Daisy Workshop

copenhagen november 5th 2021
(slideset at https://bit.ly/daisywrapper)

A (Java) Daisy execution wrapper

A wrapper to run Daisy in parallel on remote HPC hardware or Google Cloud.
No hardware required, easy setup and usually low to no cost.

Open Source - https://github.com/nitrogensensor/daisywrapper

Maintainer/supporter: jacob.nordfalk@gmail.com

Climate-KIC is supported by the EIT, a body of the European Union

Agenda

- Wrapping Daisy executions to run remotely on a HPC server
  - Short demo
- How it works
- Google cloud deployment - automated execution and scaling
  - Calculation times and cost considerations
- If time permits
  - Live demo of running 1400 simulations in Google Cloud
  - Live demo of running your own server
  - Parameters to make the wrapper work for you

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper

Slides https://bit.ly/daisywrapper

Why use a wrapper to execute Daisy remotely?

- No need to setup Daisy
- No need for a powerful PC
- Parallel execution - using all available CPU cores
- An execution cache permits re-use of old execution results
- Easy setup and use from R, Matlab, Python etc

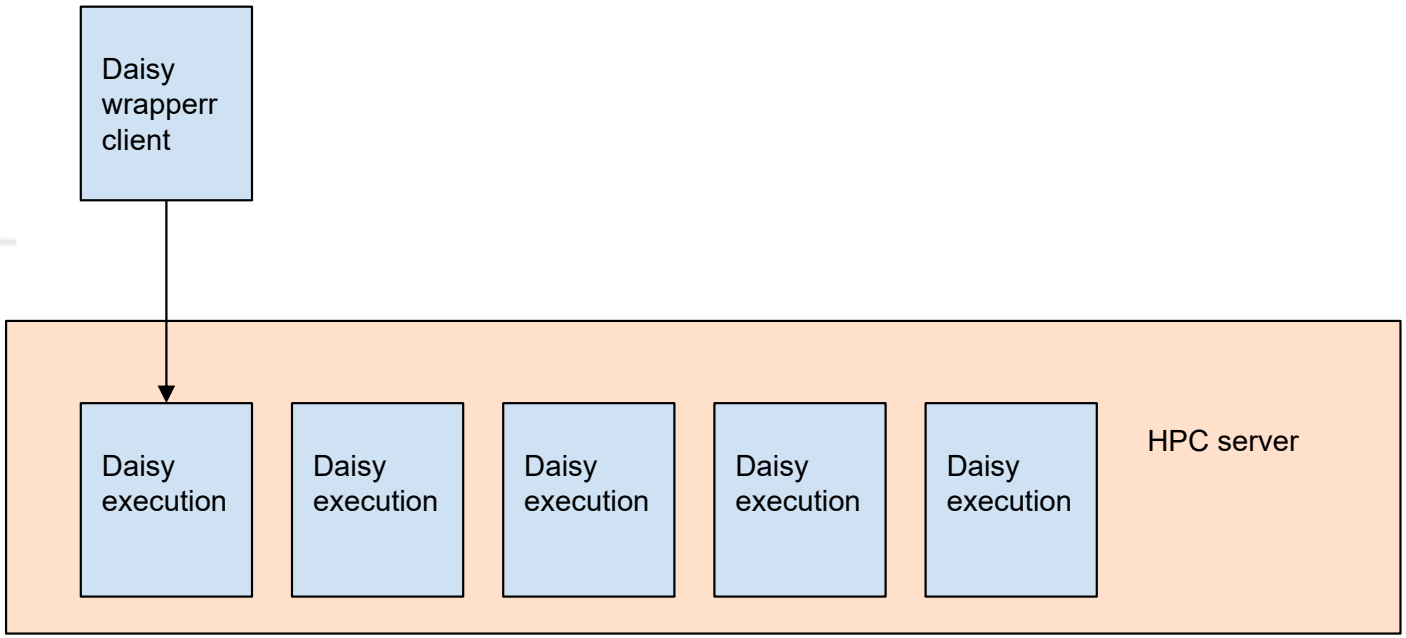

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper

Slides https://bit.ly/daisywrapper

How to use the wrapper

You do not need to install Daisy. Instead of

```
daisy.exe -d daisyInput Sim1.dai
daisy.exe -d daisyInput Sim2.dai
daisy.exe -d daisyInput Sim3.dai
```

