Nikhef 50

PDP – Physics Data Processing

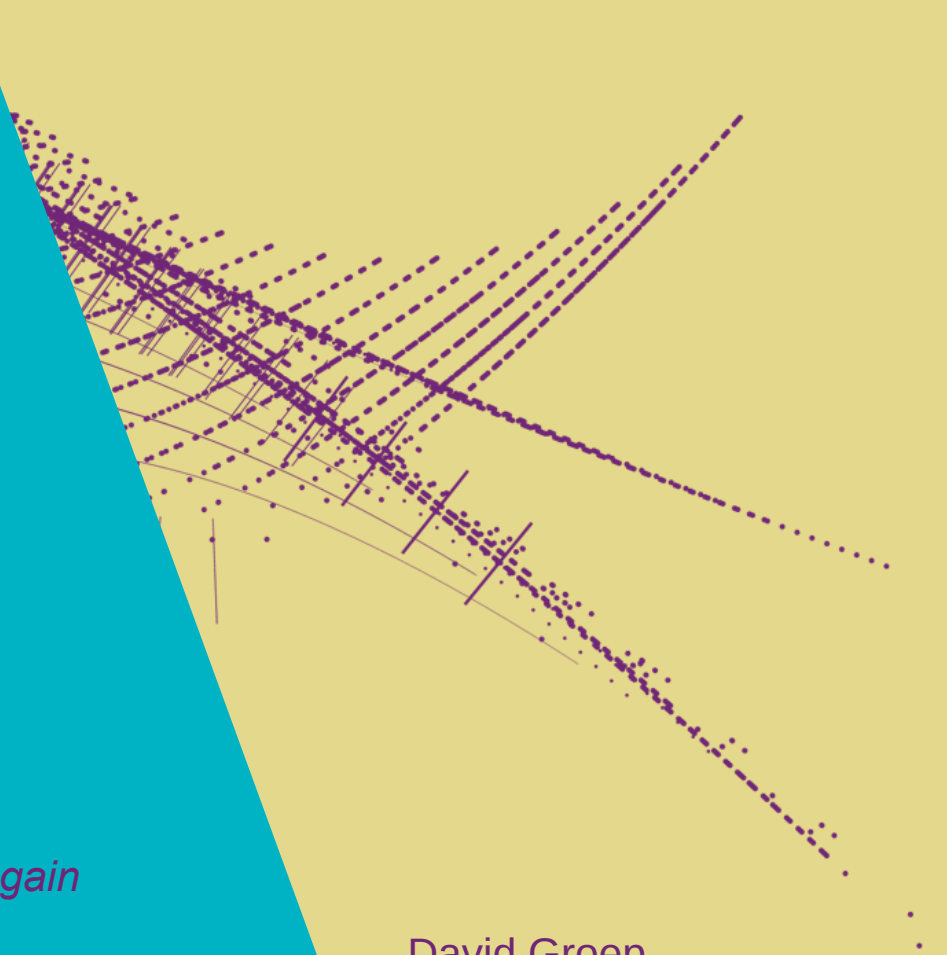

From when computing was scarce … and back again

Computing for the 50s!

