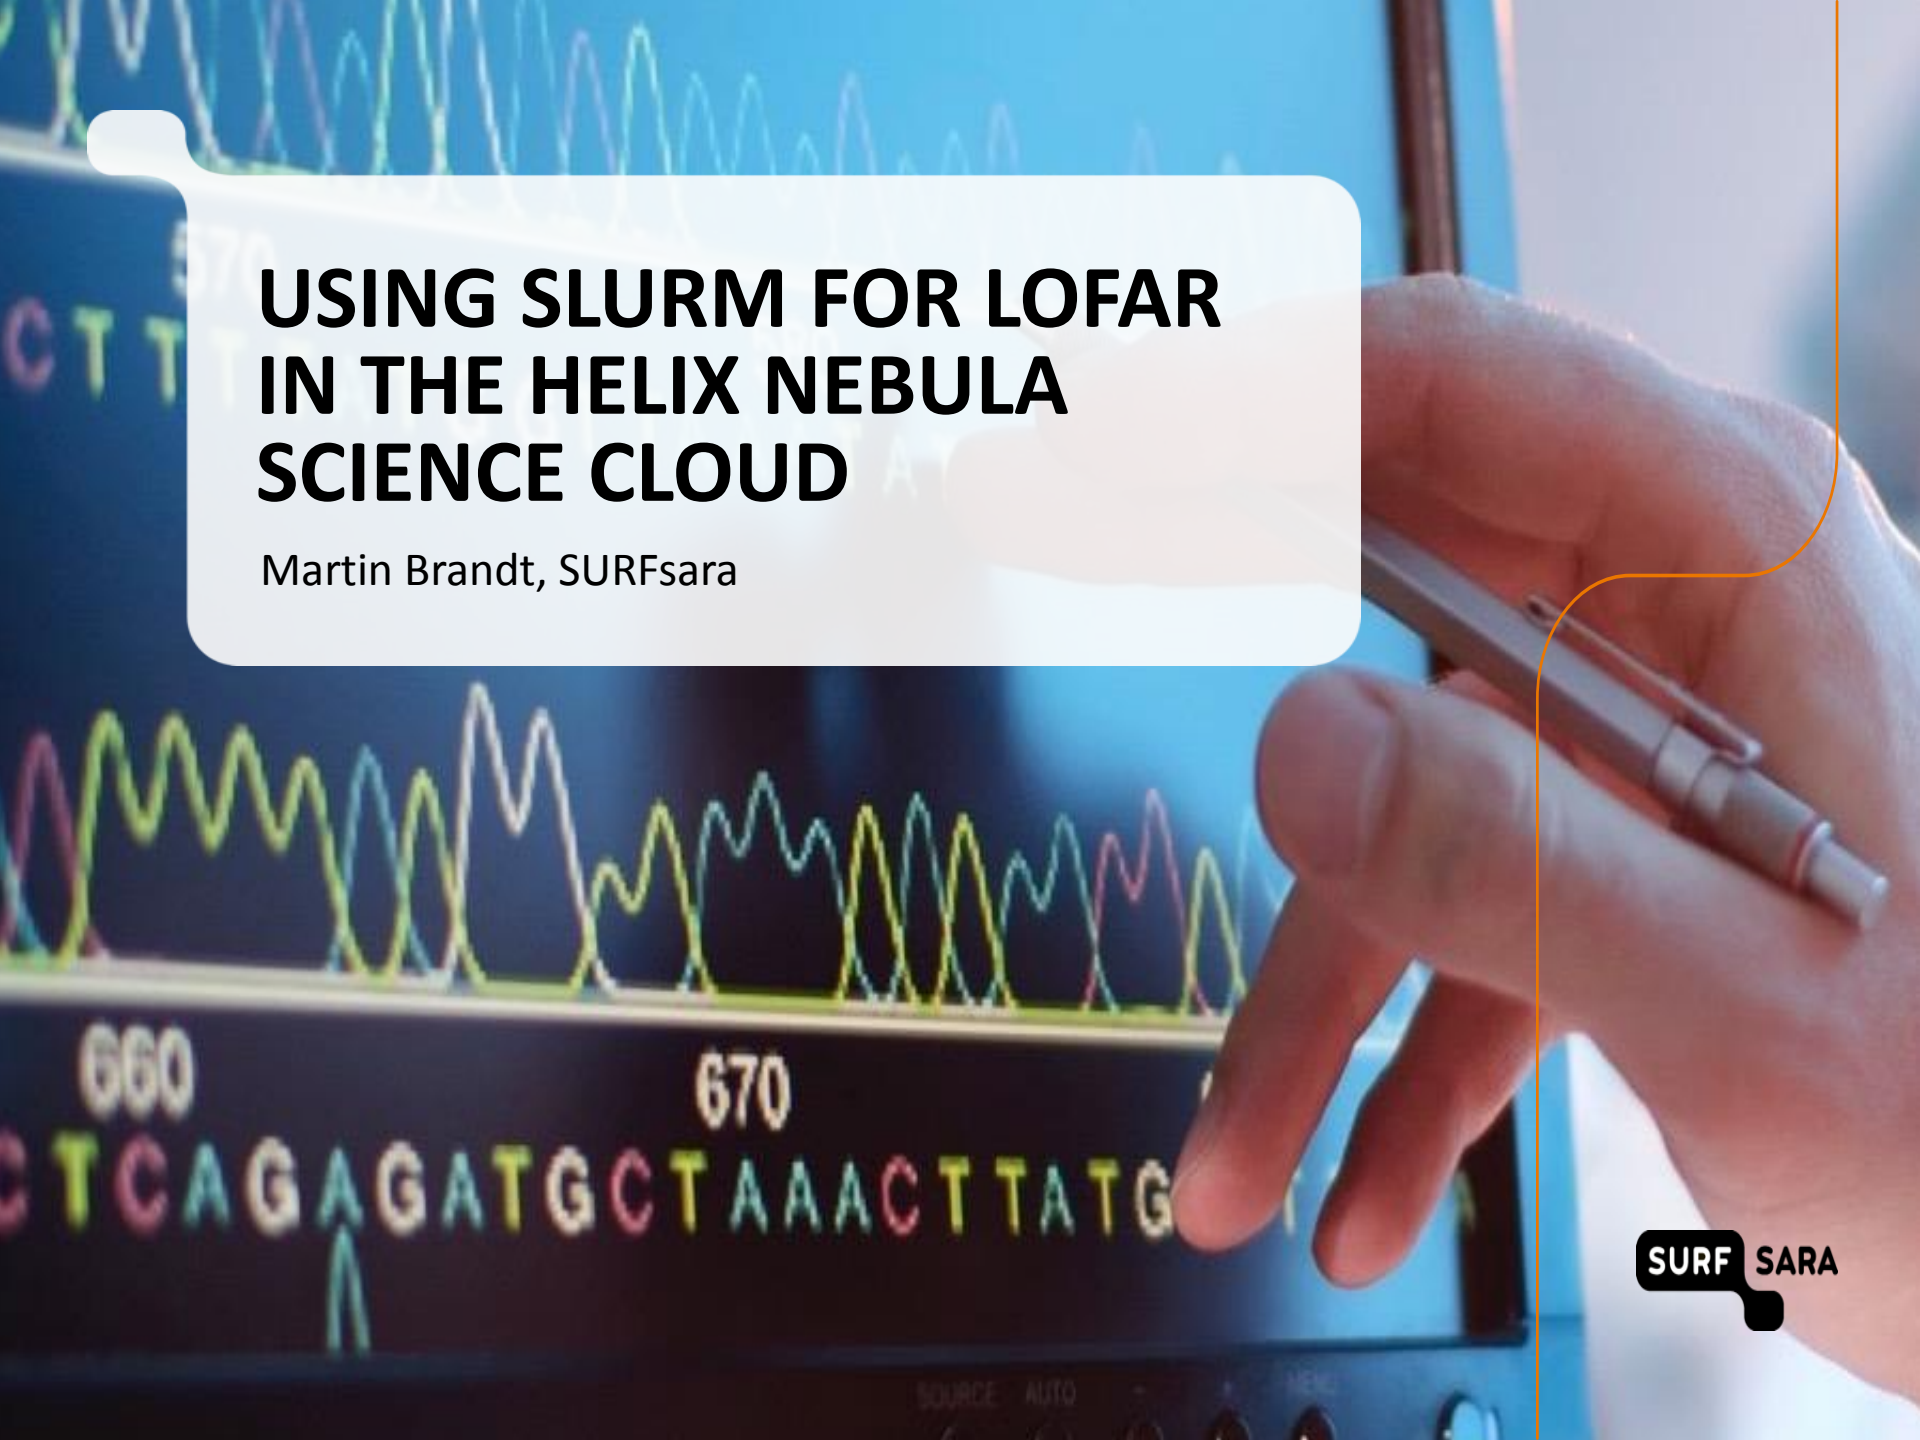

USING SLURM FOR LOFAR IN THE HELIX NEBULA SCIENCE CLOUD

Martin Brandt, SURFsara

SURF SARA

LOFAR and SLURM

Processing based on HEP experience but in LOFAR individual datasets are bigger. Dataset is

Pre-factor:
- *600 jobs (~4hrs)*
- *64 GB RAM*
- *0.5 TB disk*
- *Many workers*

Direction depended:
- *25 jobs 6 days*
- *256 GB RAM*
- *5 TB disk*
- *1 worker*

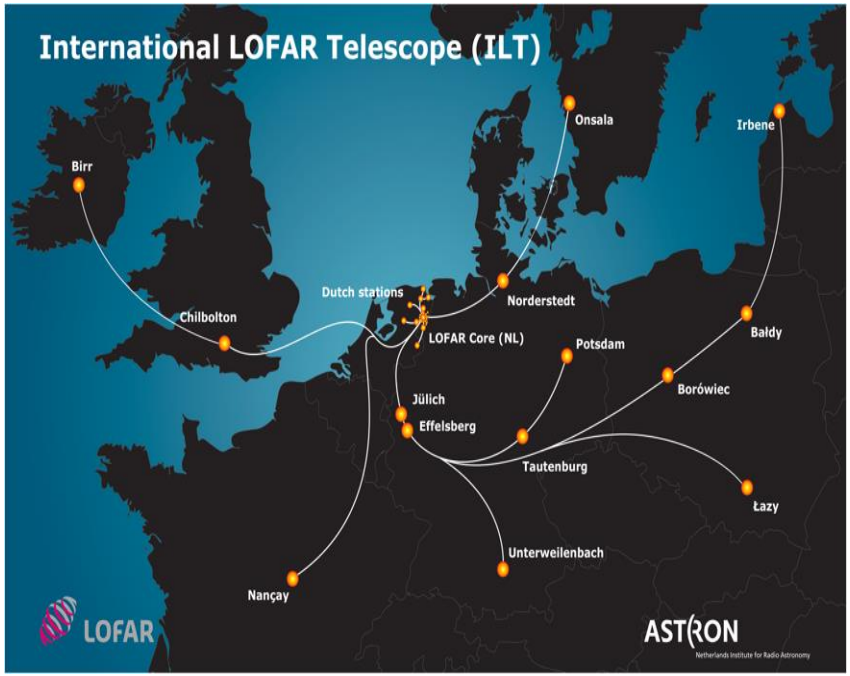

SLURM is common for large scale compute facilities, including LOFAR data sites: SURFsara, Julich, PSNC

LOFAR specific challenges:

- ***High Throughput Computing*** not ***High Performance Computing***
- dCache based storage

Using SLURM for LOFAR in the Helix Nebula Science Cloud

- Data transparency link with dCache local storage problematic, so Grid tools used
- Two implementations are different, not directly portable
- Much cloud knowledge needed to set up, but:
- Support during Helix Nebula project from providers was good

Important that for the future: specialized support for research should be available!

Conclusions:

- Good Prove of Concept for using commercial cloud for Lofar
- Lot of synergy with EOSC Pilot / Hub

SURF SARA

Thanks to:

Raymond Oonk (SURFsara)

LOFAR

Martin Brandt

martin.brandt@surfsara.nl

www.surf.nl

SURF SARA