… you install Java and get daisy.jar executable and run:

```
java -jar daisy.jar client -d daisyInput Sim1.dai Sim2.dai Sim3.dai -o daisyOutput
```

The `daisyInput` directory is uploaded to server (the default is at DTU)
The 3 simulations are done in parallel in 3 separate directories

The *complete* content of the 3 directories are downloaded and stored locally
(in `daisyOutput/Sim1` , `daisyOutput/Sim2` , and `daisyOutput/Sim3`)
Which files to download can be specified, like this: `-of daisy.log -of Output/lai_sim.csv`

Demo R notebooks:
https://github.com/nitrogensensor/daisywrapper/tree/master/src/test/resources/TestData https://daisy.nitrogensensor.eu/resultat/R/daisyResultVisualization/jordlag2.html

Climate-KIC is supported by the EIT, a body of the European Union

Scaling parallel Daisy executions

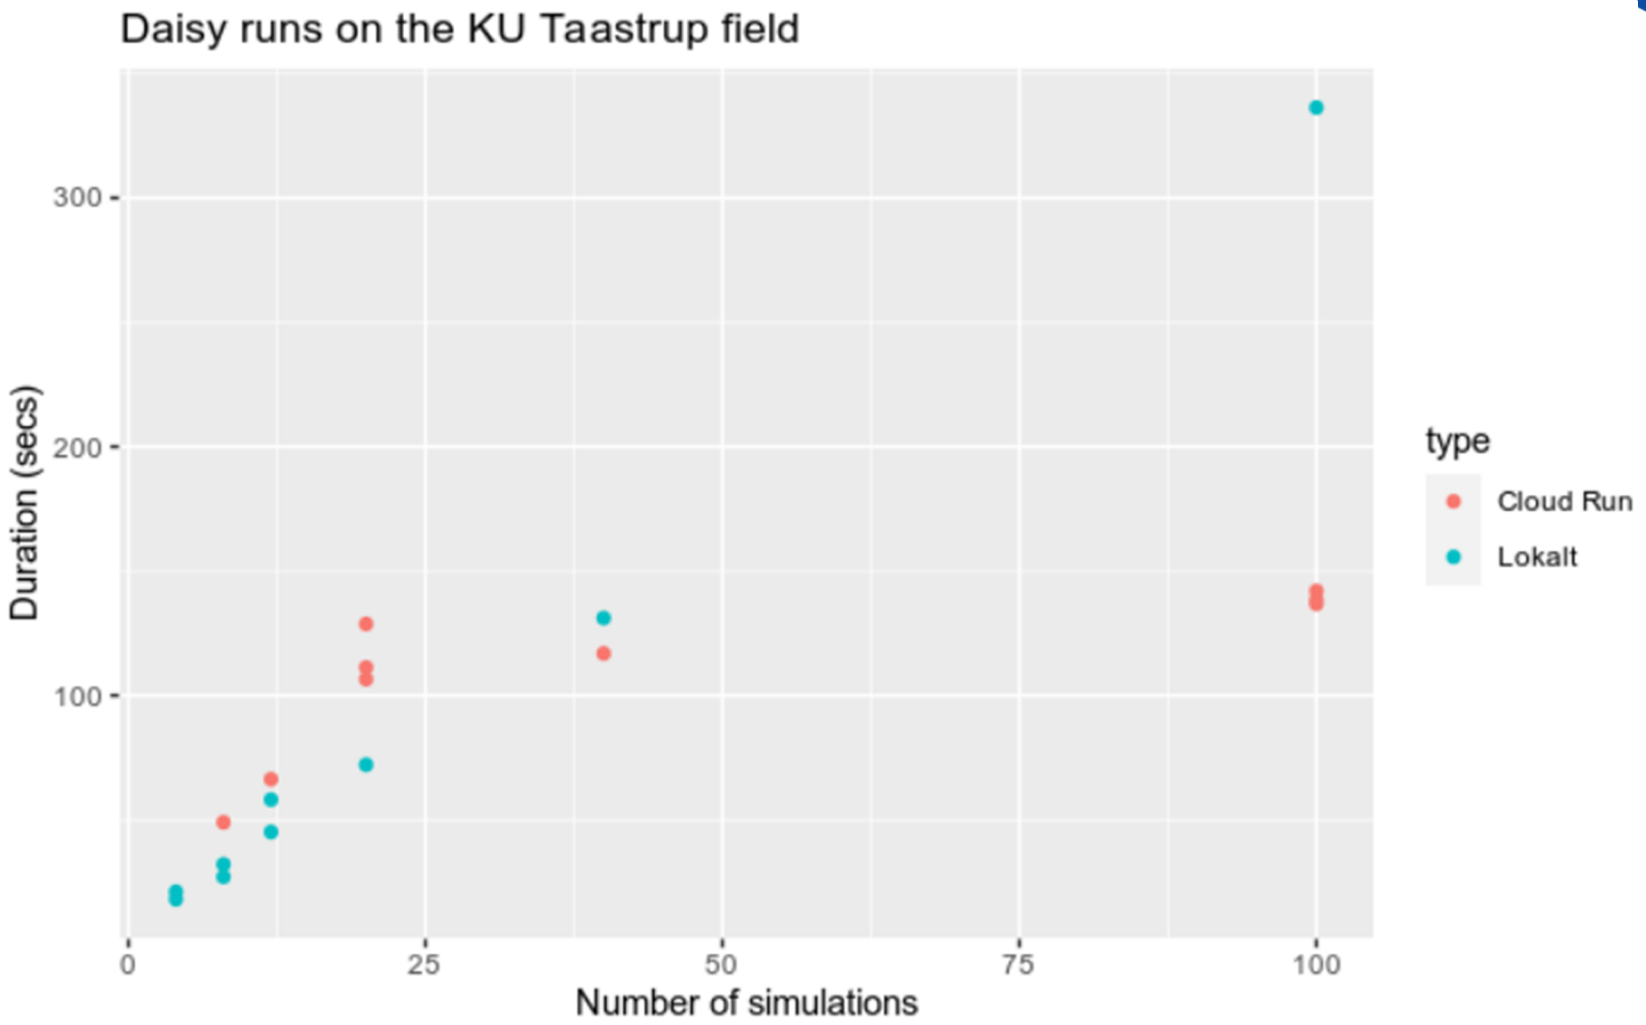

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper

Slides https://bit.ly/daisywrapper

Scaling out on Google Cloud Run

