Sustainable computing in HEP

Some local computer center views

Some local workflow views

Some global views

Martin Gasthuber, Yves Kemp, DESY IT
FH sustainability forum
DESY, 19.09.2022

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

Intro

- partition the duties – attempt from IT perspective (our – personally)
- dive into selected topics – explore dependencies, looking for options
- scenario for local changes – one of (many ?) options

- most of presented thoughts/ideas are from us (Yves, Martin) as DESY employee and worrying individuals
  - throw ‘rotten tomatoes’ to us, not the IT group ;-)

a few words about ‘players and duties’ in this game

- this event is about ‘**Sustainable computing in HEP’** - think global act local (buzzword, but)
  - real effect (reductions) come from **local** facilities (computing/networks/storage)
    - that’s where computing can be done wrong – and better of course
    - main players in the local scope
      - **management** - reflect in real doing/change, interaction with politics/funding agencies
      - **administration** - align regulations (i.e. procurement rules/strategies)
      - **we** (IT and other groups at DESY)

      } aligned decision hierarchy
  - **community** (WLCG/LHC – the ’HEP’ here/simplicity) could help with
    - changed expectations – from power (HEPSpec) to work (events processed/generated) per week/month/year
      - specify/agree on which can be done with relaxed ‘time to complete’ – adopt central scheduling
    - more efforts in efficient code – platform agnostic and platform specific (today all is X86)
  - **you** – the (main) initiatiors of computing activities
    - lot’s of ‘good practices’ are mentioned already

⟹ revision of plan & decide & daily use

short glimpse into local history

- until the very recent months (single digit !) nothing really seriously (from IT perspective)
- personal experience
  - triggering this topic since ‘mid of 90’ leads always to ‘unambiguous reflection’
  - main argument – looking at our power costs (only Euros) don’t even start thinking about that.
    - GHE – no subject
    - missed opportunity to start conversion in periods of lowest power costs - not a real subject of discussion
- decisions future invest (machines or infrastructure) are done in classical mode – lowest (real) costs drives the decision at all levels (user, planer, administration, management)
- efficiency gains never leads to lower resource consumption (and mostly to higher GHE)
- what thoughts/ideas were born at other DESY groups already we might never heard of ?
  - extend this forum ?

We are from the computer center.
We have always cared about sustainability

… or what?

Efficiency in IT … and DESY IT / Computing Center

- Always a trend in IT: Optimize efficiency
- Efficiency = Output / Input
  - Ouptput might be scientific output, or more simple CPU cycles or TB stored, with a certain quality level
  - Input usually money (purchase, operational cost, infrastructure cost, personel cost, …)
- Multi-dimensional optimization, with fuzzy parameters

- Usually, we had increasing input (=resources) available
  - → make the most efficient usage out of it
  - Do a lot more with a bit more
- Sustainability will mean: Use decreasing input (=resources)
  - Do the same with less
  - Do a bit less with a lot less

Becoming more sustainable: Where to start?

- Good engineering practice:
  - Analyse the consumption
  - Start optimizing with the biggest consumer
  - You might find some low-hanging fruits … yes, but do not fool yourself

- We will not talk about:
  - Printing less
  - Switch off your desktop over night / during the weekend
  - Add an disclaimer to the email signature that you should send less emails
  - ...

  → You can do this, and you will be more sustainable … but compared to the big consumer, this is epsilon

Some facts about the DESY computer center

- Two computer rooms in building 2b
- ~4000 physical servers for different applications, plus ~1200 virtual machines
  - about 750 physical server for NAF&Grid, about 750 physical server for Maxwell HPC
  - about 1000 physical server for dCache, of which ~500 for XFEL
- Roughly 1.3 MegaWatt power consumption (no cooling incl.)
- Two power feed, each backed by battery
- Most systems currently in watercooled racks:
  - “Cold” water from central cool water ring to back door
  - “Warm” water from back door to central warm water ring

  → Helps heating buildings and HERA tunnel
- Server lifetime 5-8 years

… what have others done?

