Ecological impact of astronomy

The Ecological Impact of High-performance Computing in Astrophysics

Simon Portegies Zwart

[1] *Leiden Observatory, Leiden University, PO Box 9513, 2300 RA, Leiden, The Netherlands* [1]

[1] Non-anonymous Dutch scientists.

ArXiv:2009.11295; Nature Astronomy vol.4, 819 (2020)

An astronomical institute's perspective on meeting the challenges of the climate crisis

Knud Jahnke[1,*,+], Christian Fendt[1,+], Morgan Fouesneau[1,+], Iskren Georgiev[1,+], Tom Herbst[1,+], Melanie Kaasinen[1,+], Diana Kossakowski[1,+], Jan Rybizki[1,+], Martin Schlecker[1,+], Gregor Seidel[1,+], Thomas Henning[1], Laura Kreidberg[1], and Hans-Walter Rix[1]

arXiv:2009.11307

Miljenko Čemeljić, CAMK, Warsaw & ASIAA, Taiwan

xkcd

A webcomic of romance, sarcasm, math, and language.

Earth Temperature Timeline

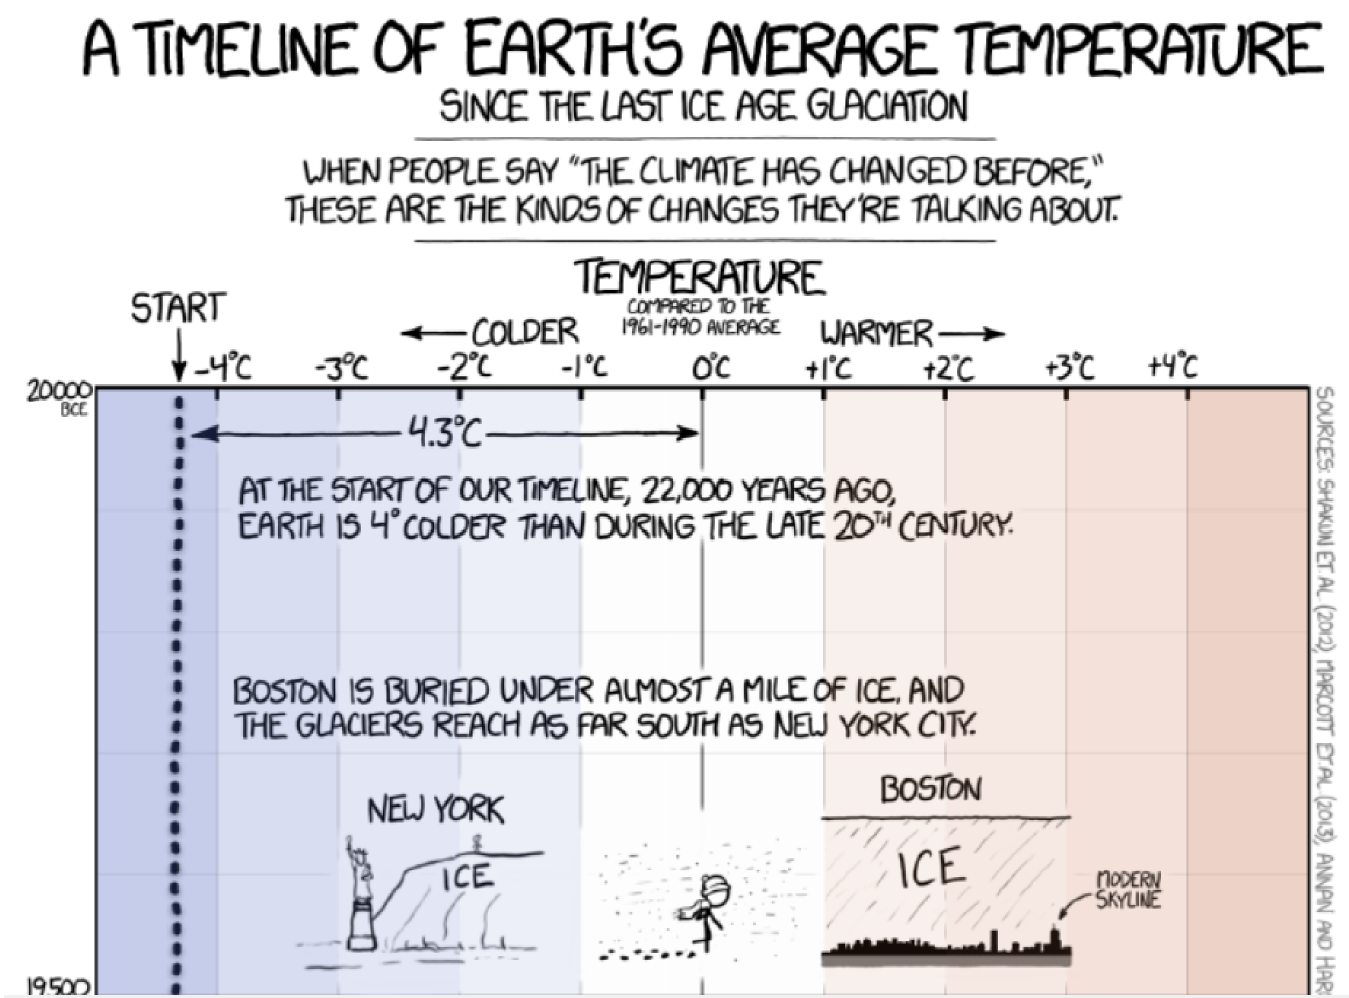

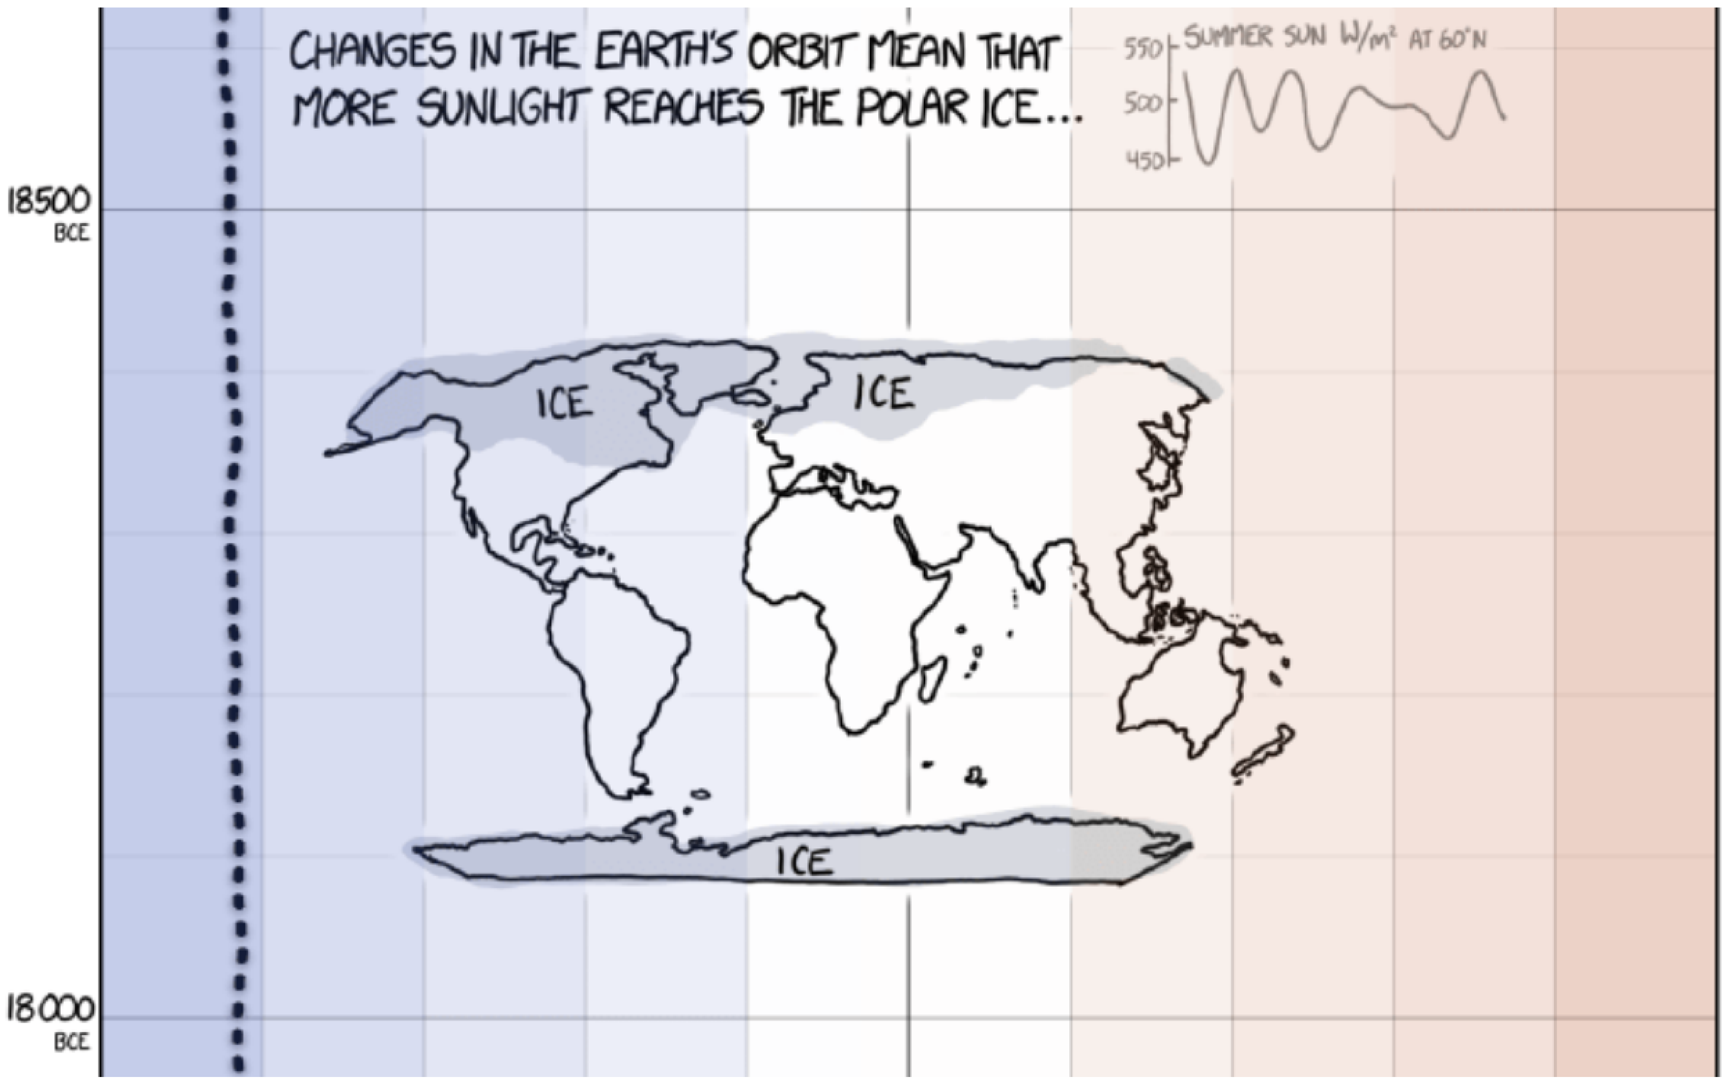
CHANGES IN THE EARTH'S ORBIT MEAN THAT MORE SUNLIGHT REACHES THE POLAR ICE...
SUMMER SUN W/m² AT 60°N
550
500
450
18500 BCE
18000 BCE
ICE
ICE
ICE

