An introduction to

Jason Maassen, Arnold Kuzniar, Jurriaan Spaaks, Stefan Verhoeven, Johan Hidding, ...

netherlands eScience center

netherlands eScience center ?

research
computer
infrastructure

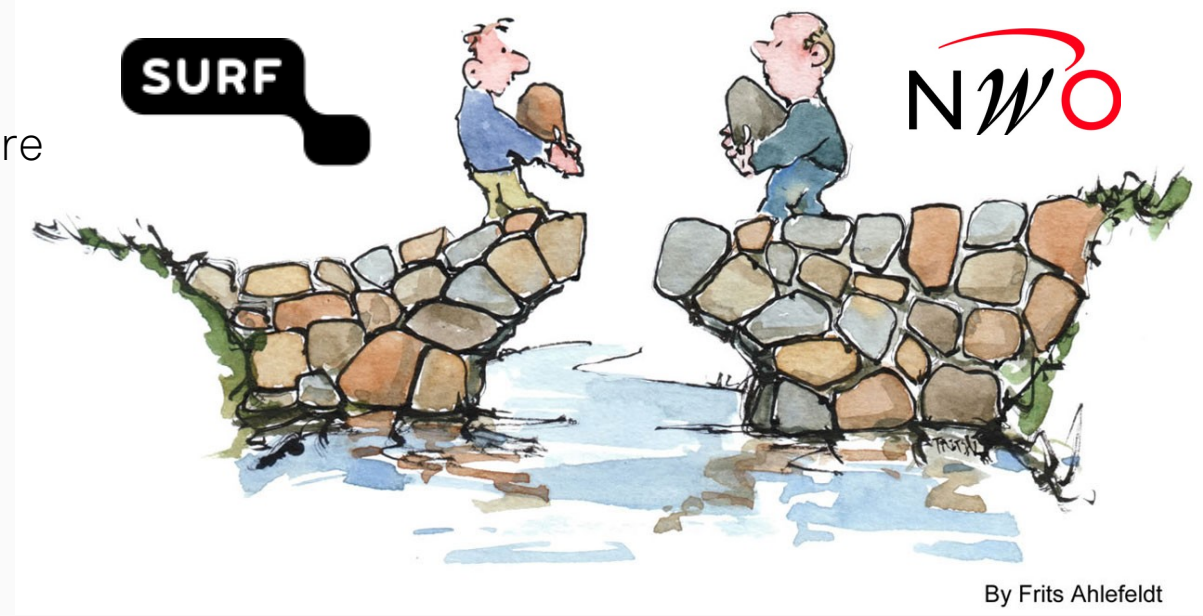

research
funding

We are an institute that provides **funding + RSEs** to research projects.

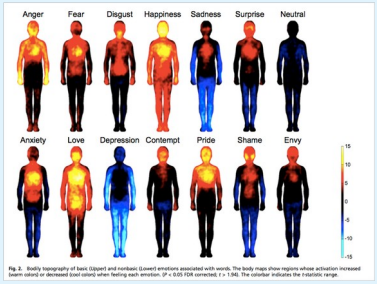

so far: ~130 projects

(on many different topics)

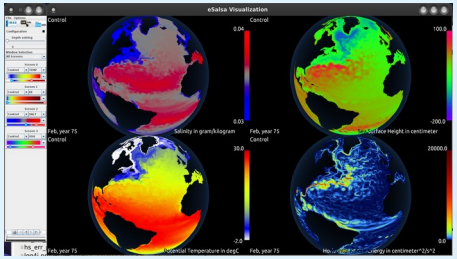

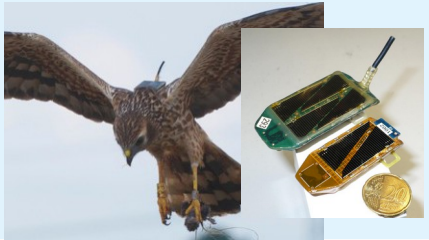

Humanities & Social Sciences

incl. SMART cities, text analysis, creative technologies

Physics & Beyond

incl. astronomy, high-energy physics, advanced materials

Sustainability & Environment

incl. climate, ecology, energy, logistics, water management

Life Sciences & eHealth

incl. bio-imaging, next generation sequencing, molecules

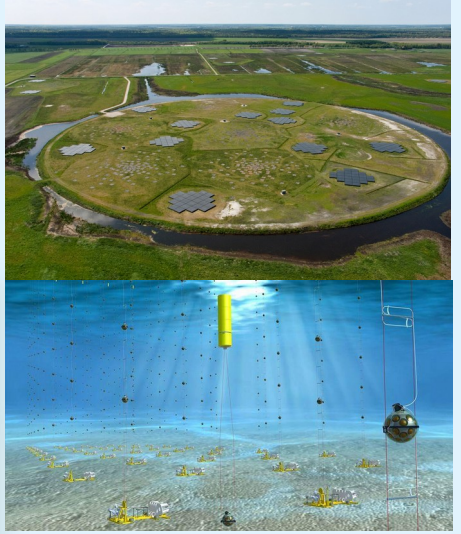

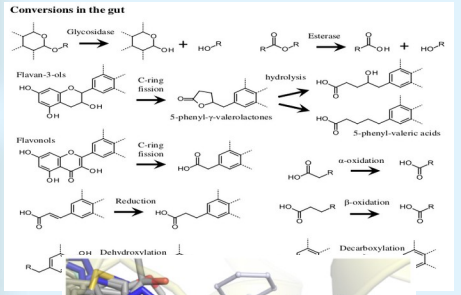

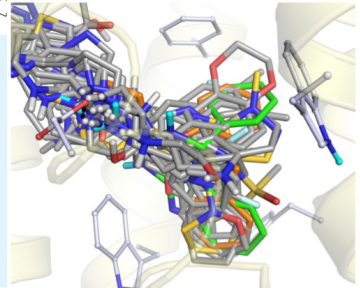

One of our goals: software re-use

We try to re-use research software between projects

- spreads cost of maintenance and development
- re-use expertise of the engineers
- increases user community
- **increases sustainability → more bang for buck!**