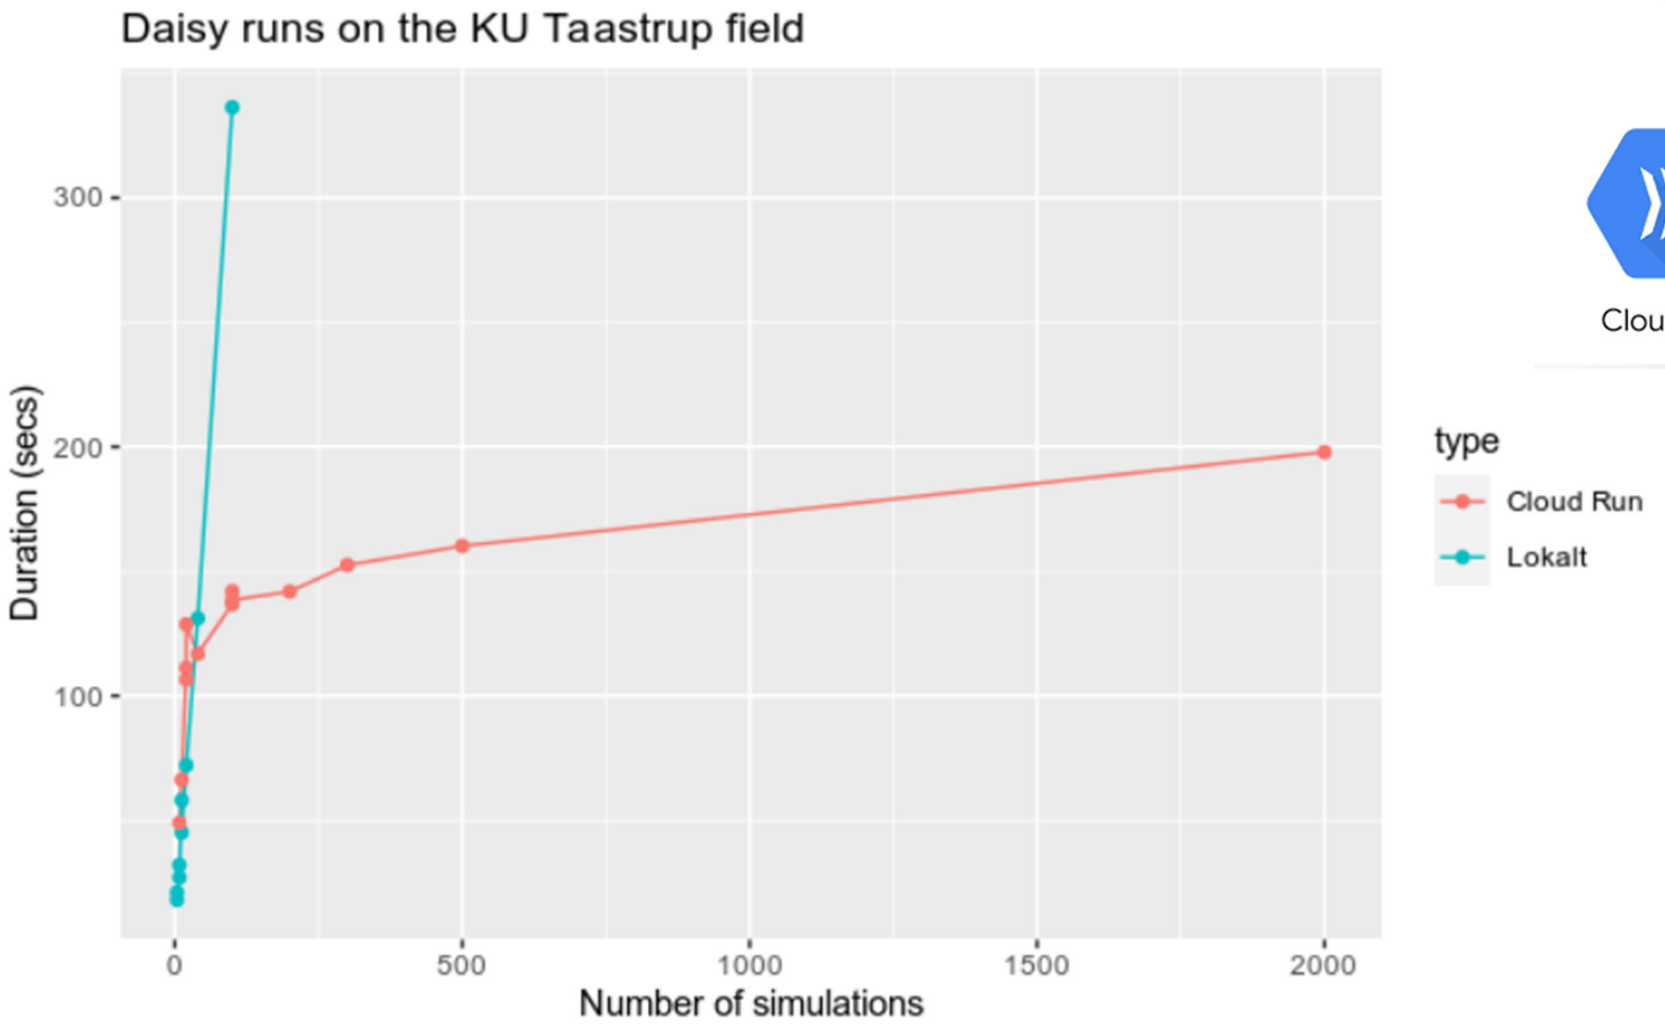

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper

Slides https://bit.ly/daisywrapper

Using Google Cloud Run

```
java -jar daisy.jar client -u https://daisykoersel-6dl4uoo23q-lz.a.run.app ...
```

- Massive parallel calculations are runned using Google Cloud Run
  - Allows serverless/scale-to-zero
  - You only pay for the CPU used during computation
- No execution cache
  - (previous local results could be reused, but this is currently turned off)