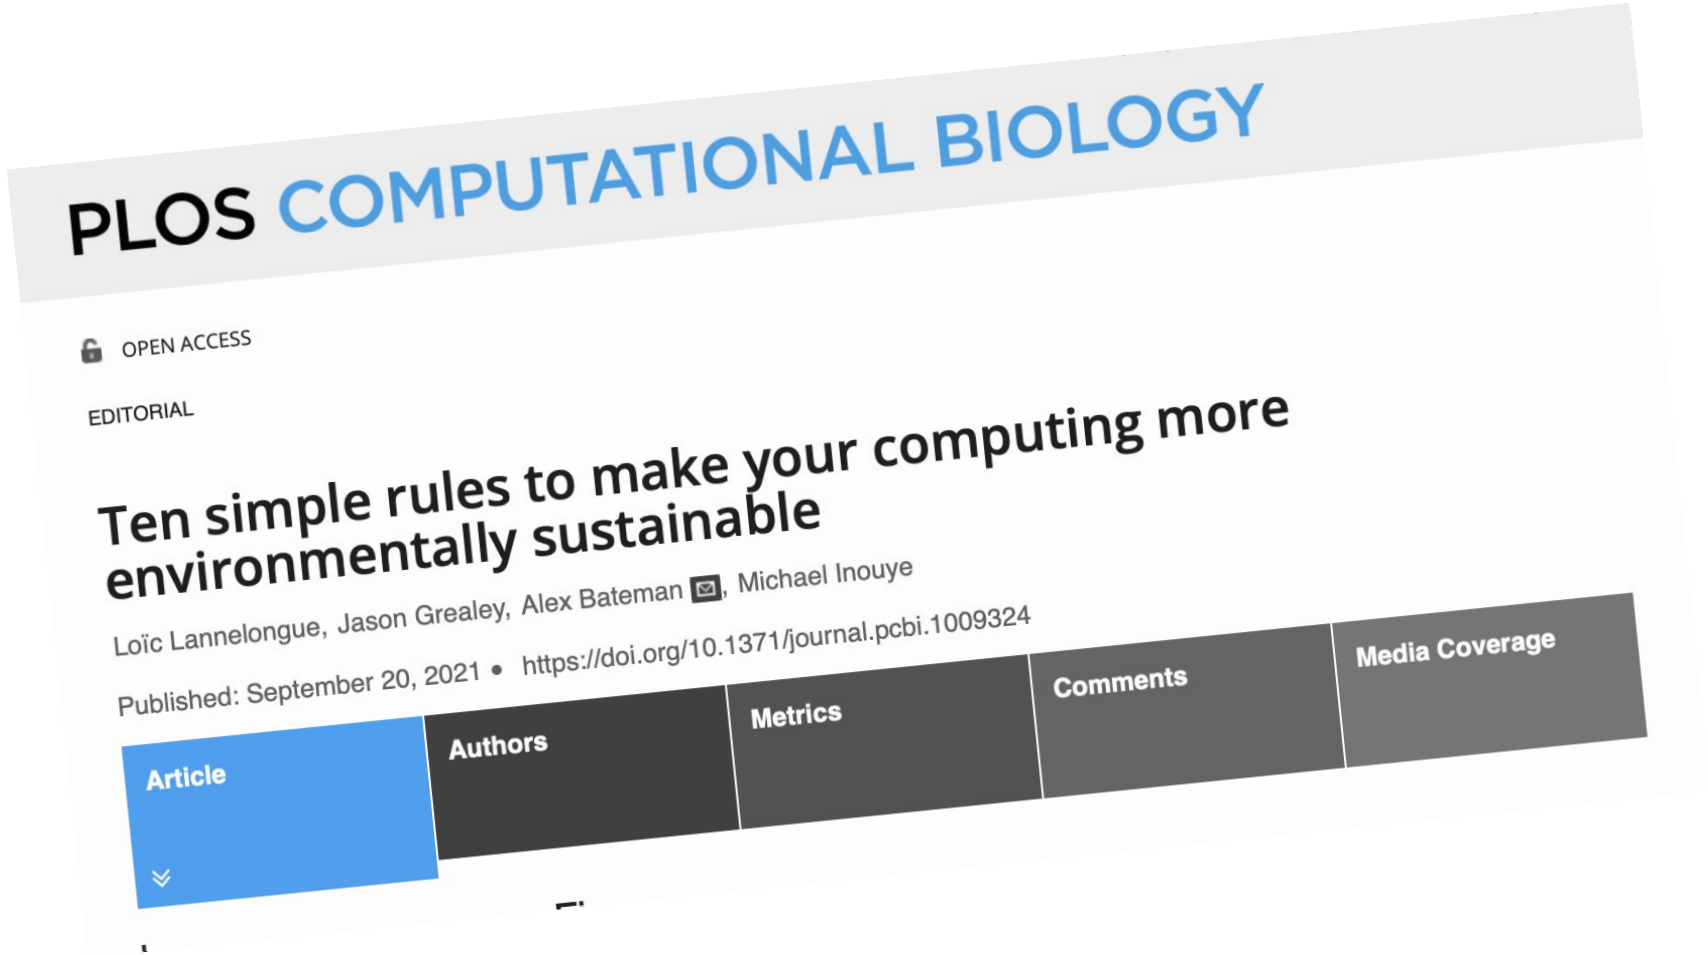

PLOS COMPUTATIONAL BIOLOGY

OPEN ACCESS

EDITORIAL

Ten simple rules to make your computing more environmentally sustainable

Loïc Lannelongue, Jason Grealey, Alex Bateman, Michael Inouye

Published: September 20, 2021 • https://doi.org/10.1371/journal.pcbi.1009324

Article | Authors | Metrics | Comments | Media Coverage

https://doi.org/10.1371/journal.pcbi.1009324

Rule 1: Calculate the carbon footprint of your work
Rule 2: Include the carbon footprint in your cost–benefit analysis