David Groep
Nikhef Jamboree 2025

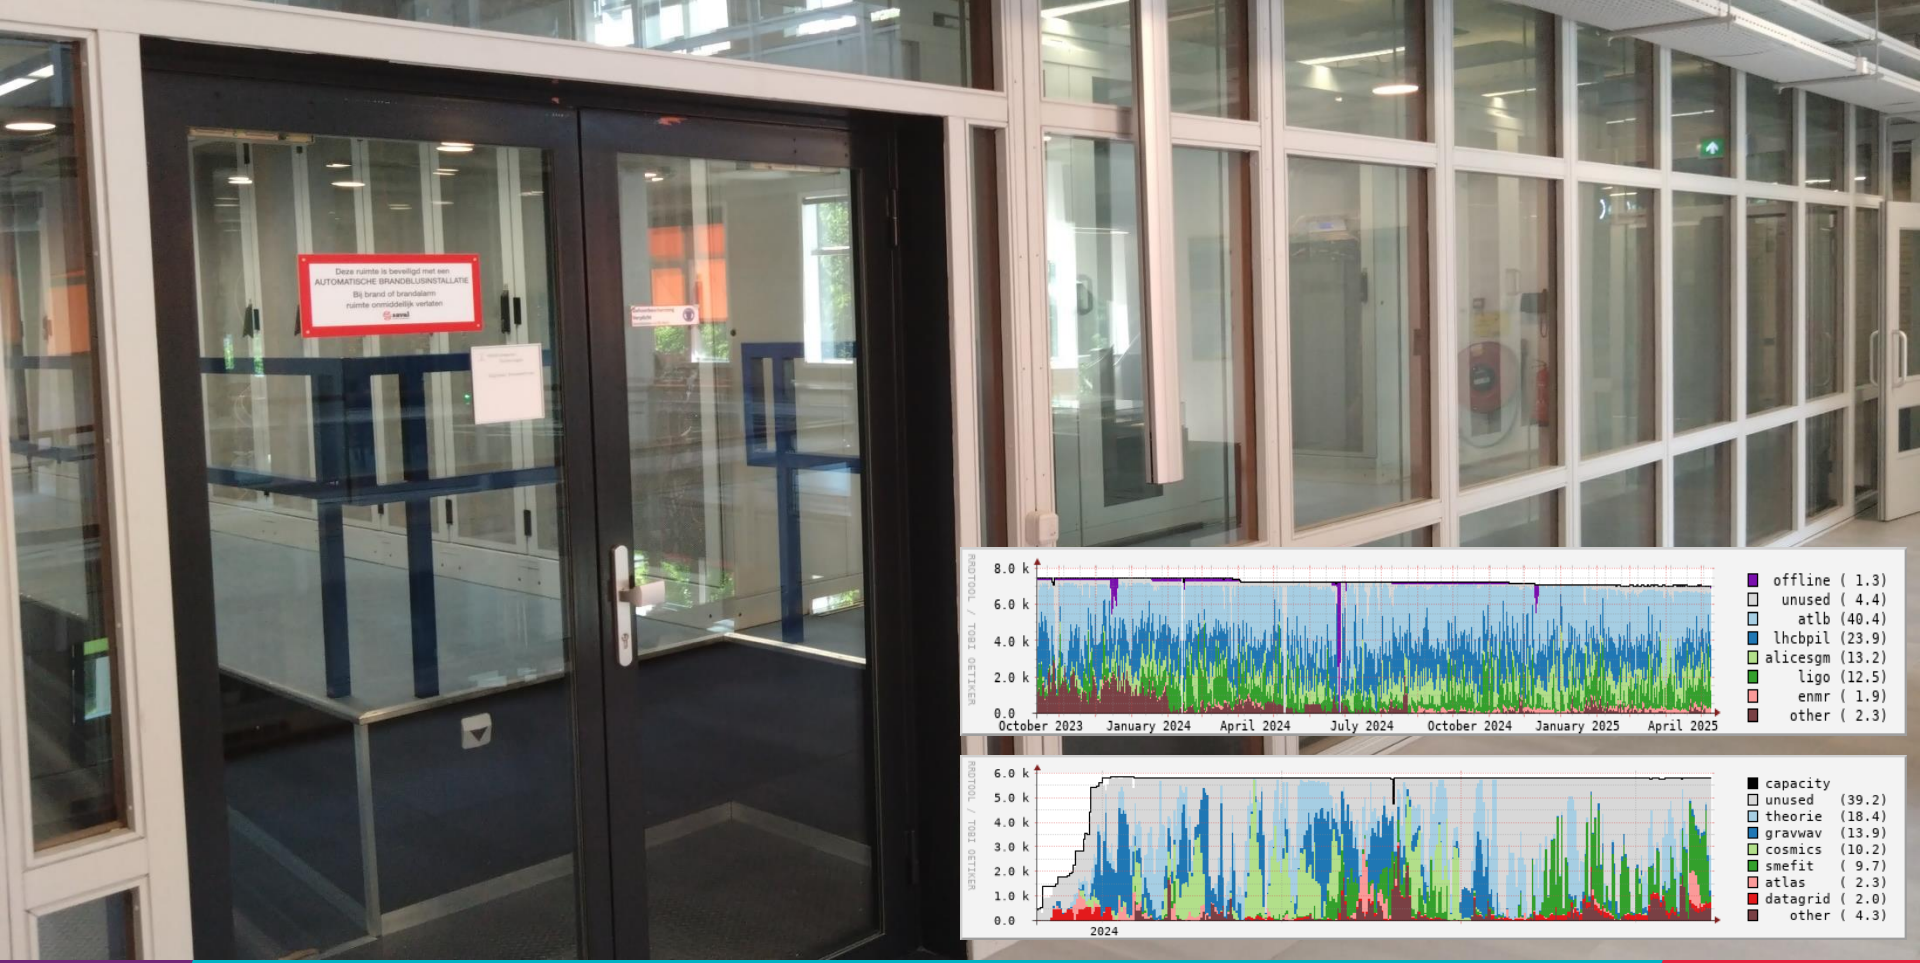

Computing for our 50s

Nikhef

‘Pillars’ of Nikhef Physics Data Processing

Infrastructure for trusted collaboration

- trust and identity (‘SSO’) for enabling communities
- managing complexity of collaboration mechanisms
- securing infrastructure for science, today & tomorrow

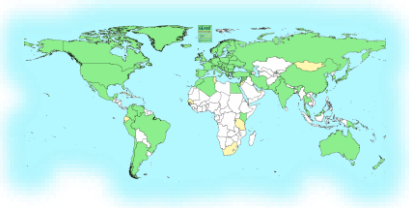

Algorithmic design patterns and software

- designing software for (GPU) accelerators, new algorithms, high-performance processors
- software design patterns for workflow & data orchestration

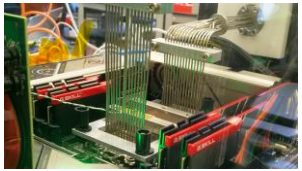

Infrastructure, network & systems co-design R&D

- building ‘research IT facilities’
- co-design & development
- big data science innovation
- research *on* IT infrastructure