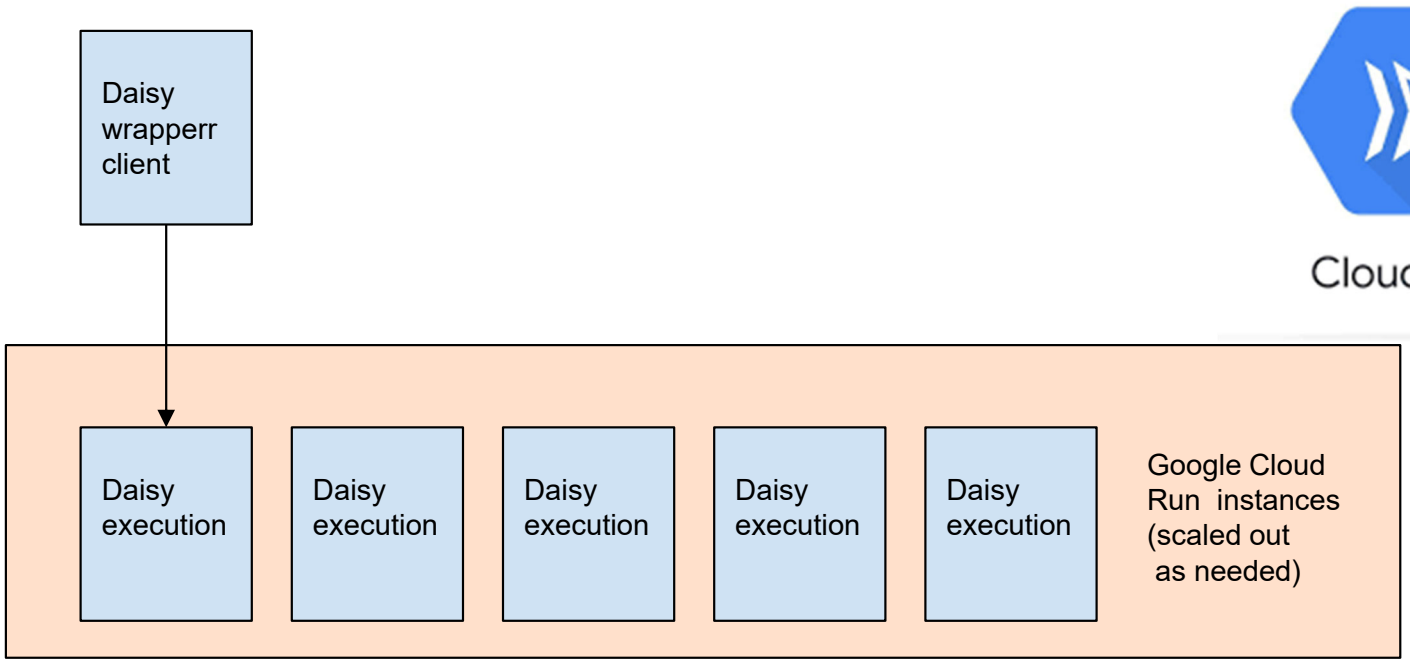

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper
Slides https://bit.ly/daisywrapper

Cloud Run Cost considerations

Each simulation minute costs approx 0.001 EUR

(1 EUR for 2000 simulations each taking 30 seconds).

Cloud Run scales down to zero when its not used. No cost when nothing is running

October 6 – November 4, 2021 (total cost)