- At DESY-IT (and probably most HEP computing centers), this has not yet been done.
- It is not easy … need to incorporate e.g. building, production of server, power and cooling usage and infrastructure …
  - It is not just the kWh. It is also CO2, rare materials, water, …
- To do it right is more than just an afternoon one-off activity
- For HEP analysis: We would like to offer something like:
  - “Your NAF job #1234 took NN seconds, and has an energy footprint of MM CO2-equivalent”
  - … might start with a very simple model, and refine this
- Will need to be taken into account for all future projects. Petra-IV, FCC, …

Rule 3: Keep, repair, and reuse devices to minimise electronic waste

- What is the optimal lifetime of a server?
- Very very coarse ballpark number
  - “A server has a CO2-equivalent for production similar to CO2-equivalent of power usage over four years” [1]
  - ”A server consumes power over four years for the same amount of money spent for server purchase” [2]
- ... but again, different other aspects to consider:
  - Longer lifetime: More power consumption for same capacity. More floor usage in computer center
  - Shorter lifetime: Larger footprint for production. More hardware invest needed. Less maintenance
- no clear optimum for DESY so far.

- Repair, reuse: Needs manpower. Balance people invest vs. hardware invest

[1] DELL spec sheets

[2] Rule of thumb, used for first-order estimates … true for past and current electricity costs (orginally formulated by Knut Woller)

Idea: Computer Center in a Carport – more later

- It is (probably) not economical to keep old computers in the expensive computer center and offer them good quality electricity and cooling

→ Second life: Low-cost, low-QoS, no-CO2 Monte-Carlo Facility

→ maybe with on-site produced green energy?

Rule 4: Choose your computing facility

- Chose countries with low carbon intensity power
- ... also holds for locations within a country
- Works with the WLCG model