From when efficient computing was absolutely vital …

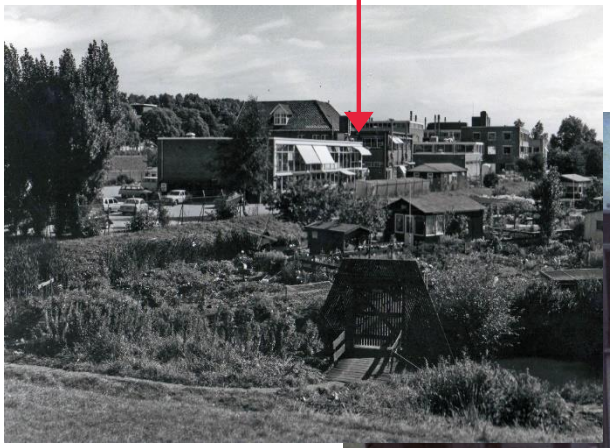

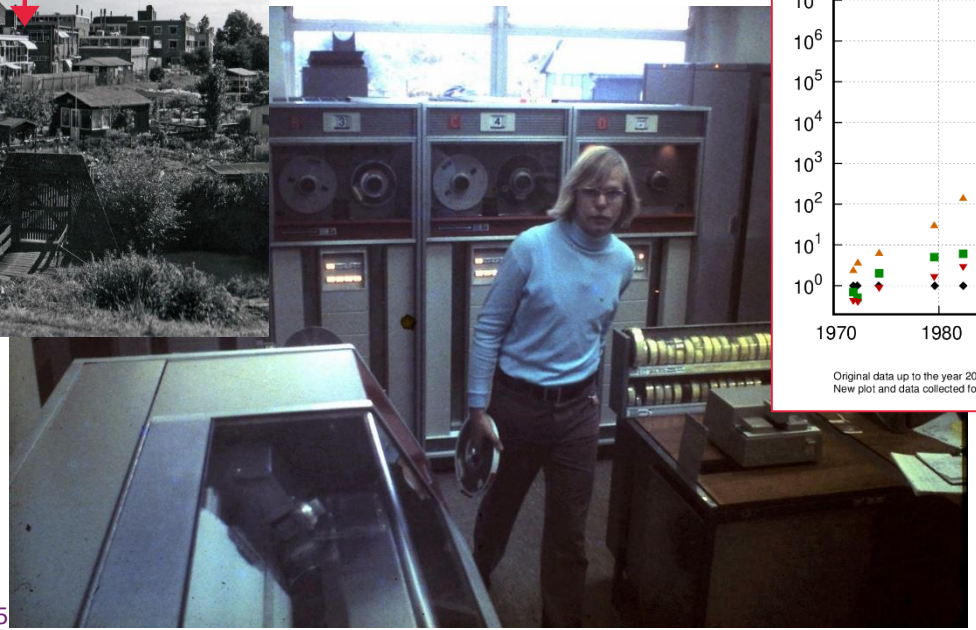

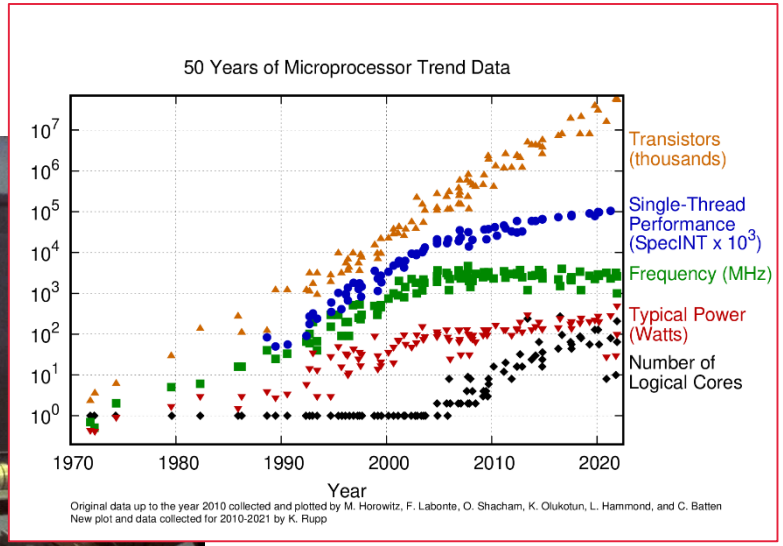

Nikhef 50

computing at Nikhef, ca. 1975

Image processor trend data: https://github.com/karlrupp/microprocessor-trend-data, by K.Rupp *et al.*, CC-BY-4.0

Stoomboot still is mostly CPU, but it's accelerating!

and you can also use CPU processing very effectively

```
In [23]: per_user_sums = df.groupby(['Owner']).agg({'corehours': ['sum'],
                                                    'occcorehrs': ['sum'],
                                                    'cputime': ['sum'],
                                                   'ExitCode': ['count']})


         per_user_sums['coreyrs'] = per_user_sums['corehours','sum']/(365*24)
         per_user_sums['efficiency'] = per_user_sums['cputime','sum'] / per_user_sums['occcorehrs','sum']

         # Sort the DataFrame by the sum of corehours in descending order
         sorted_df = per_user_sums.sort_values(by=('corehours', 'sum'), ascending=False)
         print(sorted_df.round(2).head(10))
```

	corehours	occcorehrs	cputime	ExitCode	coreyrs	efficiency
	sum	sum	sum	count		
Owner						
kchemina	4548645.06	4550042.12	2224954.74	124474	519.25	0.49
mbeekvel	3795796.35	3796607.73	2983553.85	471542	433.31	0.79
roystege	3618986.95	3628300.35	1864292.12	1300211	413.13	0.51
tsaracco	3228912.40	3327070.87	2940410.86	547799	368.60	0.88