Works best for software low on the software stack

- generic tools are easier to reuse

Examples applications

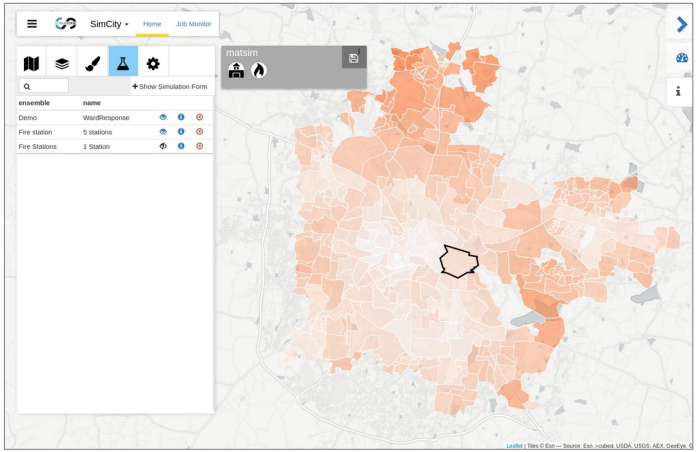

SIM-CITY

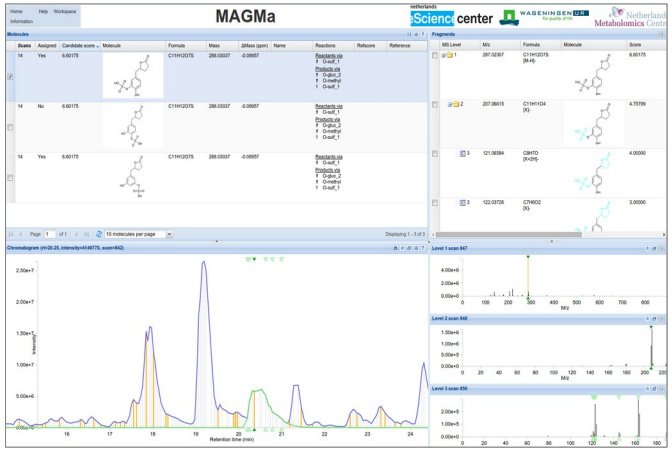

MAGMa

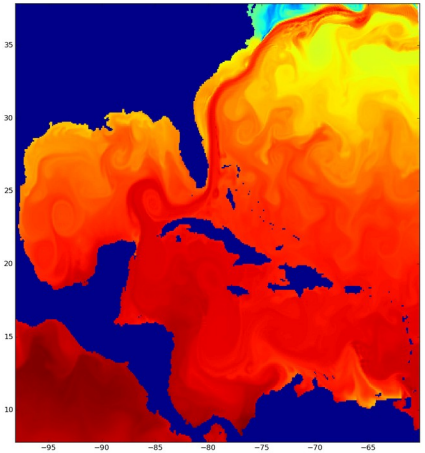

AMUSE / OMUSE

A recurring theme in many of our projects is **easy access to (remote) compute and storage.**

What is the problem ?

For many applications we need to copy data and submit jobs to remote systems. There are many ways to do this:

FTP, SFTP, WebDAV, S3, Hadoop, GridFTP, iRods, ...
Slurm, GridEngine, PBS, Torque, Amazon-Batch, ...

There are libraries and CLIs for each of these, however:

- you have to figure out how they work
- picking one will lock-in your solution
- using all of them is way too much work!

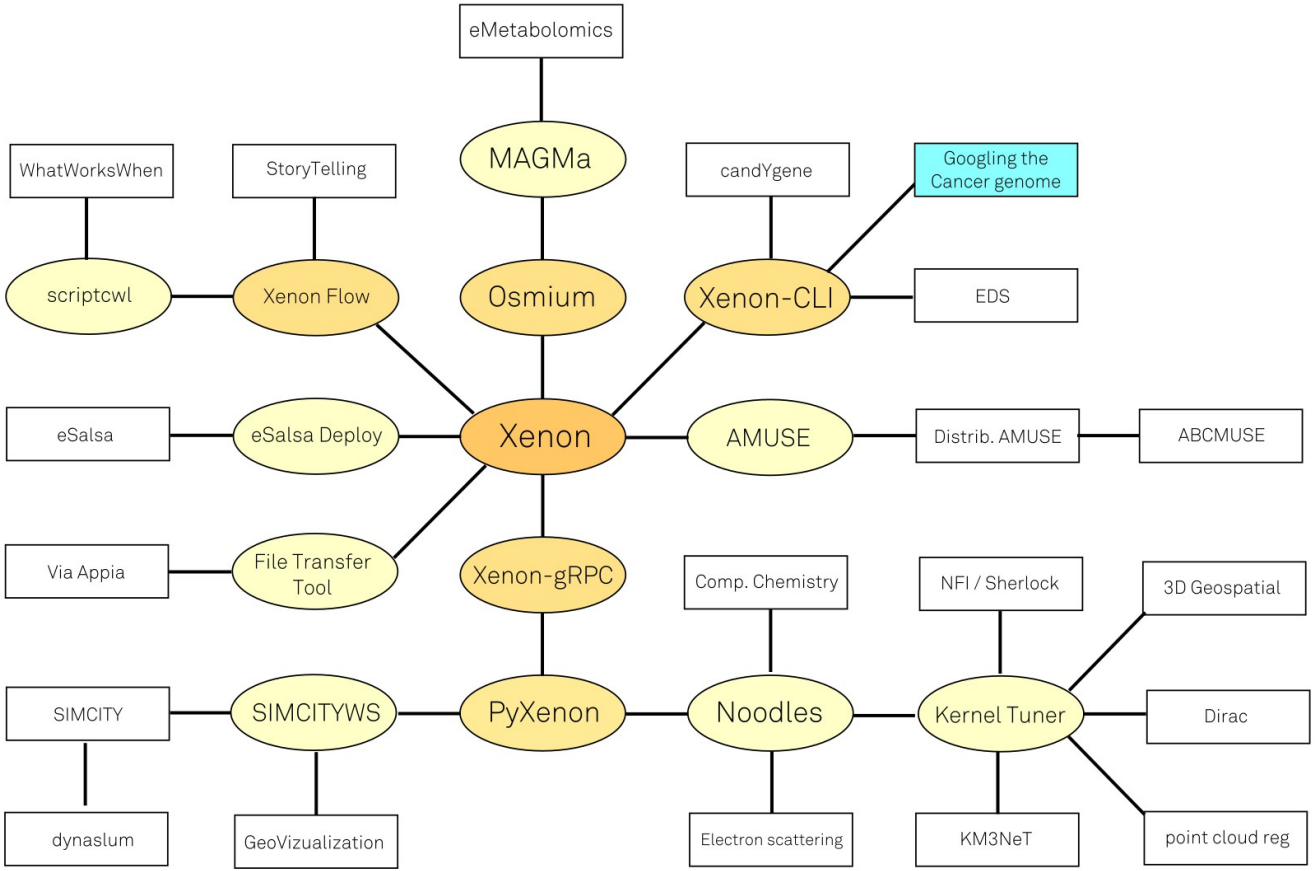

xenon

(and friends)

A collection of libraries and tools that provide easy access to (remote) compute and storage resources

Used via-via in many other tools and projects.

xenon

(and friends)