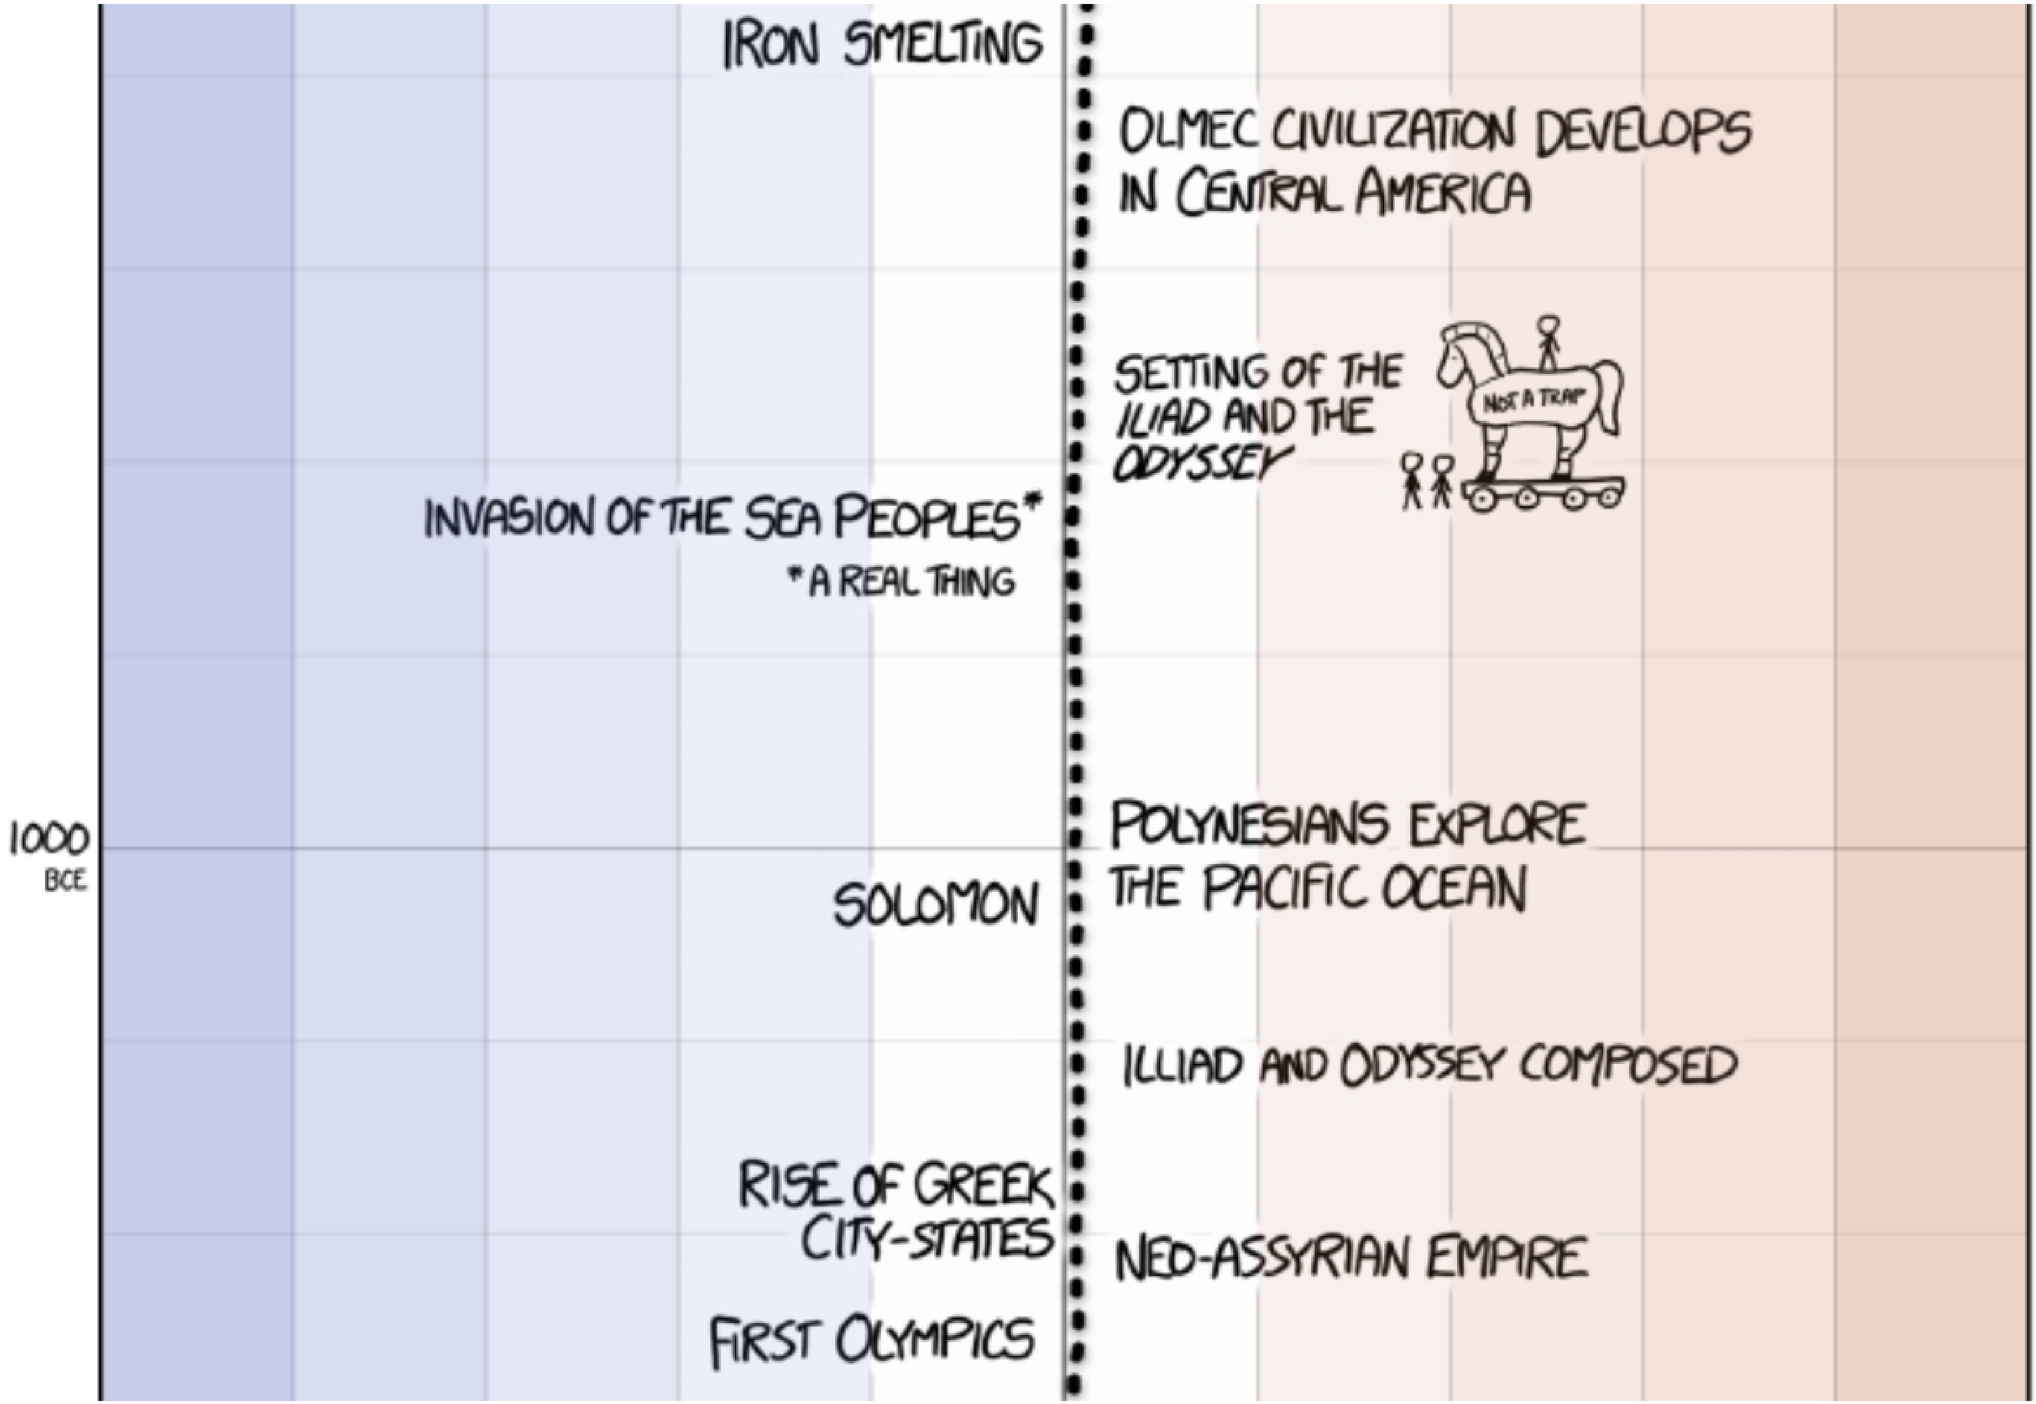
IRON SMELTING
OLMEC CIVILIZATION DEVELOPS IN CENTRAL AMERICA
SETTING OF THE ILIAD AND THE ODYSSEY
NOT A TRAP
INVASION OF THE SEA PEOPLES*
*A REAL THING
1000 BCE
POLYNESIANS EXPLORE THE PACIFIC OCEAN
SOLOMON
ILLIAD AND ODYSSEY COMPOSED
RISE OF GREEK CITY-STATES
NEO-ASSYRIAN EMPIRE
FIRST OLYMPICS