Congrats to kchemina – and to tsaracco for being the most efficient user in the top-4

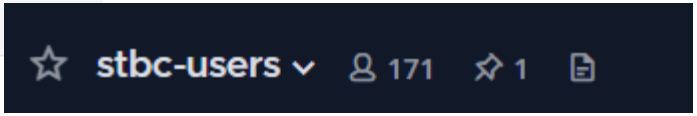

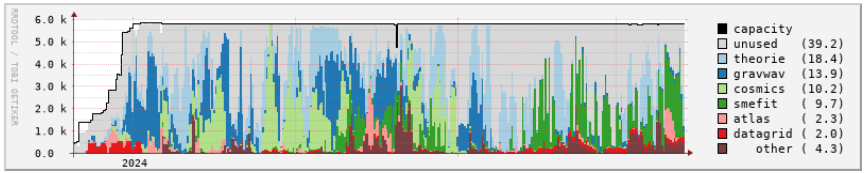

Stoomboot usage May 2024 – May 2025, analysis: Jeff Templon
See also https://www.nikhef.nl/pdp/stats/stbc/intern/stbc_summ_plots

But who is the most energy efficient?

someone discovered an empty Stoomboot on New Years day ☺

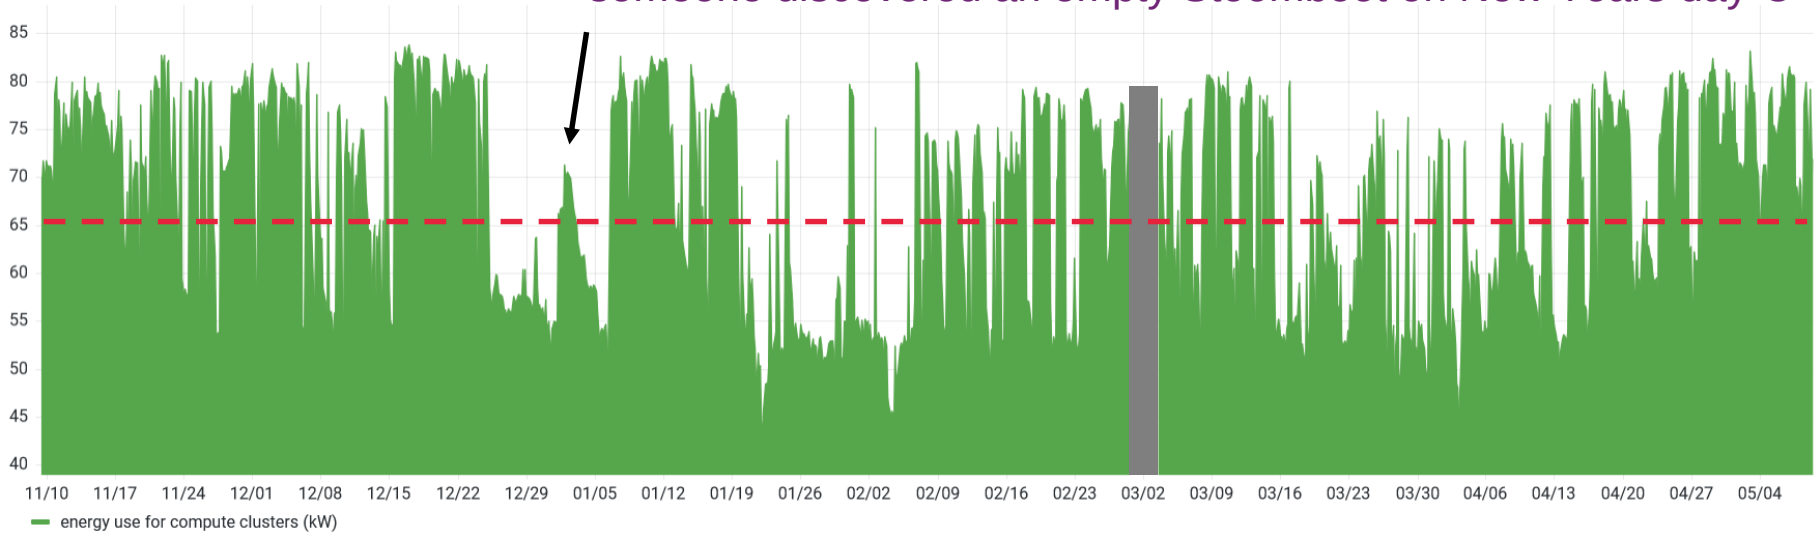

Data: aggregate energy use of all compute clusters (stoomboot + grid) at Nikhef from November 2024 till May 2025.
data from steker.nikhef.nl: `sum(ipmi_dcmi_power_consumption_watts)/1000`
See also https://go.nikhef.nl/stbc-energy for just the Stoomboot usage