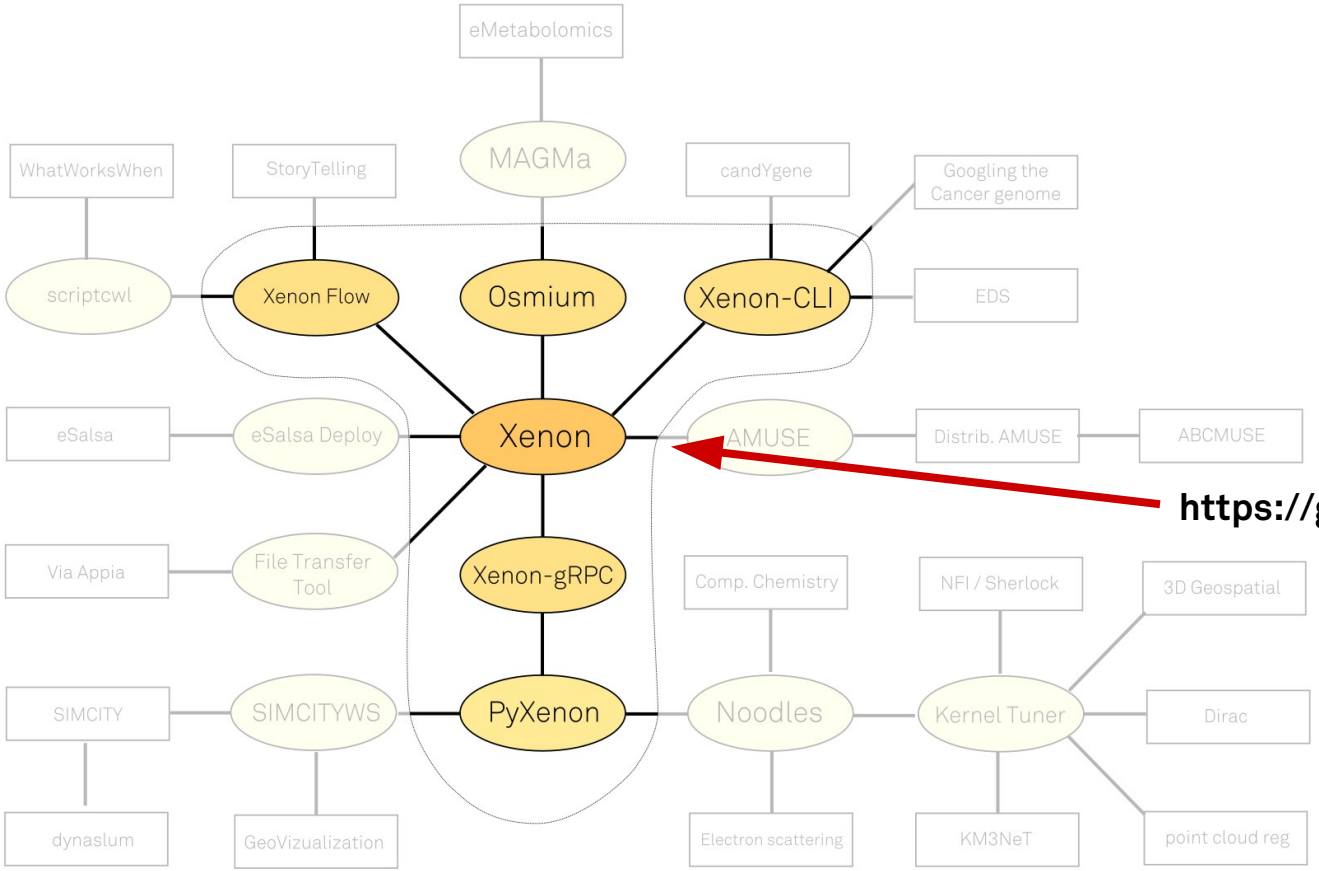

Core packages at:

https://github.com/xenon-middleware

How does xenon help ?

Xenon offers simple programming interfaces to access remote file systems and schedulers.

Adaptors implement the functionality for different backends.

A CLI allows for easy use in scripts.

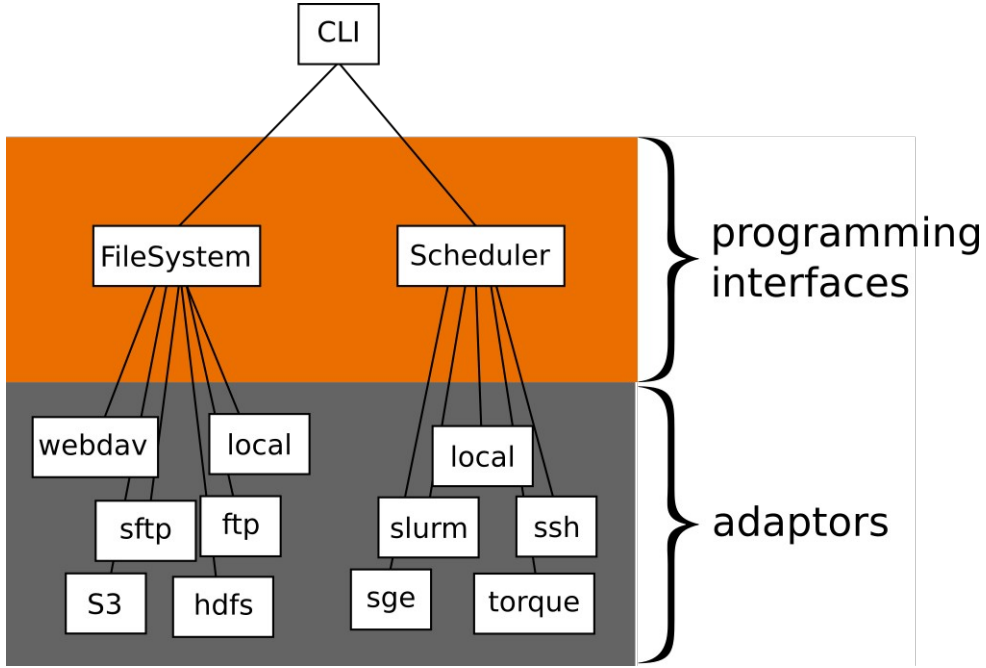

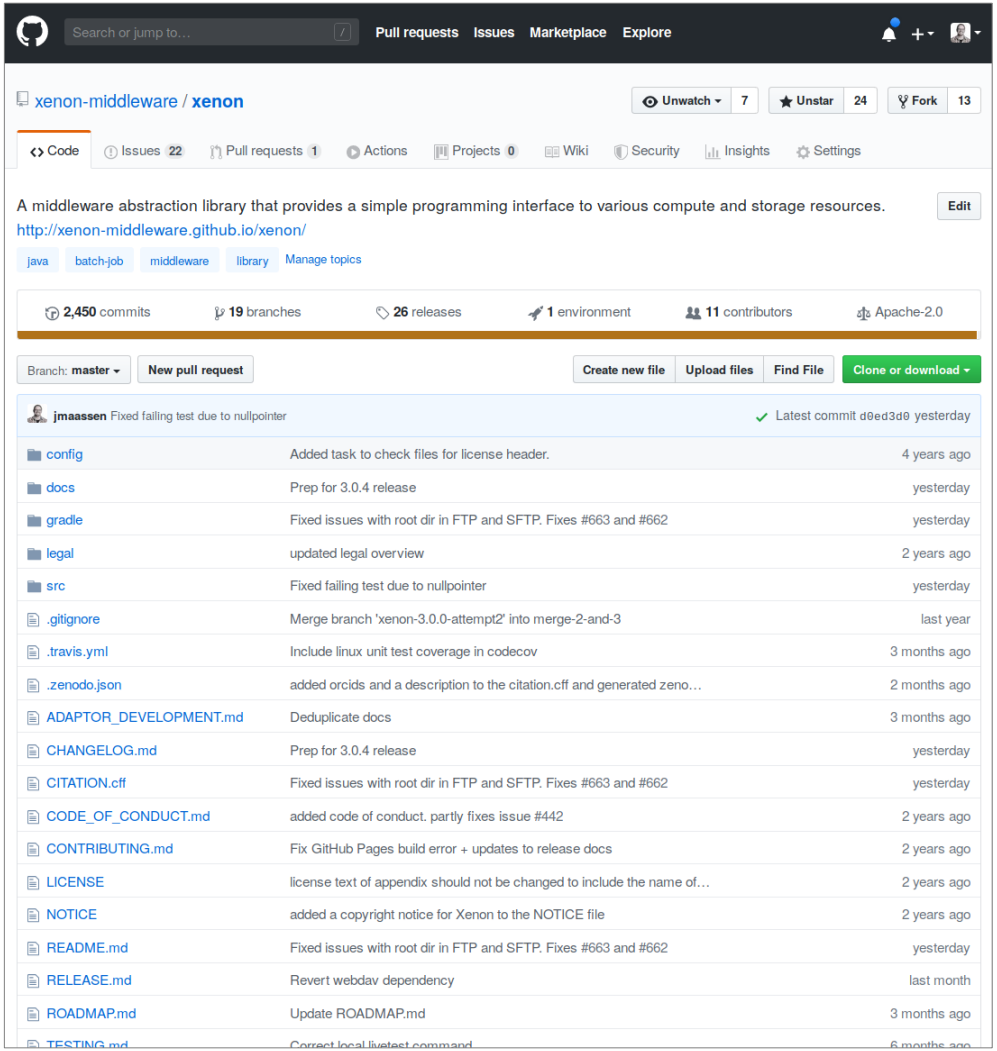

xenon

(the library)

Target audience is developers creating tools and (other) libraries.