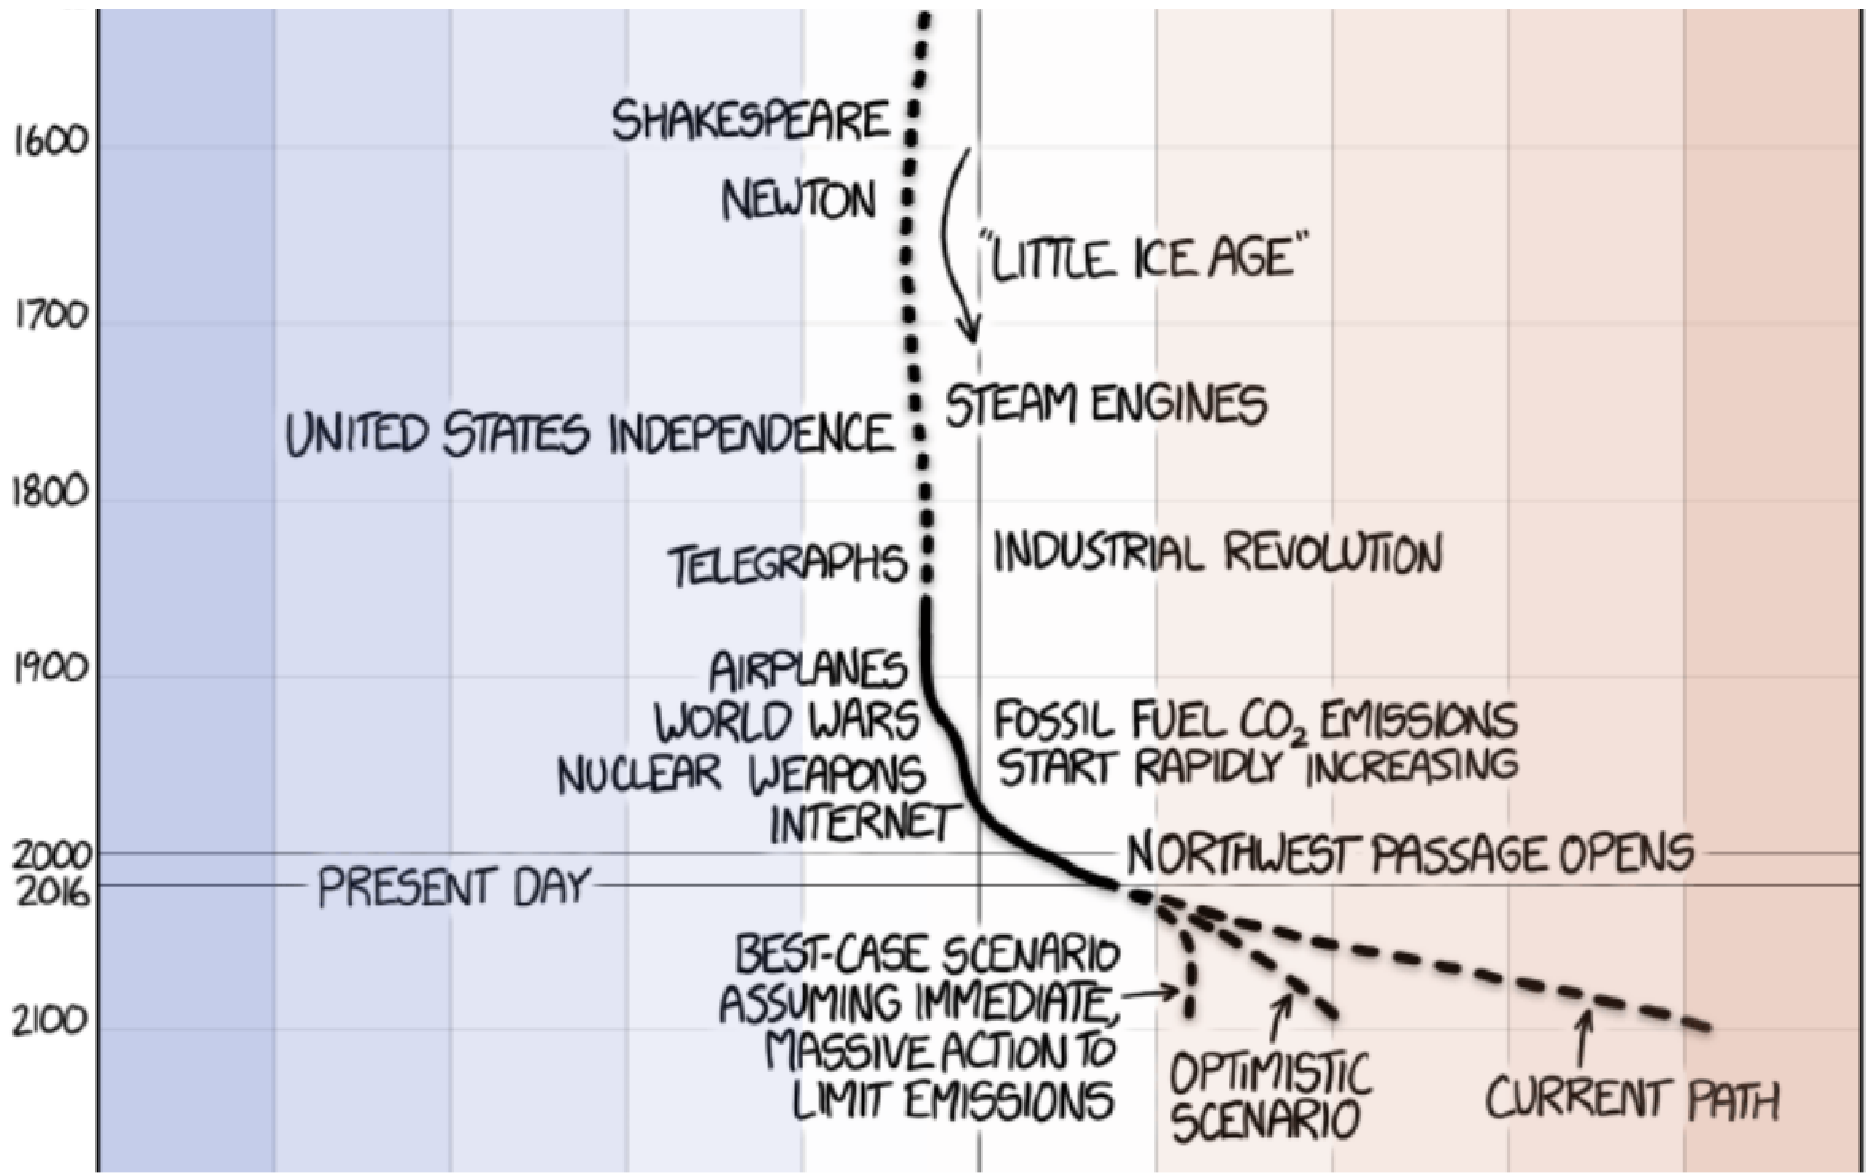
1600
1700
1800
1900
2000
2016
2100
SHAKESPEARE
NEWTON
"LITTLE ICE AGE"
STEAM ENGINES
UNITED STATES INDEPENDENCE
TELEGRAPHS
INDUSTRIAL REVOLUTION
AIRPLANES
WORLD WARS
NUCLEAR WEAPONS
INTERNET
FOSSIL FUEL CO_2 EMISSIONS START RAPIDLY INCREASING
NORTHWEST PASSAGE OPENS
PRESENT DAY
BEST-CASE SCENARIO ASSUMING IMMEDIATE, MASSIVE ACTION TO LIMIT EMISSIONS
OPTIMISTIC SCENARIO
CURRENT PATH