The real cost of computing

This ~65 kW average for computing is the equivalent of 235 households

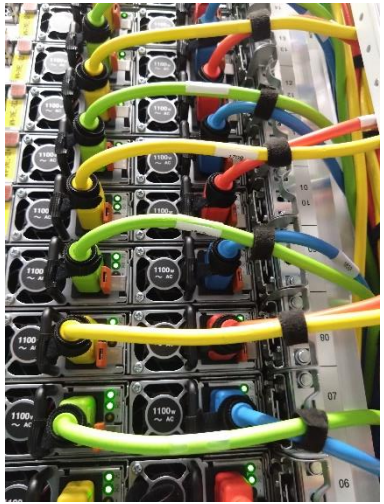

Compute only, all of NDPF (DNI-HTC and Stoomboot). Sources:
https://www.milieucentraal.nl/energie-besparen/inzicht-in-je-energierekening/gemiddeld-energieverbruik/ - 2420kWh/yr per household
Adobe Firefly: “a village with 235 houses, some of which are semi-detached, some council housing, and a scattering of villas, surrounded by fields and nature, with one overhead H/T power line running towards the village” – generating 4 images takes ~20Wh (charging a smartphone 4 times, https://www.technologyreview.com/2023/12/01/1084189/making-an-image-with-generative-ai-uses-as-much-energy-as-charging-your-phone/

Which is an expensive exercise …

And you have to cool the heat away from H234b, adding +30% (“PUE 1.3”)

- so ~~65~~ → 84.5 kW
- cost is ~ 15 ct/kWh
- so *just* the cluster compute alone is 111 033 €/year

- and you will need data from disk, which is equally power-hungry
- the typical load of our science data centre is 175kW, or 230 000 €/year

15ct/kWh: WCW rolling average energy pricing 2023-2025 including marginal grid connection costs and taxes

banknote image: https://pxhere.com/nl/photo/1327791

Now we *do* procure 100% green electricity at Nikhef …

Actual effective electricity mix different from the zero-emission we buy

Dutch average source mix

270 gr CO2e per kWh
(scope-1)

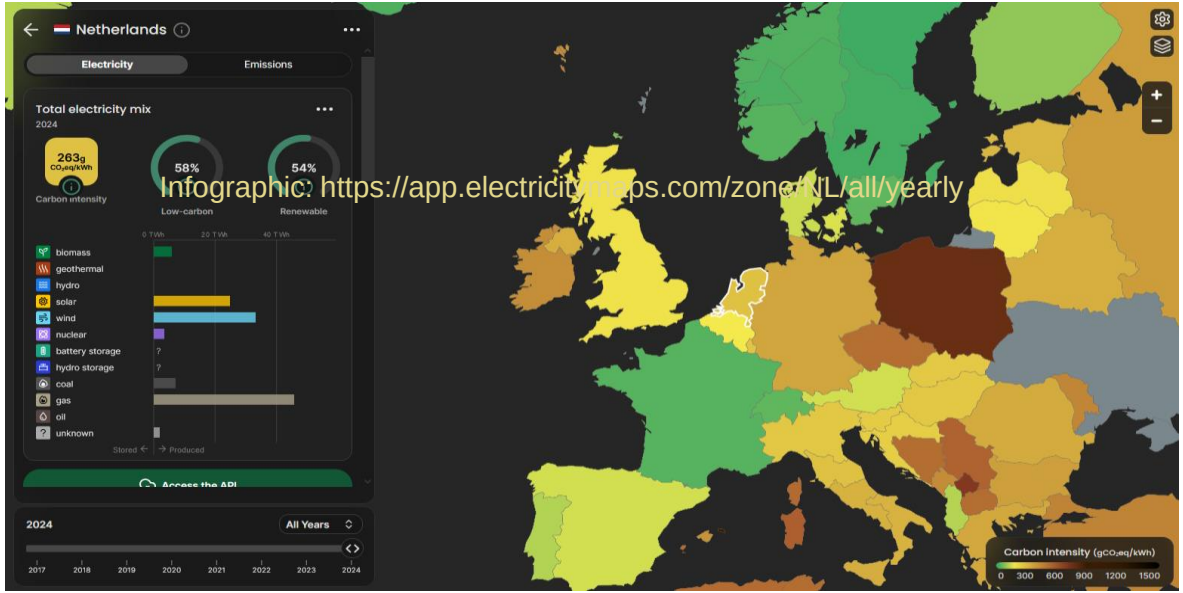

Infographic: https://app.electricitymaps.com/zone/NL/all/yearly

GHG data from *Methodiek CO2 emissiefactoren elektriciteit*, MilieuCentraal, https://co2emissiefactoren.nl/media/sources/240214-Notitie-CO2emissiefactoren-elektriciteit-pdf.pdf, data for 2023 (last year again better)