Programming interface is kept simple, just focus on basic tasks:

- submit a job
- copy files

Good enough for 90% of the cases

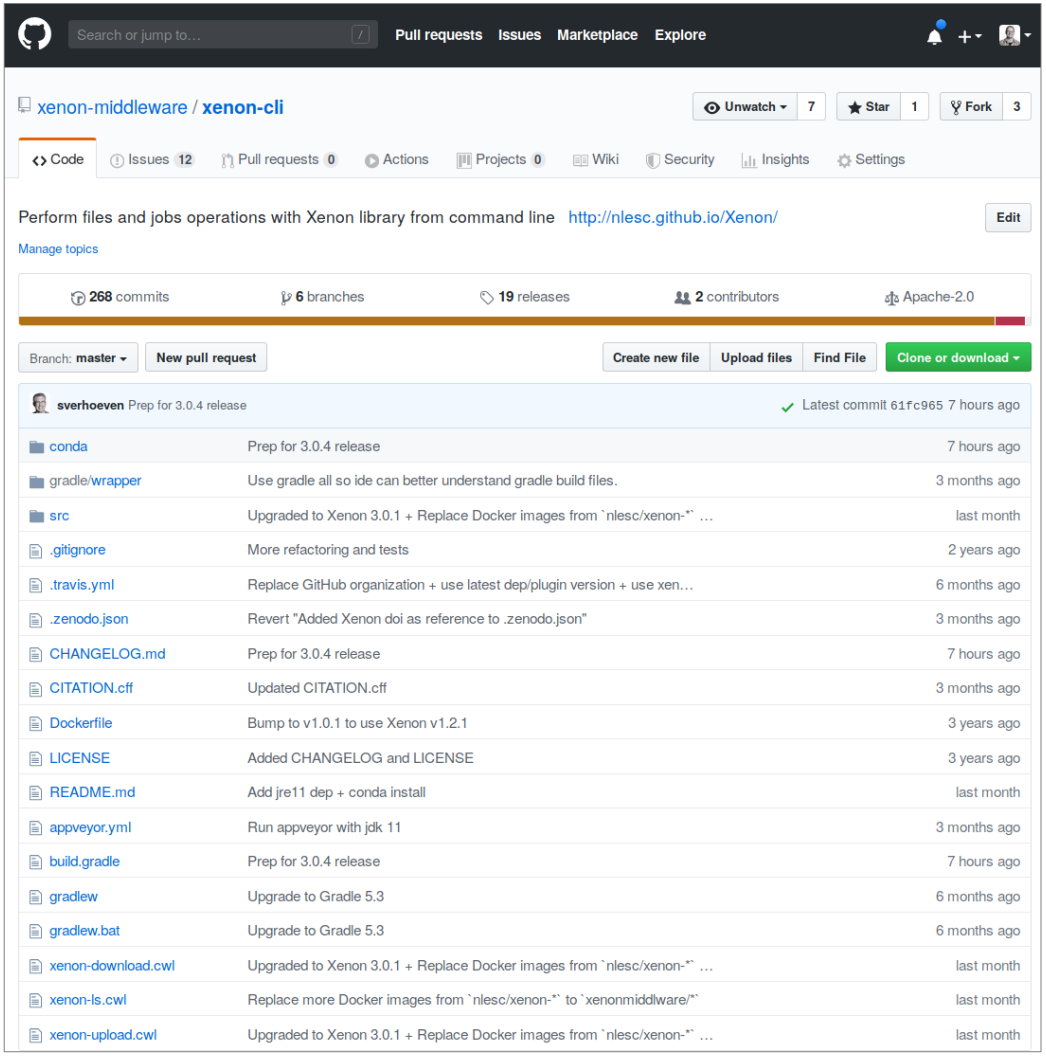

xenon-cli

(the command line tool)

Target audience is users creating scripts and workflows

Offers uniform syntax to use different platforms with focus on basic tasks:

- submit a computation
- copy files

Good enough for 90% of the cases

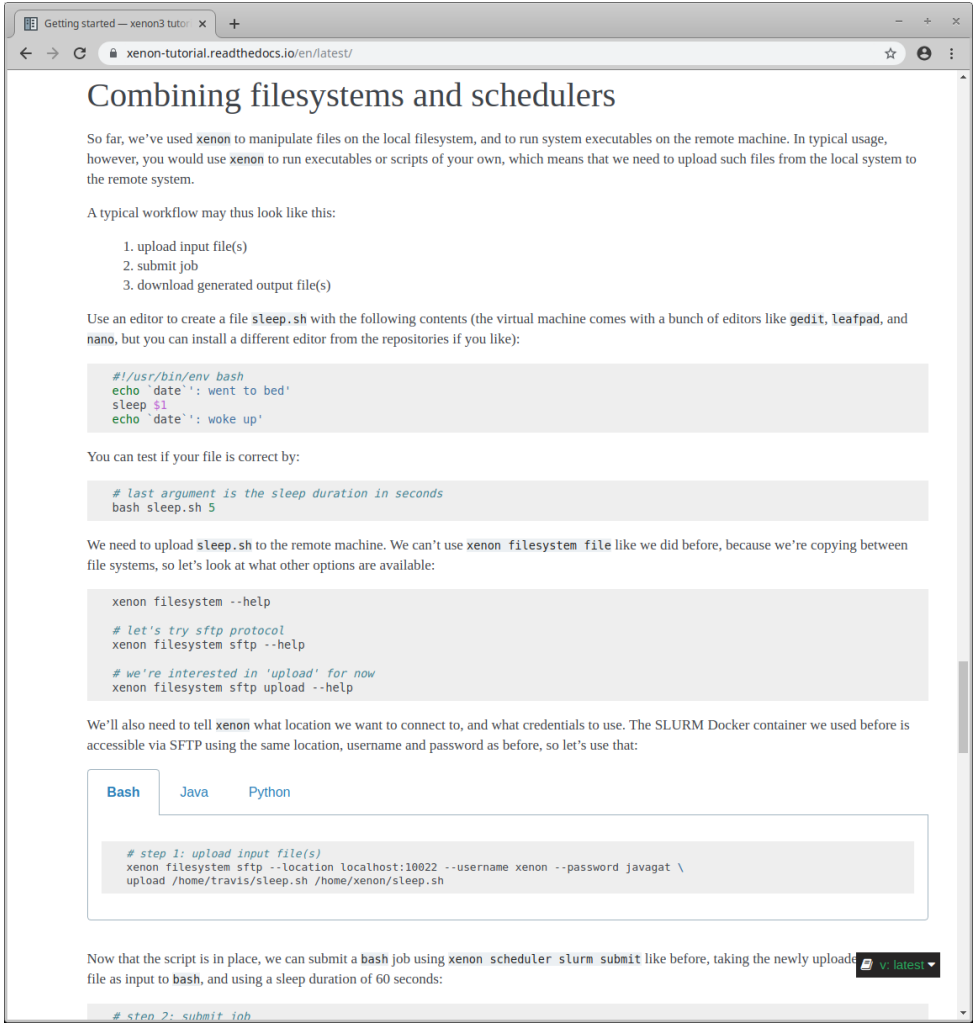

tutorial

https://xenon-tutorial.readthedocs.io

Explains how to use the
- Python API
- Java API
- Command Line Interface

Provide Linux VM image with
- pre-installed xenon
- docker containers (slurm, sftp)

Bash | Java | Python

```
xenon filesystem file list /home/travis/fixtures
```

Bash | **Java** | Python

```java
package nl.esciencecenter.xenon.tutorial;

import nl.esciencecenter.xenon.filesystems.FileSystem;
import nl.esciencecenter.xenon.filesystems.Path;
import nl.esciencecenter.xenon.filesystems.PathAttributes;


public class DirectoryListing {

    public static void main(String[] args) throws Exception {

        String adaptor = "file";
        FileSystem filesystem = FileSystem.create(adaptor);
        Path directory = new Path("/home/travis/fixtures");
        Boolean recursive = false;
        Iterable<PathAttributes> listing = filesystem.list(directory, recursive);

        for (PathAttributes elem : listing) {
            if (!elem.isHidden()) {
                System.out.println(elem.getPath());
            }
        }
    }
}
```