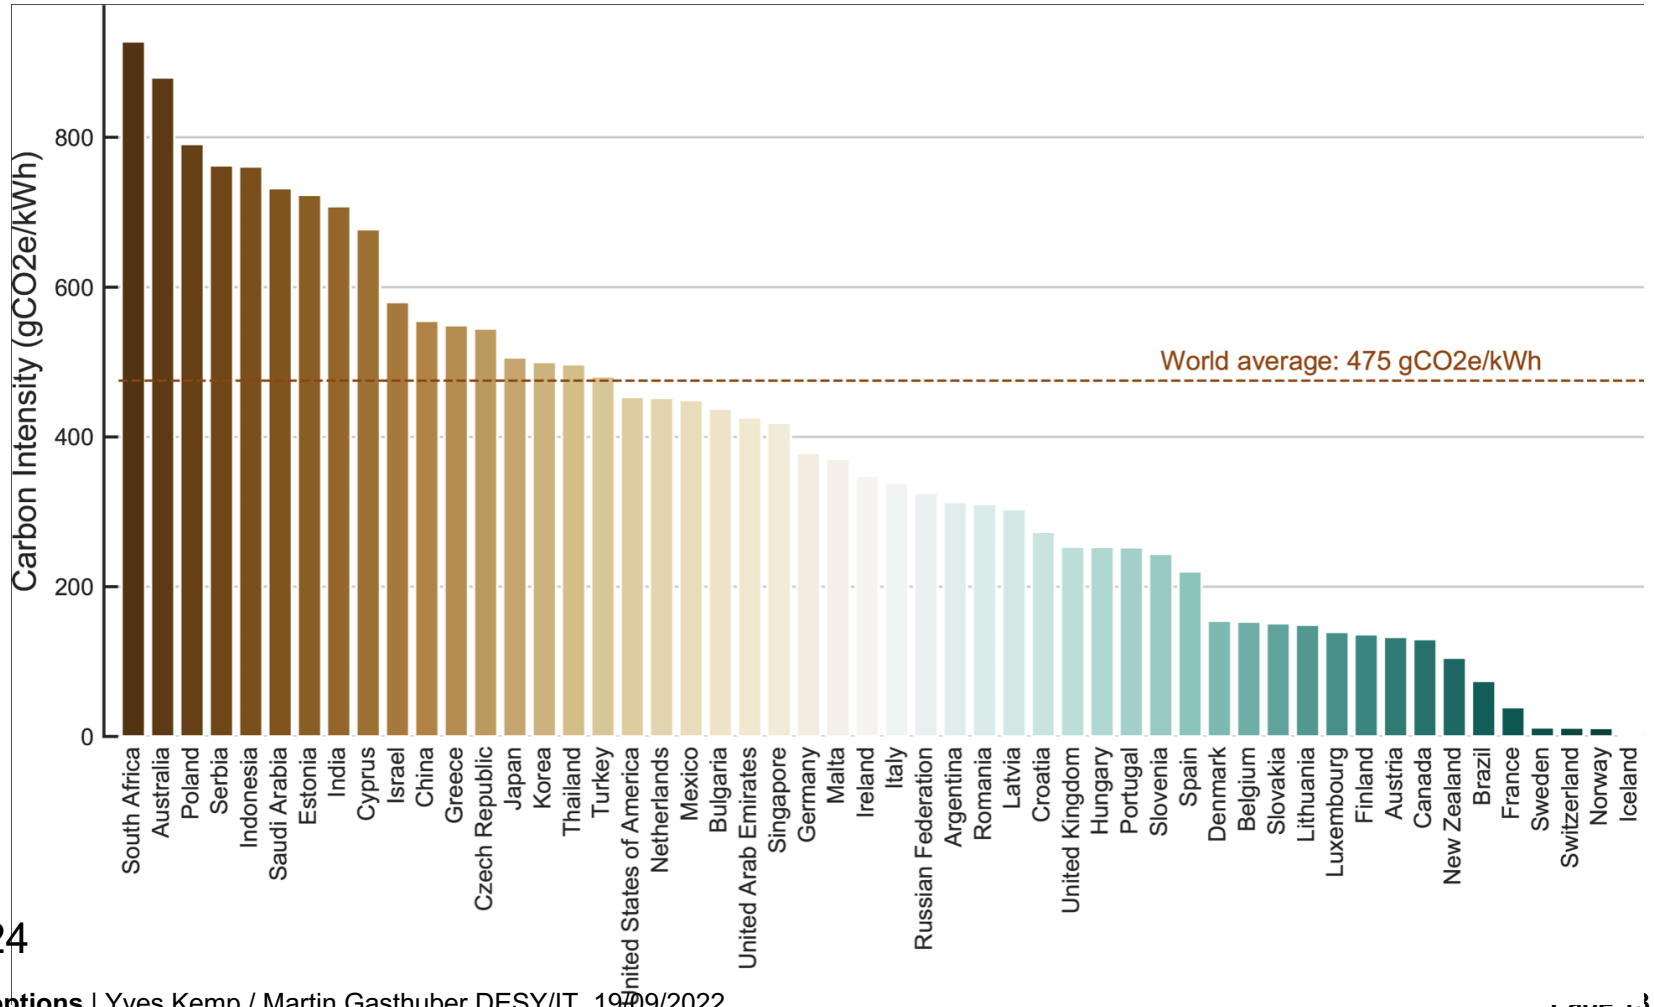

Figure: https://doi.org/10.1371/journal.pcbi.1009324

Let’s look at Germany

- Most of the 38% wind energy is produced in the northern half of Germany
- Reduce electricity transport, and place compute centers accordingly

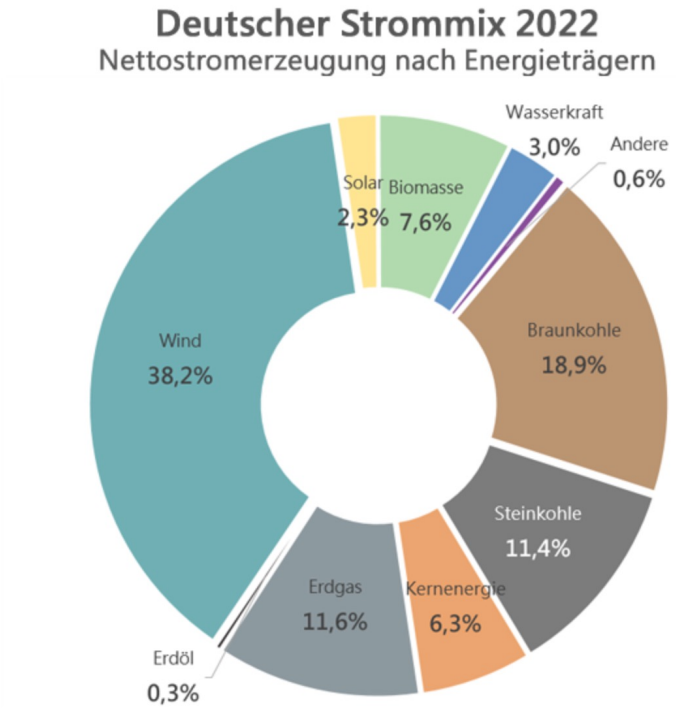

https://de.m.wikipedia.org/wiki/Datei:Windkraftanlagen_in_Deutschland.png
https://www.check24.de/strom/strommix-deutschland/ (nach Fraunhofer/ISE)