banknote image: https://pxhere.com/nl/photo/1327791

Let’s look at stoomboot, since that is where *you* can help

Summary: Yearly view for the Stoomboot NG

Top 8 Groups over Period

Metric	Trend
Running Jobs By Group	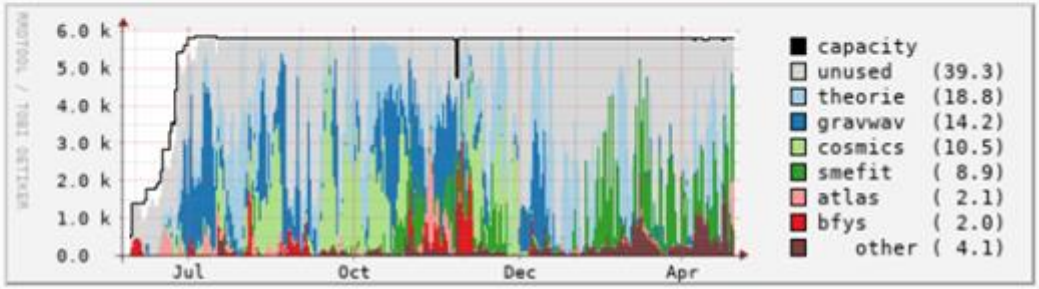

https://go.nikhef.nl/stbc-energy

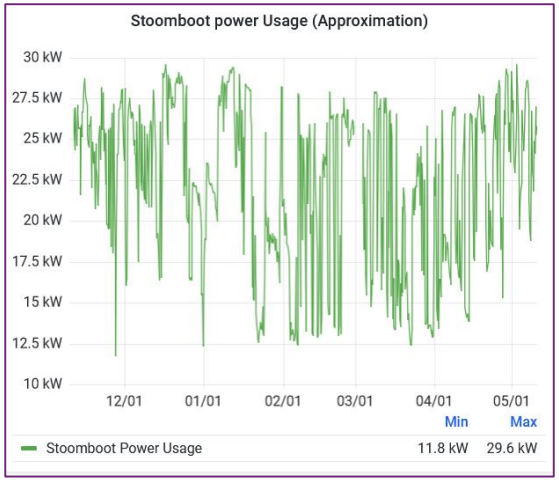

Stoomboot is not always full, by design, since otherwise you wait too long for your results (we see queueing complaints regularly …)
The grid side is more efficient, since it is 96% full all the time

The impact of the *stoomboot*

Emily's power plot for Stoomboot shows **11.8 kW idle**, with PUE 1.3 makes **103 MWh per year**

- at 15ct/kWh, this costs **20158 €/year** just for energy to start with…
- with 270g CO2e/kWh this is **36282 kgCO2e**
- to put that in context: AMS-GVA is 171.3 kgCO2e
  so **save at least 106 return trips** for enjoying having stoomboot ☺

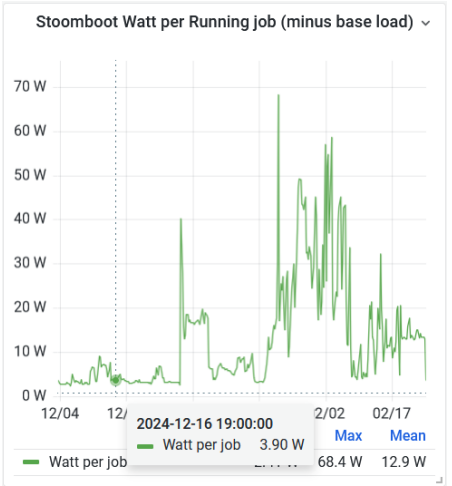

But **peak usage** power is **29.6 kW**
that is where you can help maximize the physics return thereof

- power used in jobs 29.6-11.8 = 17.8 kW for 5774 cores * 1.3 (PUE)
- computing you initiate takes **4 W/core**
- 100 jobs of 4 cores for 8 hours **3.46 kg CO2e** (and costs 2 € extra)
- put that in context, that is a **120 minute hot shower** (at 29 gCO2e/minute)

Pricing: WCW 2023-2025, GHG data MilieuCentraal, flight emission data https://curb6.com/footprint/flights/ https://www.bbva.es/en/general/sostenibilidad/soluciones-para-personas/huella-de-carbono-personas/repositorio/ducharse-a-diario.html gives 90-200g/5min, so on average 145/5 = 29 gCO2 per minute. Stoomboot energy use: Emily Kooistra, https://go.nikhef.nl/stbc-energy

So you want to be more efficient!

So: by thinking before submitting you can … shower longer?

or if you take the train for your longer trips …
why not *also* use that time to optimize your code?

or do you want to stand a chance for the future data rates?