Bash | Java | **Python**

```python
import xenon
from xenon import Path, FileSystem


def run_example():

    xenon.init()

    filesystem = FileSystem.create(adaptor='file')
    path = Path("/home/travis/fixtures")

    listing = filesystem.list(path, recursive=False)

    for entry in listing:
        if not entry.path.is_hidden():
            print(entry.path)

    filesystem.close()


if __name__ == '__main__':
    run_example()
```

Portable HPC workflows based on

Snakemake + CONDA + XENON

Arnold Kuzniar

netherlands eScience center

Googling the cancer genome project

Identification and prioritization of cancer-causing structural variants (SVs)

UMCU team

Jeroen de Ridder (PI)

Wigard Kloosterman

Luca Santuari

Carl Shneider

eScience team

Lars Ridder (coordinator)

Arnold Kuzniar

Sonja Georgievska

Jason Maassen

Stefan Verhoeven

UMC Utrecht

netherlands eScience center

Challenge: reliable detection of SVs

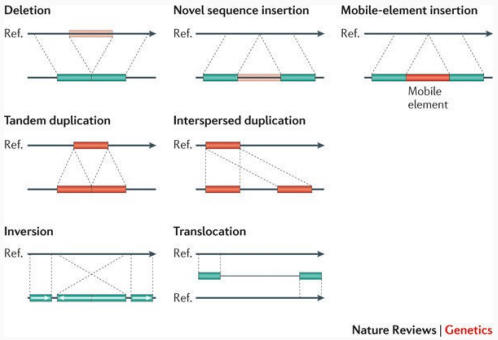

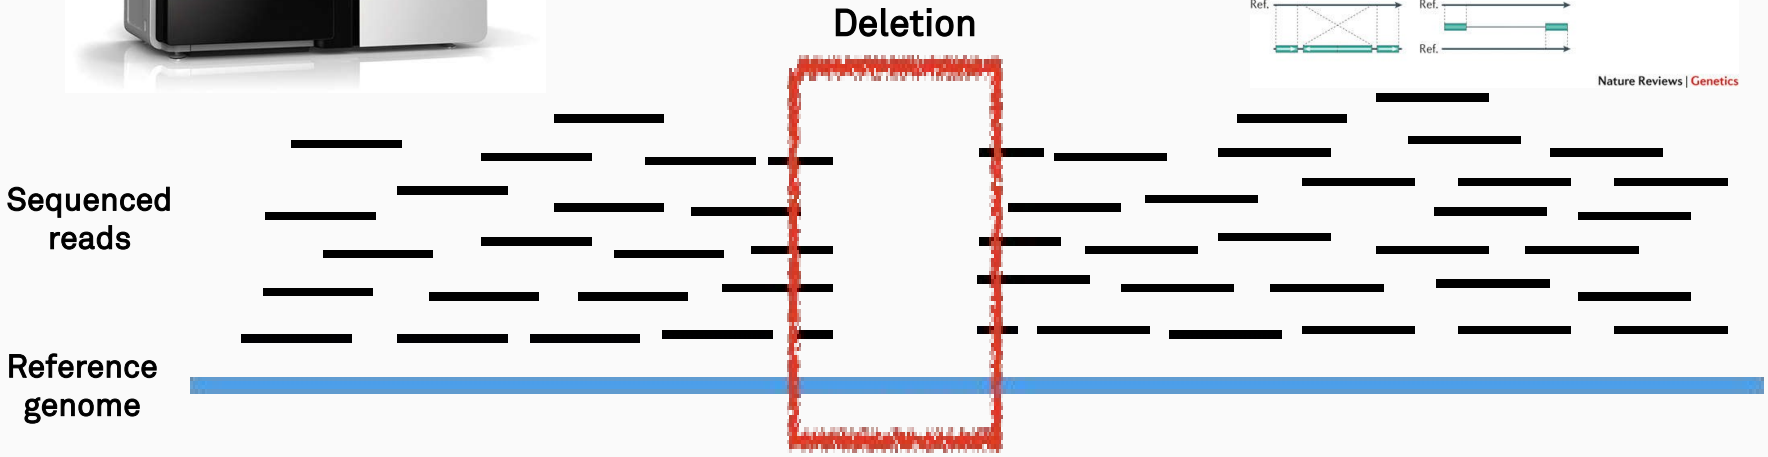

netherlands eScience center

sv-callers workflow

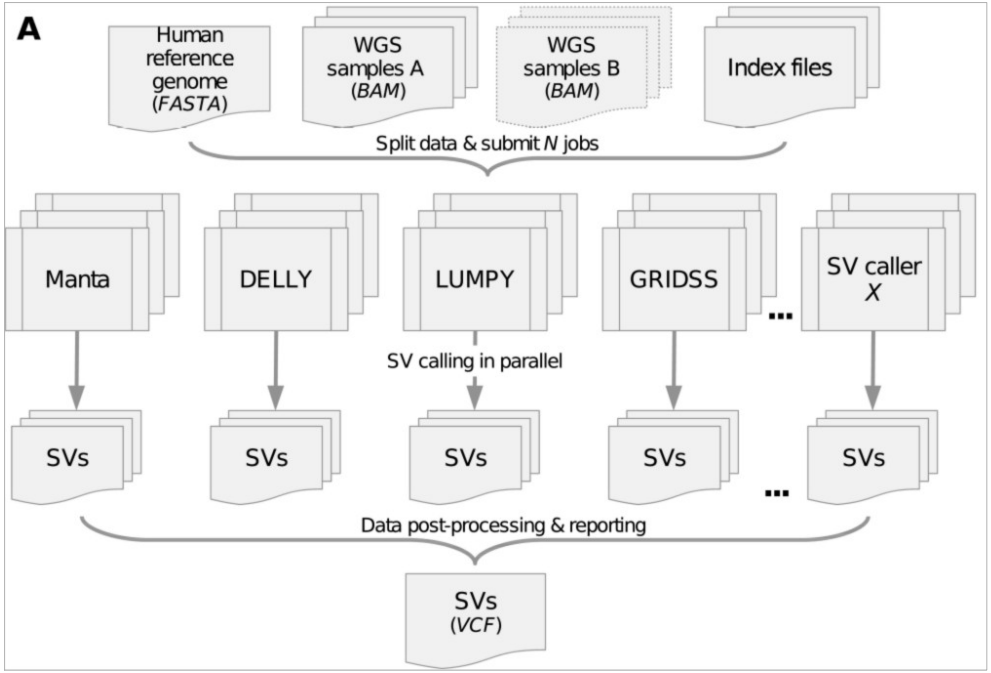

Software	Implementation	Parallelism
Manta	C++, Python	pyFlow[1] tasks, SIMD[2]
DELLY	C++	OpenMP[3] threads
LUMPY	C/C++, Python	not supported
GRIDSS	Java, R, Python	Java threads, SIMD[2]