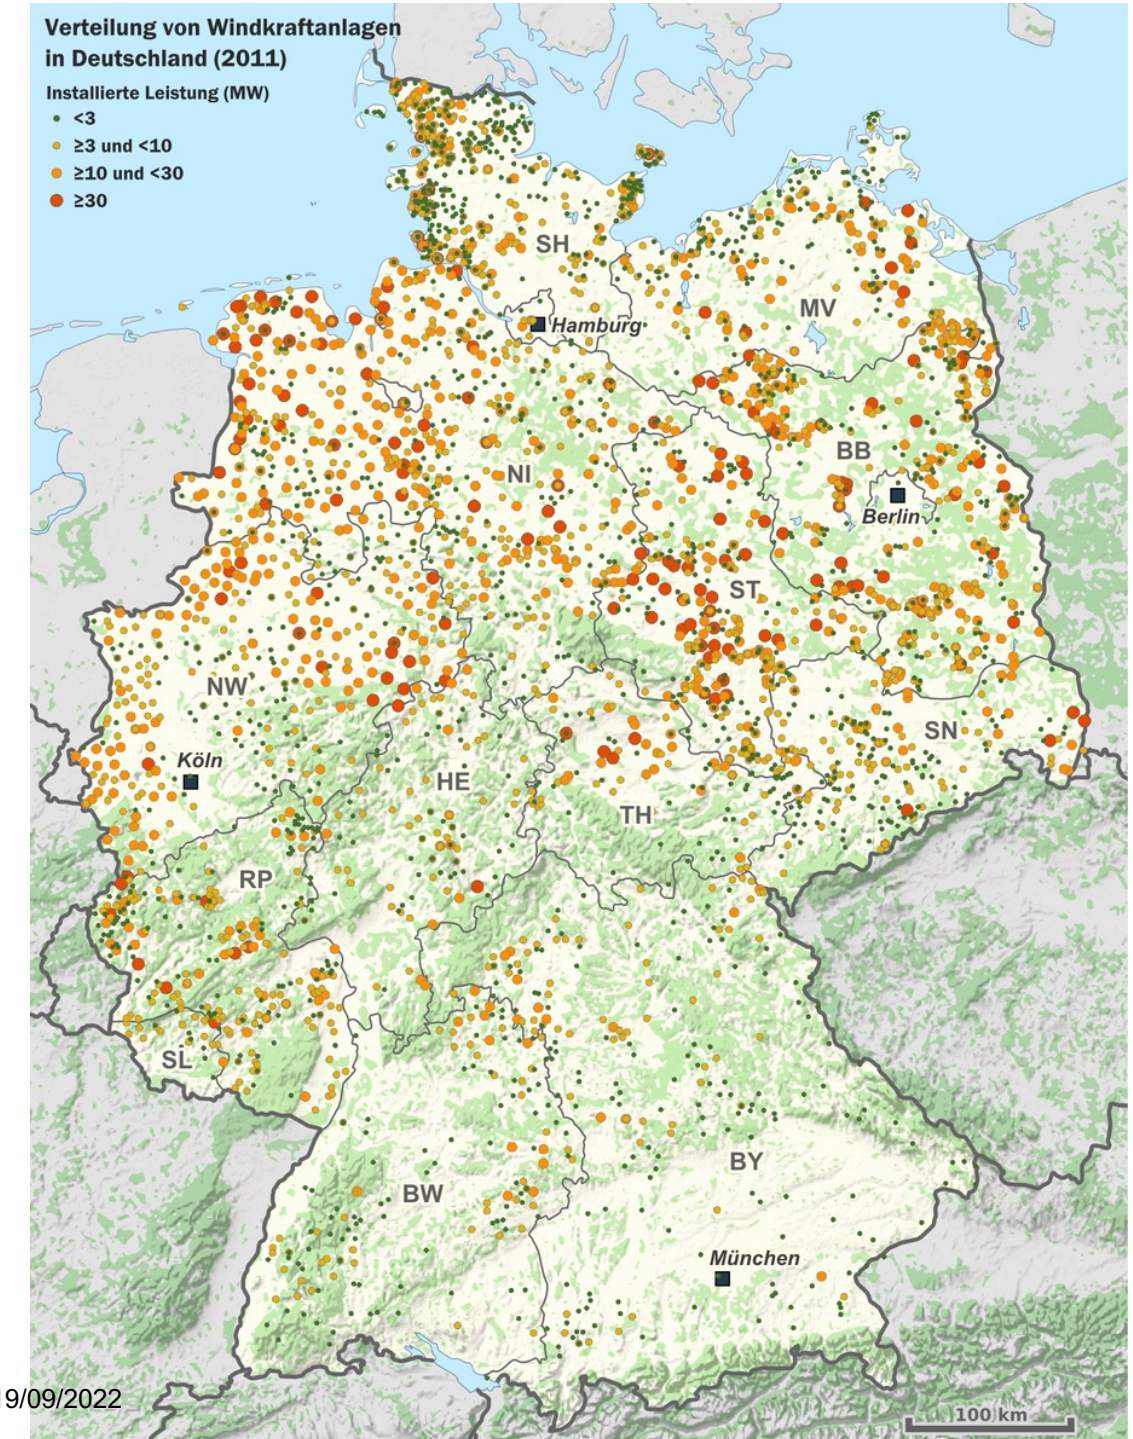

Smart and Agile

- Parts of the German “Energiewende” consists of using electricity in a smart way.
- This means: Consume electricity when it is there, do not consume electricity when there is none.

- We need to adapt our computing centers, our infrastructure, our workflow, our code, our habits!

Short-Term ideas for reducing power consumption

- ~~Reduce power consumption~~
- Use power, when there is (green) power
- Power down (parts of) batch farm if no (green) power
- E.g. combine prediction from e.g. https://gruenstromindex.de with up-time of batch servers
- If done clever, no job gets interrupted