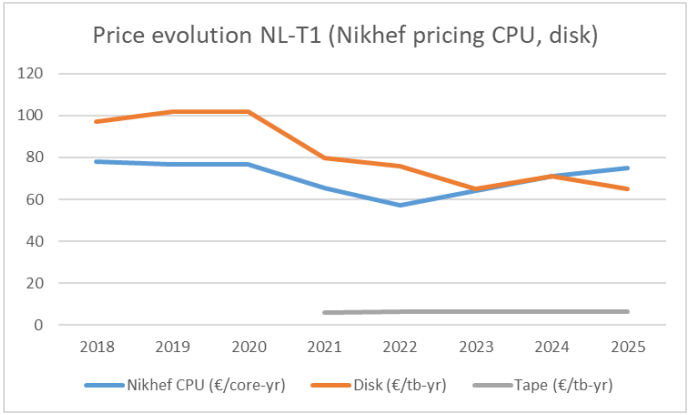

corrected for the (slight) performance increase per core

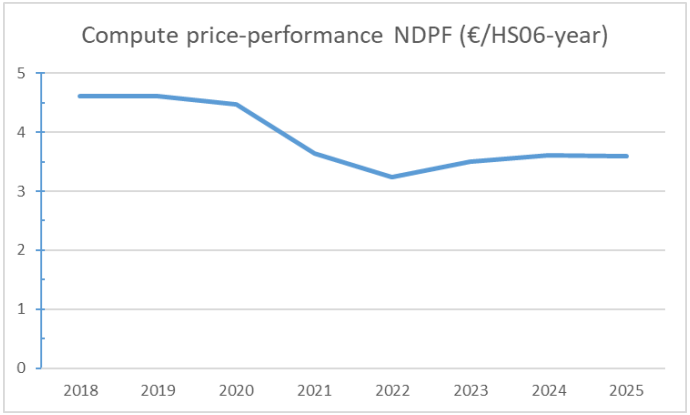

Image: Wikimedia. Price evolution: NDPF TCO NLT1-Nikhef-price-evolution-2018-2025 Effective TCO for stoomboot 156 €/used-core-year (due to 56% occupancy and support components)

There are quick wins …

If some cases it is Really Simple™ …

`-march=x86-64-v3` ➧ 15% improvement*

+ simple code optimisation gave additional 14% gain*

And use the proper ROOT version (6.32+)
for your processing, where vectorisation is the default

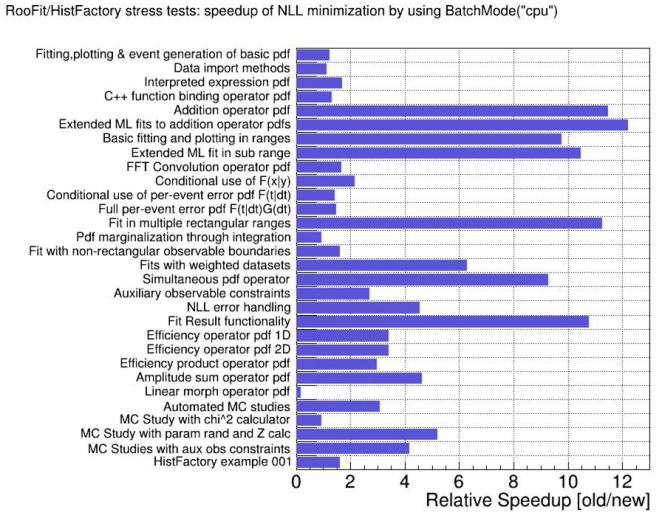

- Efficient use of scarce resources
- More sustainable in every sense (money and environment)
- **Faster physics results!**

* for the same amount of energy. optimised arch and code for Finesse library SuiteSparse KLU by Emily Kooistra.
Speedup graph: https://indico.cern.ch/event/1484669/contributions/6456378/attachments/3061358/, Jonas Rembser at the HSF workshop '25. Speedup measured with Root 6.28.04, see CHEP2023

PDP-CT – all about collaboration, so help us to help you!