Combines existing SV detection tools into **a portable HPC workflow**

sv-callers workflow

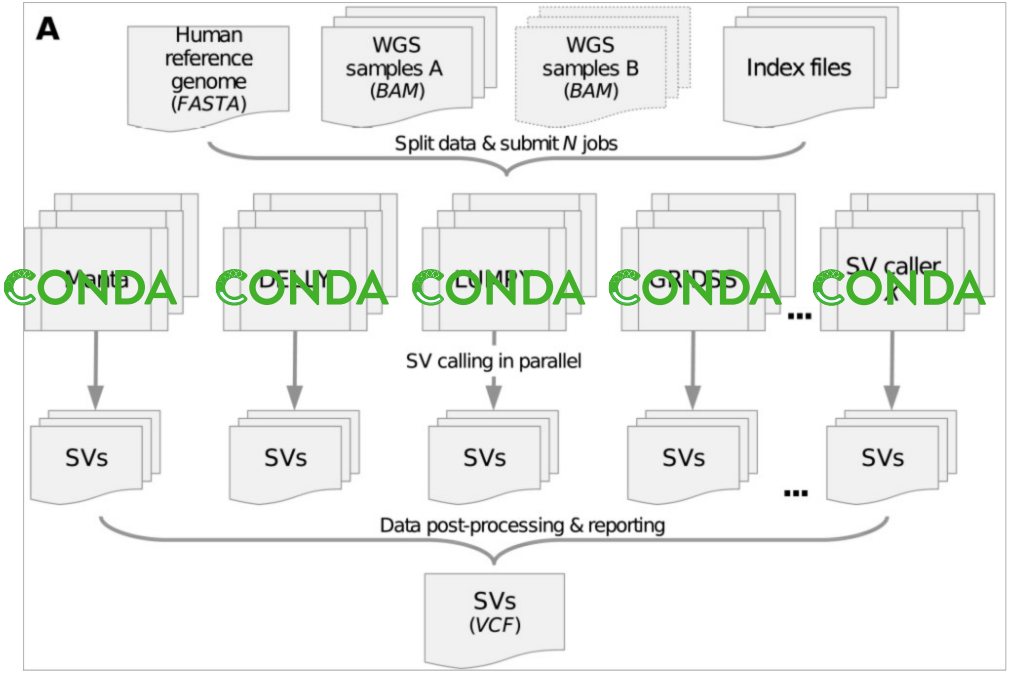

Conda is used to enable automatic install of the tools on the target system

sv-callers workflow

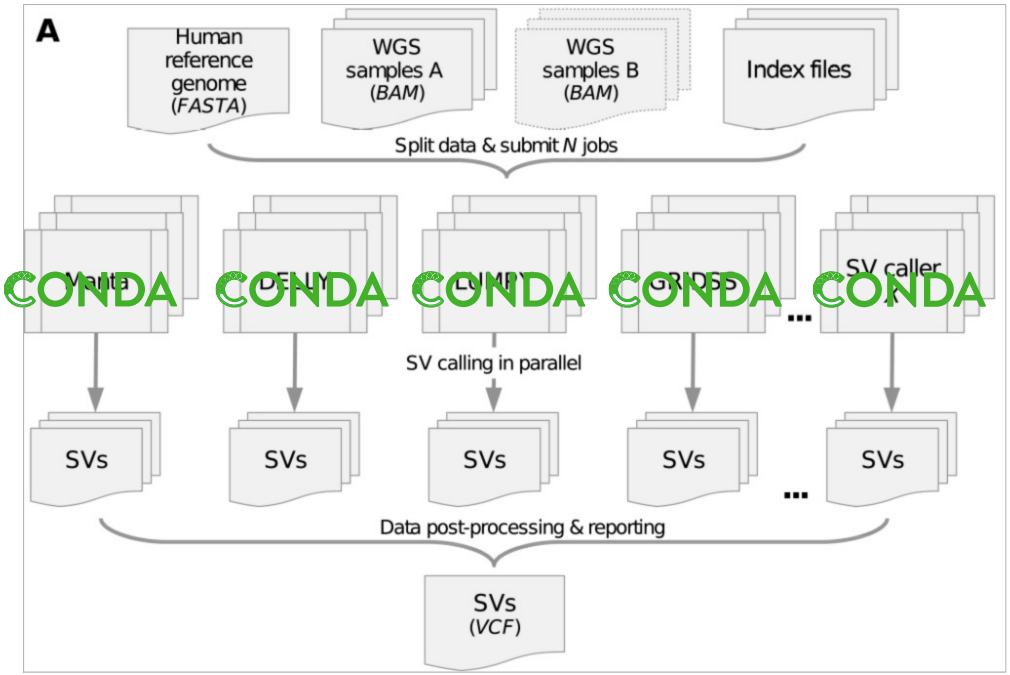

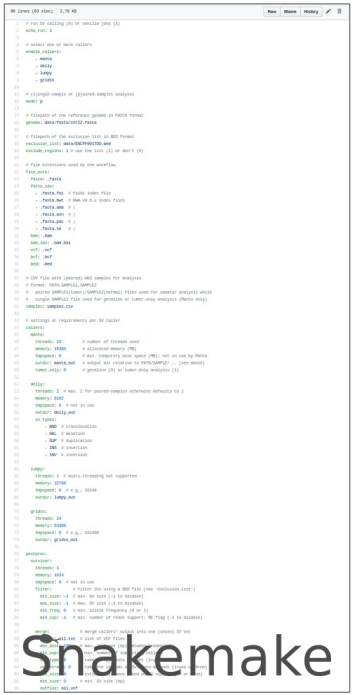

Snakemake is used to define the overall workflow

sv-callers workflow

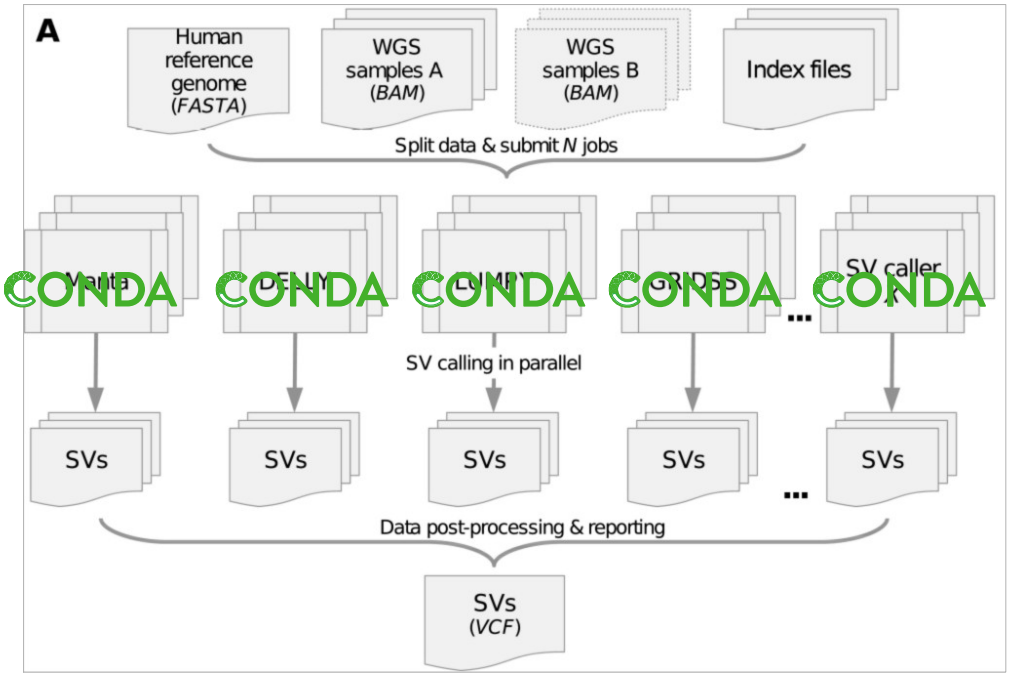

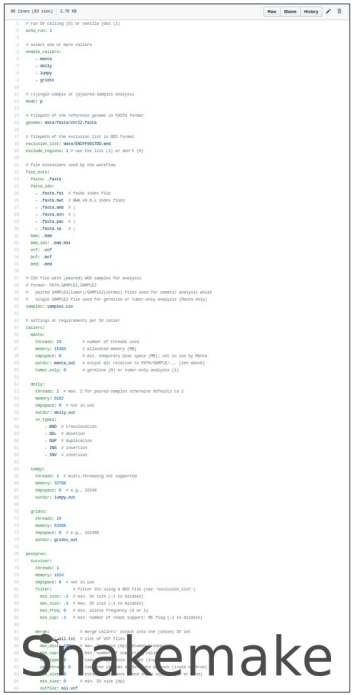

XENON

DAS 5

Utrecht Bioinformatics Center

SURF SARA

Snakemake then uses xenon to deploy on the available resources

Snakemake → Xenon → GridEngine

```
snakemake -C echo_run=1 mode=p
enable_callers="['manta','delly','lumpy','gridss']"
--use-conda --latency-wait 30 --jobs 9
--cluster 'xenon scheduler gridengine
--location local:// submit --name smk.{rule} --inherit-env
--procs-per-node {threads} --start-single-process
--max-run-time 1 --max-memory {resources.mem_mb}
--working-directory . --stderr stderr-%j.log
--stdout stdout-%j.log' &>smk.log&
```

