch Dienst
German Academic Exchange Service



Bundesministerium
für Bildung
und Forschung



Bundesministerium
für Wirtschaft
und Technologie



Alexander von Humboldt
Stiftung/Foundation