- (Status: Investigation of feasibility)

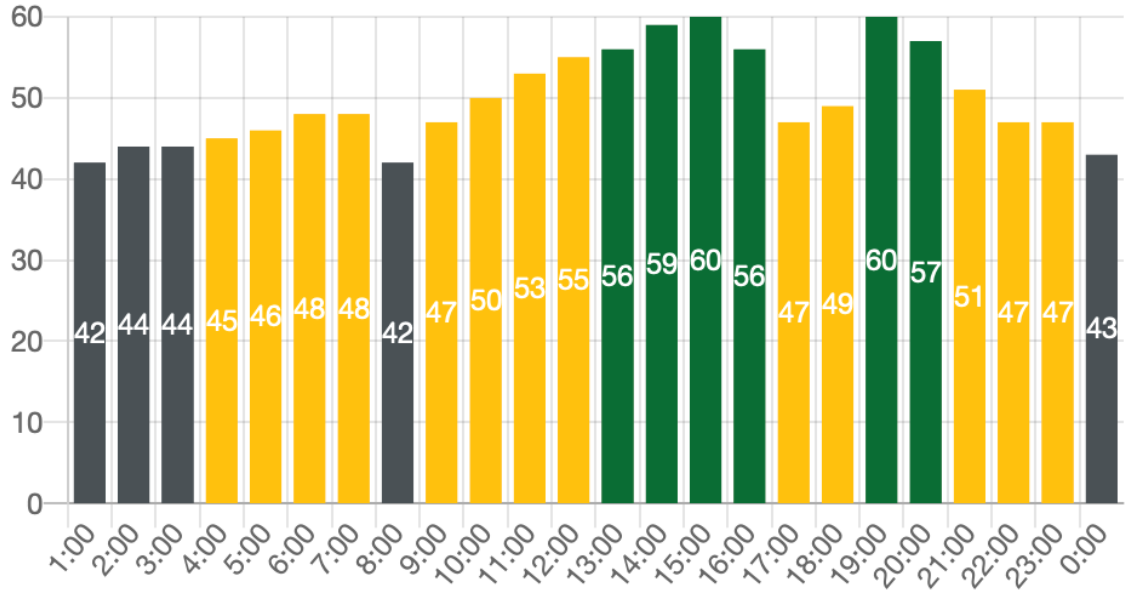

Forecast https://gruenstromindex.de
for 18.9.2022

- In the long, jobs need to become flexible:
  - React to short-term variations
  - Preemptible jobs that can be interrupted and relaunched

Rule 5: Choose your hardware carefully

- The “small” factors:
  - Chose RAM adequately: GHG scale with amount of RAM deployed
  - Wisely choose CPU: #cores, frequency, ...
- The “big” factors:
  - Choose an optimal architecture: x86 ? GPU ? ARM ? … ?
  - Computing center infractuctures: Cooling, Power supplies …

- Or do not chose any new hardware?