DKK11.93

includes -DKK30.82 in credits

↑ 2,958.97%

DKK11.54 over September 6 – October 5, 2021

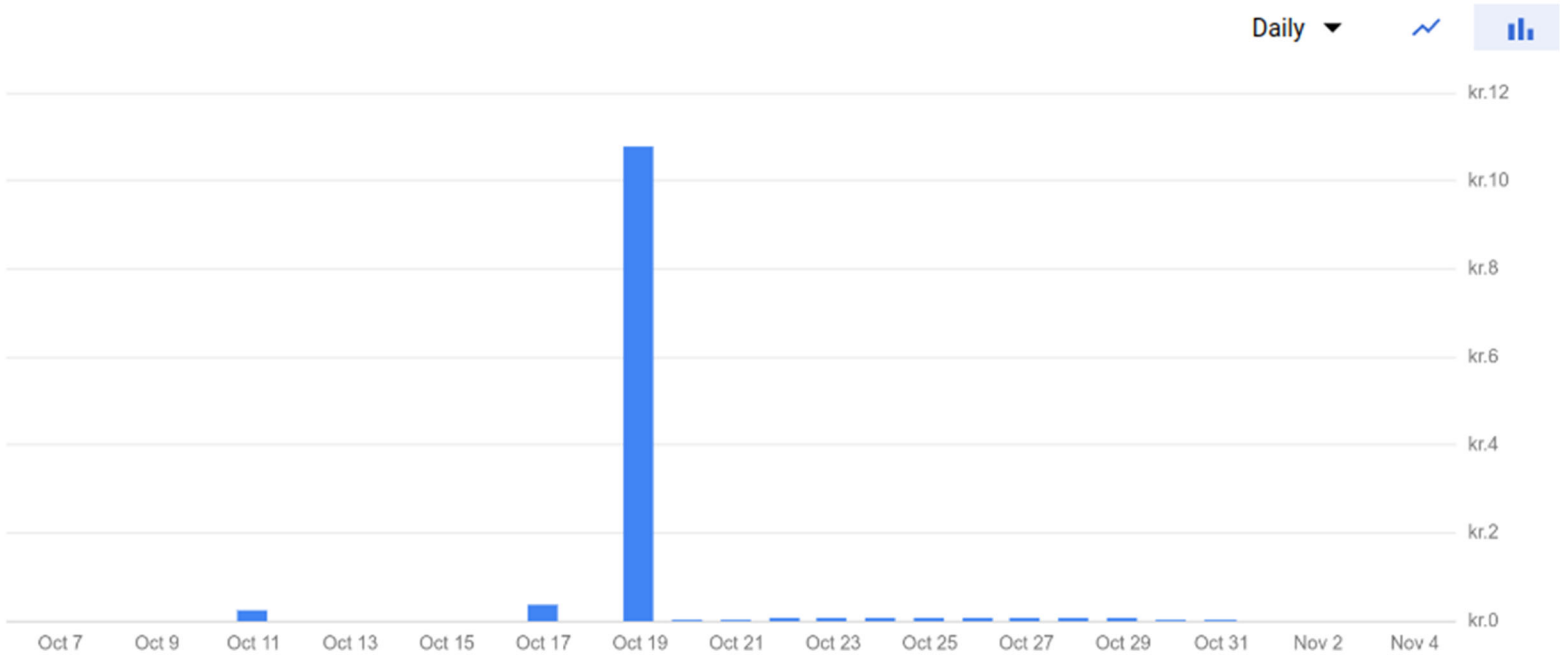

Service	Cost	Discounts	Promotions and others	↓ Subtotal
Cloud Run	kr.42.75	-kr.30.82	kr.0.00	kr.11.93

Climate-KIC

Climate-KIC is supported by the EIT, a body of the European Union

On premise use - setting up a Daisy execution server

```
java -jar daisy.jar server
```

- Starts a server at endport http:// (ip address) :3210/
  - Test with `java -jar daisy.jar client -u http://localhost:3210/ ...`
- All available cores will be utilized
- An execution cache will be set up
  - If *exactly* the same files are sent from the client a cached version will be immediately returned
  - This allows you (and others) to reproduce results without waiting for Daisy re-executions