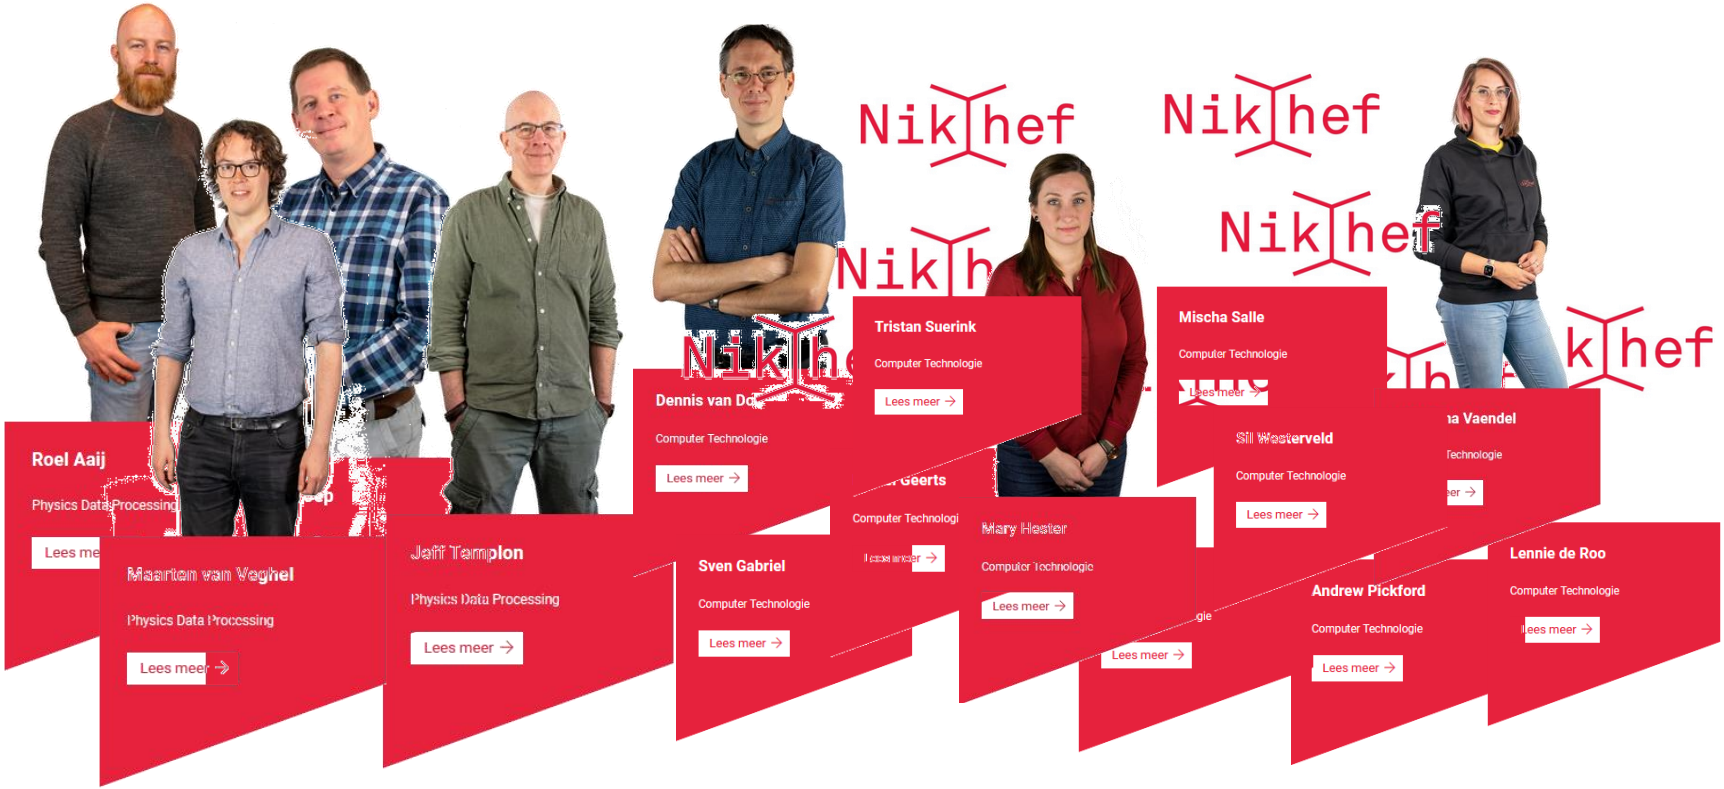

PDP and CT-PDP – meet us in Amsterdam at the science data centre, or at the Office Hours on the 1st Thursday of the month!

GPUs are more efficient as well … supported by MT & CT

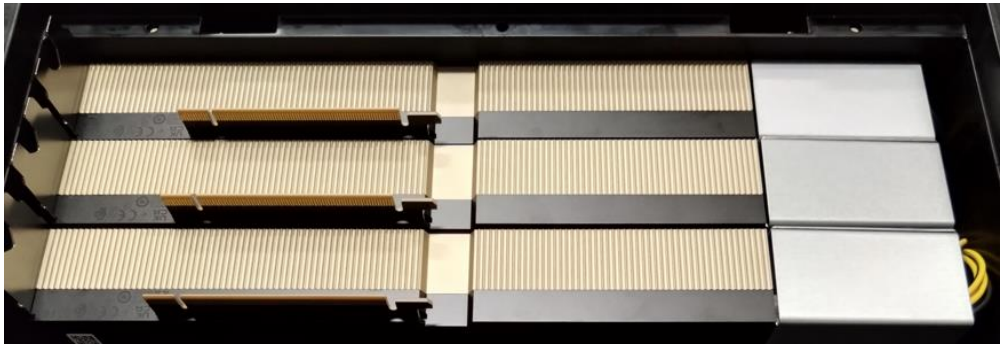

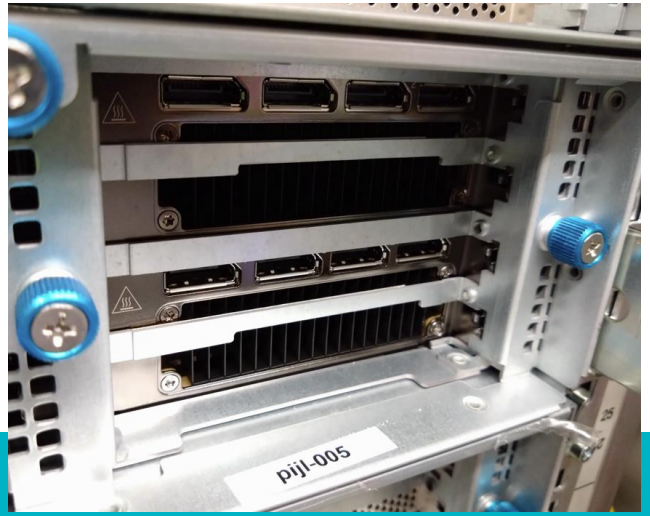

Thank Stan Heijnen for getting you safe GPU computing
(and prevent potential ‘true positive’ fire alarms from H234b)

Next up

Maarten van Veghel

Keeping it real (time): impact of PDP on LHCb

David Groep
davidg@nikhef.nl
https://www.nikhef.nl/~davidg/presentations/
https://orcid.org/0000-0003-1026-6606

Nikhef