- **Need good communication triangle Scientists ←→ VO experts ←→ Computing Center**

Rule 6: Increase efficiency of the code

- Efficiency increase of code usually can be huge factors …
  - something we will never achieve by just looking at Computing centers!
- Needs training of all level of scientists
- Needs awareness with coding scientists, and with managing scientists!
  - Prefer code quality over development speed

- IT department can help with benchmarking and profiling
- IT department can help translate resource usage → CO2equivalent

→ Call for interdisciplinary Computational Scientists, bridging domain scientists and computing center professionals

Rule 7: Be a frugal analyst

- Think before running any job. You might be able to e.g. eliminate some phase space beforehand
- Scale up an analysis:
  - Develop with small test sample O(1%), test with larger sample O(10%), only then launch full 100% analysis
  - … you will not believe how often we see jobs in the NAF that simply cannot be successful
- Use checkpointing for long lasting computations, to reduce CPU consumption in case of failure
- Do benchmarks if you have different algorithms at hand (on small sample)

Another meaning of *frugal*

- Optimize the resource footprint of your data:
- Reuse data as long as possible!

→ **Data Preservation**

- In the HERA Data Preservation effort, we learned the importance of e.g.:
  - well documented, high-quality, portable code
  - well annotated data (“FAIR principles”)

- Thinking data preservation right from the beginning of a new project will:
  - Make you a frugal *and* successful analyst
  - Reduce the footprint of your data
  - Improve the quality and efficiency of your code (see previous rule)

Rule 8: Releasing a new software? Make its hardware requirements and carbon footprint clear

- Batch systems need to know your memory footprint. Try to give a good estimate:
  - Too little: Your job might crash due to resource limits → waste
  - Too high: You will be allocated memory which others cannot use → waste
  - Developers should indicate memory consumption … or you should make some tests

- Developers should include carbon footprint, so that users can compare:
  - Is there really an efficiency increase w.r.t another product?
  - Is there really an efficiency increase w.r.t to the previous version?

Rule 9: Be aware of unanticipated consequences of improved software efficiency

- Computing as well as HEP analysis has a history of having increased efficiency dramatically over the last decades
- Still, we consume always more power in computing: Efficient algorithms enable more scientific questions, thus trigger more computation

→ rebound effect

- … remember rule 7: Be a frugal analyst?