eScience center

Snakemake → Xenon → Slurm

```
snakemake -C echo_run=1 mode=p
enable_callers="['manta','delly','lumpy','gridss']"
--use-conda --latency-wait 30 --jobs 9
--cluster 'xenon scheduler slurm
--location local:// submit --name smk.{rule} --inherit-env
--procs-per-node {threads} --start-single-process
--max-run-time 1 --max-memory {resources.mem_mb}
--working-directory . --stderr stderr-%j.log
--stdout stdout-%j.log' &>smk.log&
```

eScience center

Results

(A) Germline and (B) somatic SVs detected in the benchmark and in the cell lines samples, respectively, using Manta, DELLY, LUMPY and GRIDSS. Most SVs are caller-specific, followed by SVs common to three of the four callers. SVs detected by the callers were filtered and merged into one set (see the Methods section). Note: The Venn diagrams include the largest GRIDSS sets as the GRIDSS output varies slightly each run using the same input. Fig. S1-2 show the comparisons across sample copies.

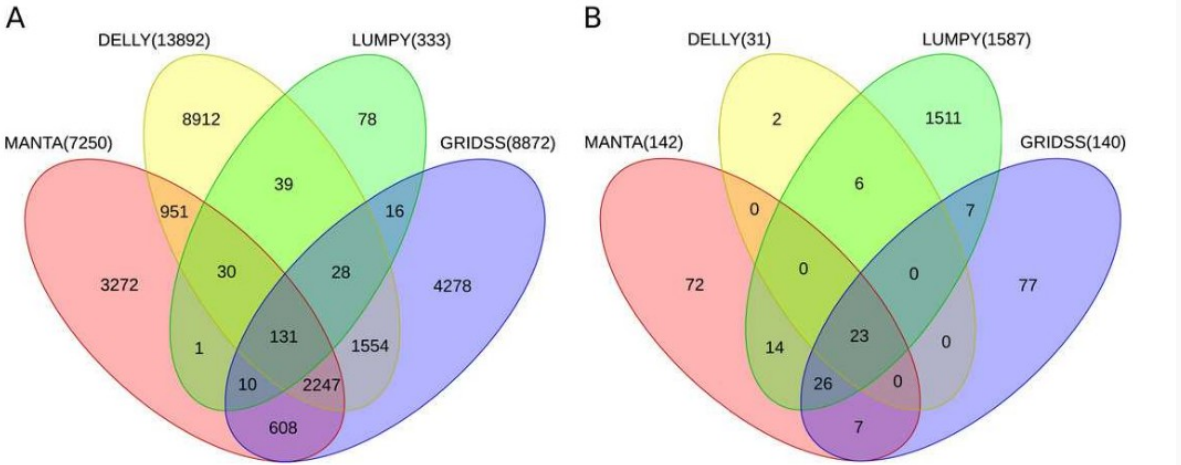

https://github.com/GooglingTheCancerGenome/sv-callers

“sv-callers:
a highly portable parallel workflow for structural variant detection in whole-genome sequence data”

Arnold Kuzniar, Jason Maassen, Stefan Verhoeven, Luca Santuari, Carl Shneider, Wigard P. Kloosterman, Jeroen de Ridder

Accepted for publication in PeerJ

eScience center

Xenon Roadmap

almost done: GridFTP, amazon-batch

in progress: azure-batch

ideas: iRODS (?), google storage (?), ...
language bindings for C++, C#, Go, etc.
better integration in snakemake,

suggestions are welcome!

Thanks!

https://github.com/xenon-middleware
https://xenon-tutorial.readthedocs.io
https://github.com/GooglingTheCancerGenome/sv-callers

xenon@esciencecenter.nl
j.maassen@esciencecenter.nl

www.esciencecenter.nl

Interested in Research Software ?

The Netherlands eScience Center is the Dutch national center of excellence for the development and application of research software to advance academic research.

Join the team !

n.renaud@esciencecenter.nl

netherlands eScience center

+31 (0)20 460 4770

www.esciencecenter.nl

blog.esciencecenter.nl

Extra slides

SIM-CITY

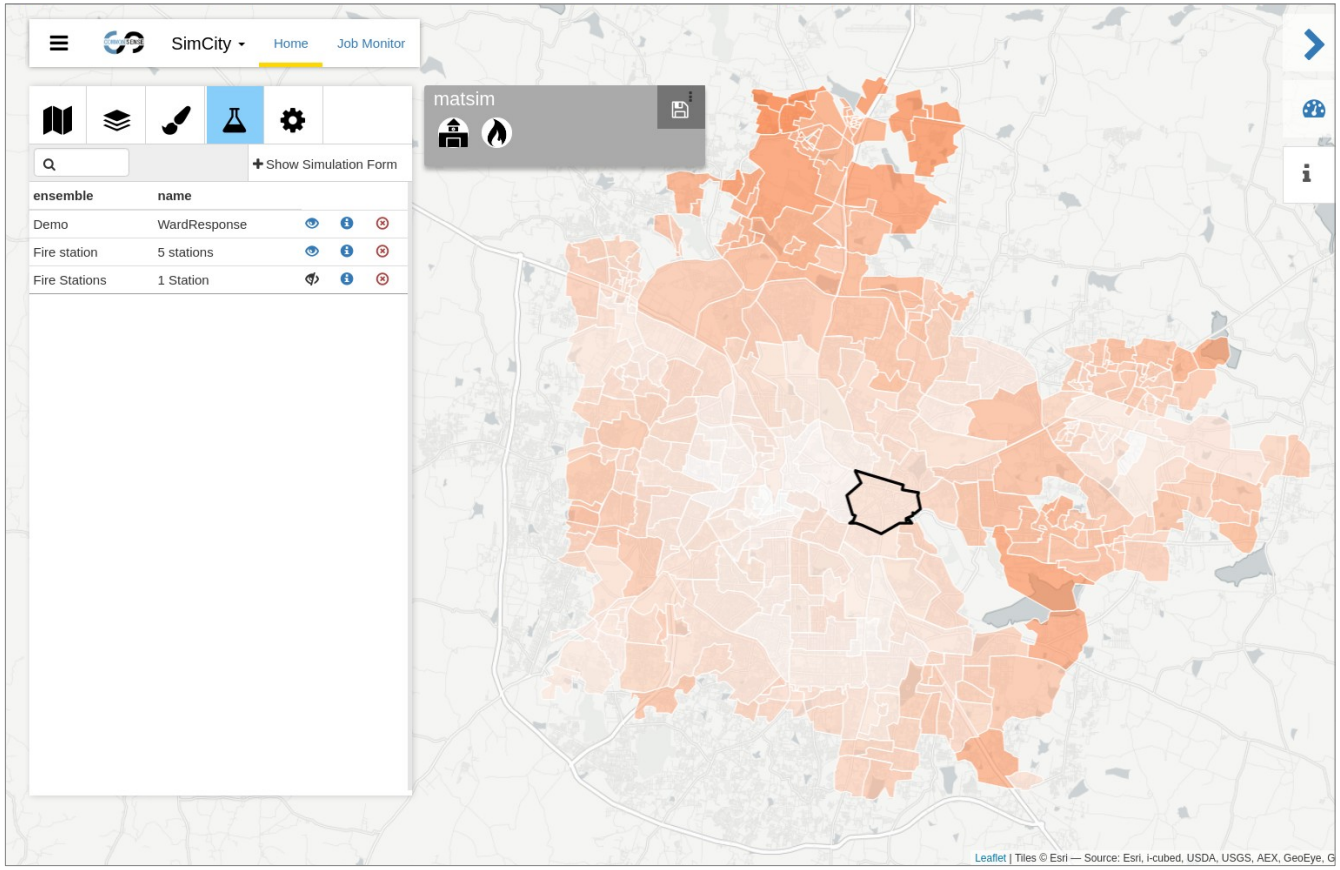

Goal: decision support for urban social economic complexity

Run urban planning simulations from an interactive user interface.