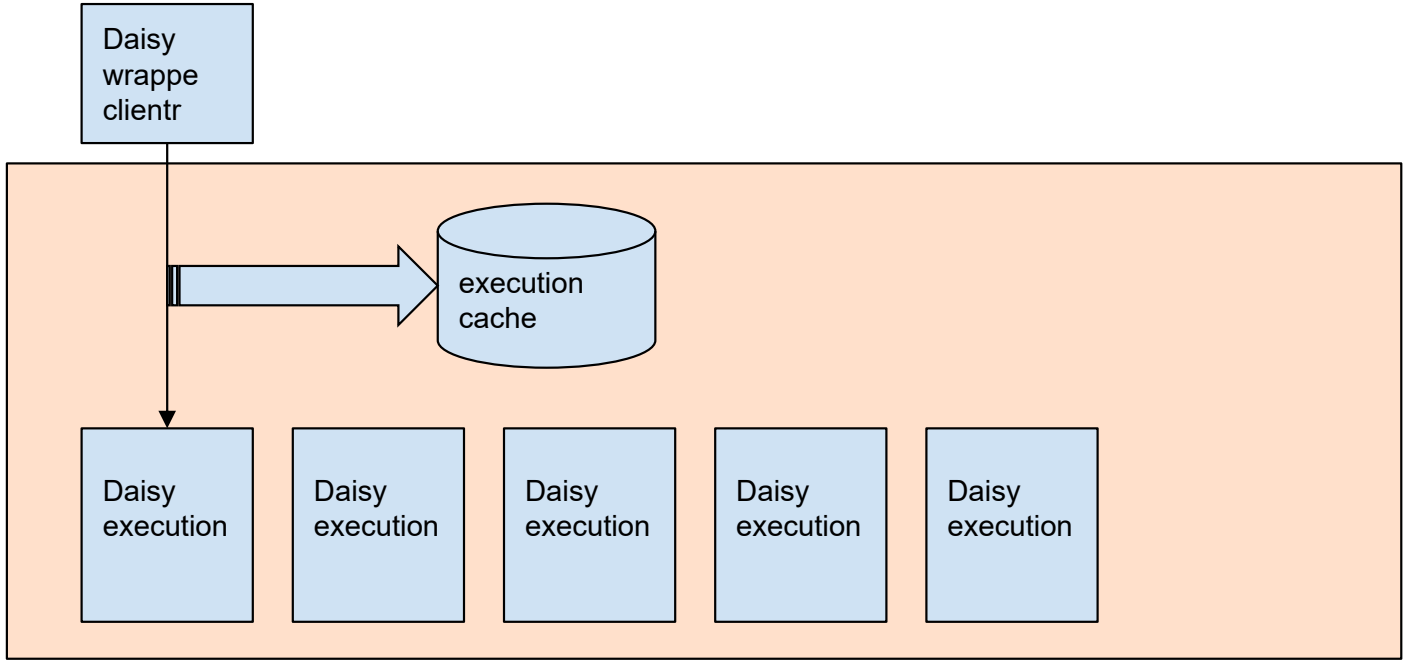

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper
Slides https://bit.ly/daisywrapper

What does the users say?

We have been using the wrapper intensively in NitrogenSensor project

Simulation for each 10x10 meters

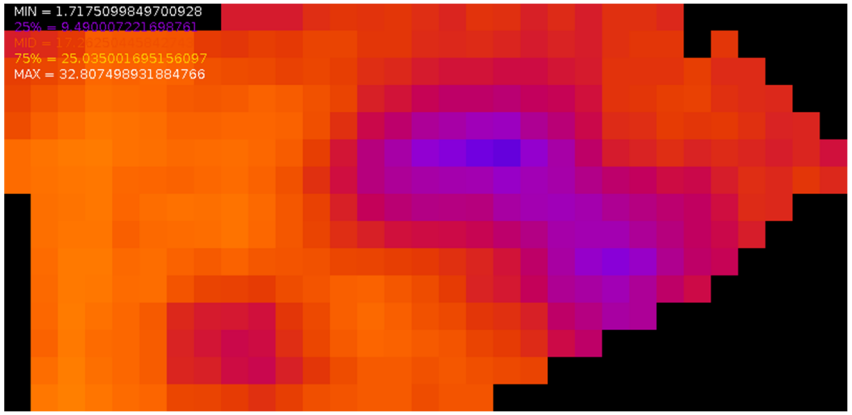

“I have been using the Daisy Wrapper a lot in the calibration work for spring barley in climate kick. Super nice work. It helped a lot.“

Simon Fiil Svane <sfs@plen.ku

Try it out! Contact Jacob Nordfalk for support (jacob.nordfalk@gmail.com)

DTU Diplom
Center for Diplomingeniøruddannelse

dinero.

https://github.com/nitrogensensor/daisywrapper

Slides https://bit.ly/daisywrapper

eit Climate-KIC DTU

Climate-KIC is supported by the EIT, a body of the European Union

Making the wrapper work for you

There's a replacement system

```
java -jar daisy.jar client -r '(stop *),(stop 2008 8 20)' -d TestData
Exercise01.dai
```

The replacement system allows you to re-use the *same* dai file for multiple runs.
E.g.
-r _sand_,37.1 replaces '_sand_' in the .dai file with '37.1'

-r _sand_:_humus_,10:90,20:80,30:70,40:60,50:50
gives 5 runs where sand rises from 10 to 50 and humus falls from 90 to 50 in steps of 10

Climate-KIC is supported by the EIT, a body of the European Union

https://github.com/nitrogensensor/daisywrapper
Slides https://bit.ly/daisywrapper