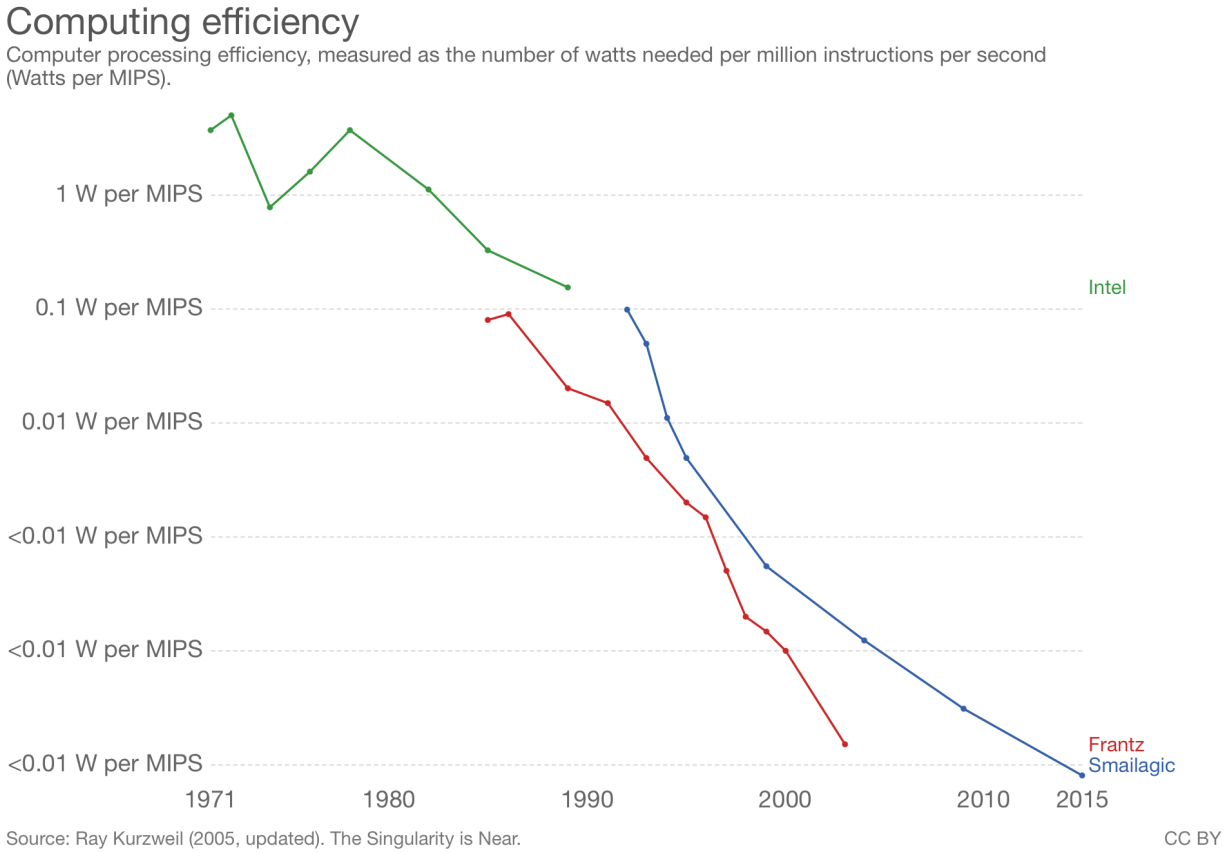

https://ourworldindata.org/grapher/computing-efficiency

Rule 10: Offset your carbon footprint

- e.g. plant trees, donate to initiatives, …
- IMHO, this is something for management, and should be a last resort if greenhouse footprint needs to be reduced

- First, try to *really* reduce your footprint *locally*

be more tangible

the ‘act local’ part of the story

- explore local option – the ‘where and when’
- better understand power market
  - aim at collaborating, not just producer/consumer relation
- identify tasks to be addressed by different groups @DESY
  - IT, administration, MKK, users, WLCG/LHC-Computing and the ‘heads’
- why not wait until SCOPE 2 emissions got zero (sometimes ;-) - @DESY: our external electricity power
- risks and opportunities on recent power price explosion
  - lowering the ‘budget for right invest’
  - too many ‘nervous’ discussions and decisions
  - linked to brown energy consumption – if we save costs (Euros) by using less ‘power everytime’ now, we reduce GHG emissions immediately
  - without more invest (money and brain) there will be no effective results on the large scale
    - management duty ?
  - booster to made right decisions ?

dive into (I)

- supply power – known universalities
  - zero GHE generation means wind & solar (exclude the discussion about ‘sustainability of nuclear power)
    - huge fluctuations in available power – the ‘when’
    - short term predictions works (i.e. 15 minutes ahead)
  - any grid (including large scale storage/buffering) is and will be expensive (compared to generation costs)
    - be as close as possible to the power generation
    - reduce GHE for building and maintaining – none of them are earth-friendly
  - we are close to the North Sea and related wind parks – close enough ?
- computing & collateral infrastructure
  - extend lifetime of components – we already have a decent number of ‘too old to be efficient’ – and growing
    - IT task – make HW components more standardized and adoptable to varying lifetime
  - minimize DC like building expenses - go for simplest environment – i.e. the large & old experimental halls
  - no power for cooling - from day 1 / look for simple usage of ‘pre-heated’ air (~30 degree) (i.e. tomato plant on roof ;-)
    - no cooling infrastructure even in the planning – except ’open windows/walls’ – free cooling

dive into (II)

- computing & collateral infrastructure (continued)
  - make power demand adjustable (i.e. 15 minutes interval) – biggest challenge for IT - from 1-3% to 100%
    - explore grid stabilization opportunities – i.e. Lancium DC as stabilizing factors – keep the grid at 50Hz
    - could expect (at least) 80% of scientific tasks to be candidates for ‘make it today or tomorrow’
      - discuss and agree for new ‘service agreements’ – short: get the expectations on all parties right
    - make ‘classical’ (expecting constant available power) computing as low as possible
      - lowering that power demands is directly could to GHE/brown energy – buying the annual power as green in the current situation will not really lower brown energy (formerly know as green washing)
  - simulations (workshop on sustainable HPC at SC2021) 50/50 split HW resources shows promising result
    - we have better suited workloads, and other special challenges
- planning
  - make (mandatory) GHG emissions included in all evaluations for future invest – finding the best alternatives
    - continue with the ‘only Euro based’ evaluation, but make GHG emission an additional port - go for 200Euros/tCO2
      - train people to include always, gain experience on emission sources – we have to do it anyway soon
      - make sure we do get things right today which will run for years/decades

dive into (III)

- administration / legal
  - make sure the above topic is not violating any (critical) regulations
  - can be applied to any external ‘call for tender’ – DESY contractors have to follow the same
  - …

summary

- lots of newly discovered links/dependencies – more to discover
  - requires groups to collaborate, not done before
- large number of individuals @DESY who are ‘worrying’ are there – and growing
  - skilles are widespread but very promising to be part of
  - demanding for settling the ‘how’ and ‘who’ – management and coordination
- looks ‘doable’ – why not going the first steps (and further)