Example: location of fire stations

Simulations must be started on a remote compute cluster.

eMetabolomics

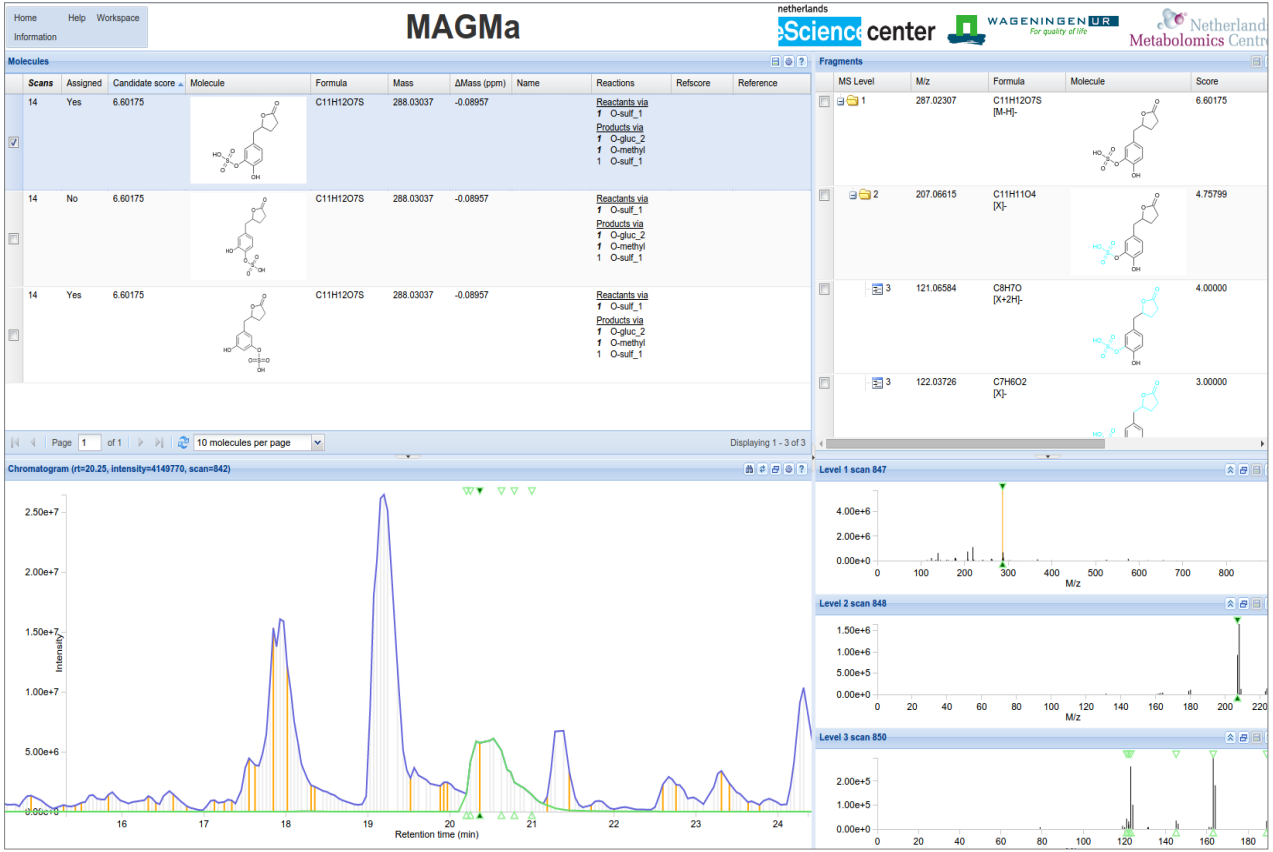

Goal: computer assisted metabolite profiling.

Determine how compounds are metabolized in the human body based on mass spectrometer data and metabolism rules.

Requires starting analysis computations on a server containing the metabolism database.

AMUSE / OMUSE

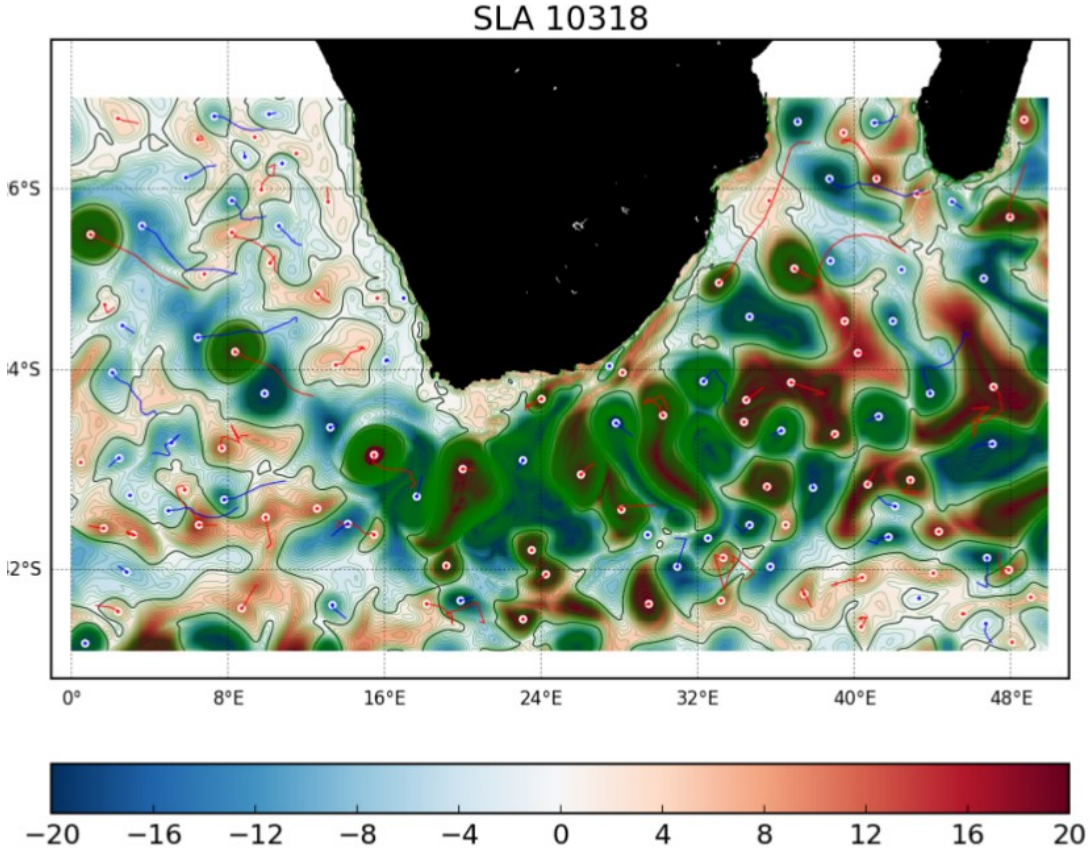

Goal: combine existing physics simulations into more complex models a simple python framework.

Originally for astrophysics (AMUSE), but now also for ocean modelling (OMUSE).

Requires starting simulations on multiple compute clusters and controlling them remotely.