Outline

- Carbon footprint of astronomy and of computing
- How to achieve optimal performance & carbon emission
- The role of computing language on the ecology

- An institute's greenhouse gas emissions
- Measures to reduce emissions

The Ecological Impact of High-performance Computing in Astrophysics

Simon Portegies Zwart

[1] *Leiden Observatory, Leiden University, PO Box 9513, 2300 RA, Leiden, The Netherlands* [1]

[1] Non-anonymous Dutch scientists.

Computer use in astronomy continues to increase, and so also its impact on the environment. To minimize the effects, astronomers should avoid interpreted scripting languages such as Python, and favor the optimal use of energy-efficient workstations.

ArXiv:2009.11295; Nature Astronomy vol.4, 819 (2020)

Carbon footprint of astronomy and computing

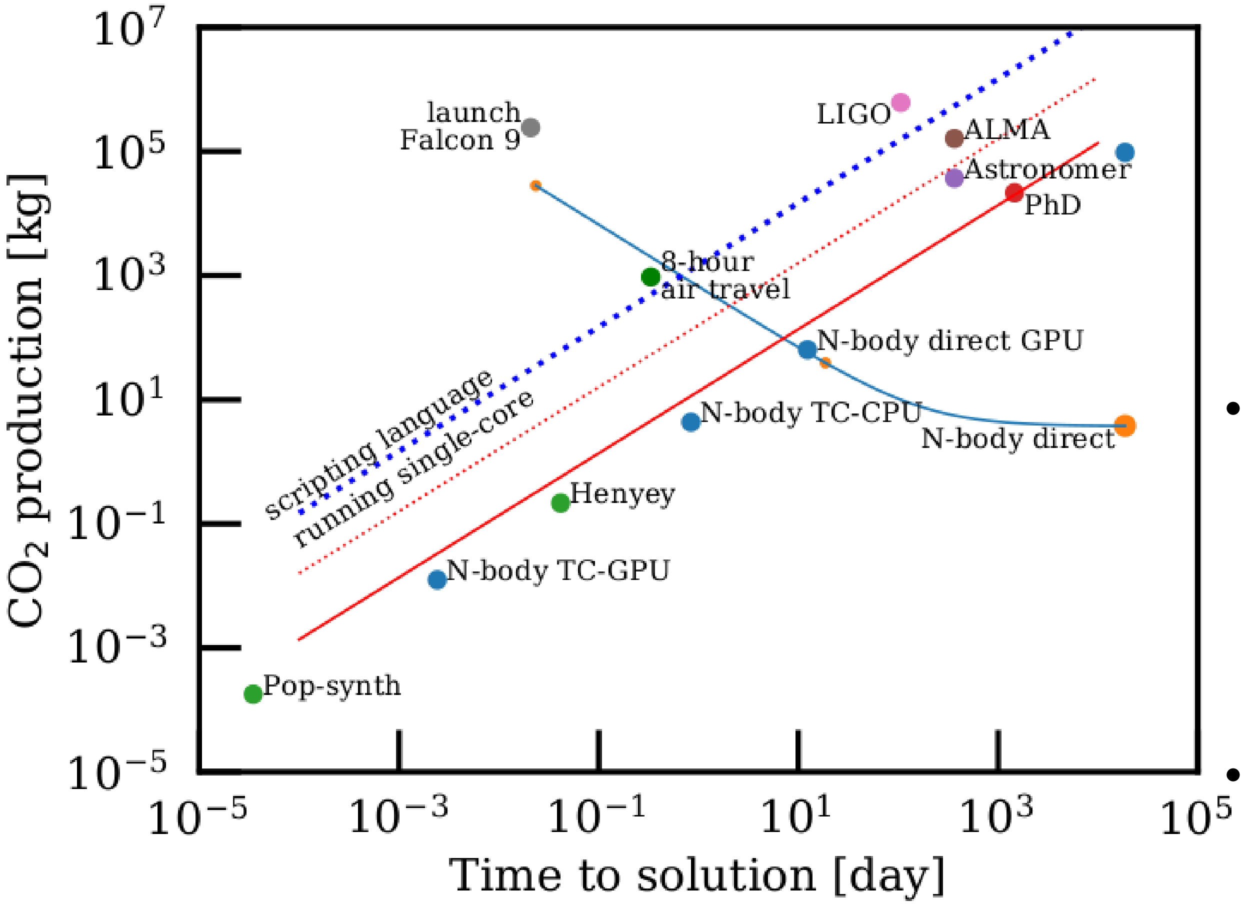

Figure 1: **CO^2 emission (in kg) as a function of the time to solution (in days) for a variety of popular computational techniques employed in astrophysics, and other activities common among astronomers [3] [4]. The solid red curve gives the current individual world-average production, whereas the dotted curves give the maximum country average. The LIGO carbon production is taken over its first 106-day run (using $\sim$ 180 kW) [5], and for ALMA a 1-year average [6]. A Falcon 9 launch lasts about 32 minutes during which $\sim$ 110 000 liters of highly refined kerosene is burned. The tree-code running on GPU is performed using $N = 2^{20}$ particles. The direct N-body code on CPU (right-most blue bullet) was run with $N = 2^{13}$ [7], and the other codes with $N = 2^{16}$. All performance results were scaled to $N = 2^{20}$ particles. The calculations were performed for 10 N-body time units [8]. The energy consumption was computed using the scaling relations of [9] and a conversion of KWh to Co_2 of 0.283 kWh/kg. The blue dotted curve shows the estimated carbon emission when these calculations would have been implemented in Python running on a single core. The solid blue curve to the right, starting with the orange bullet shows how the performance and carbon production changes while increasing the number of compute cores from 1 to 10^6 (out of a total of 7 299 072, left-most orange point) using the performance model by [10].**

- Comparison of the average Human production of CO_2 (red line) with other activities, such as telescope operation, the emission of an average astronomer, and finishing a (four year) PhD.
- The emission of carbon while running a workstation is comparable to the world's per-capita average.

Carbon footprint of computing

- The relation between the time-to-solution and the carbon footprint of the calculations is not linear. When running a single core, a supercomputer-used to capacity-produces less carbon than a workstation. More cores result in better performance, at the cost of producing more carbon.

- Similar **performance as a single GPU is reached when running 1000 cores, but** when the number of cores is further increased, the performance continues to grow at an enormous cost in carbon production.

- When **running a million cores, the emission by supercomputer** by far exceeds air travel and **approaches the carbon footprint of launching a rocket into space**.

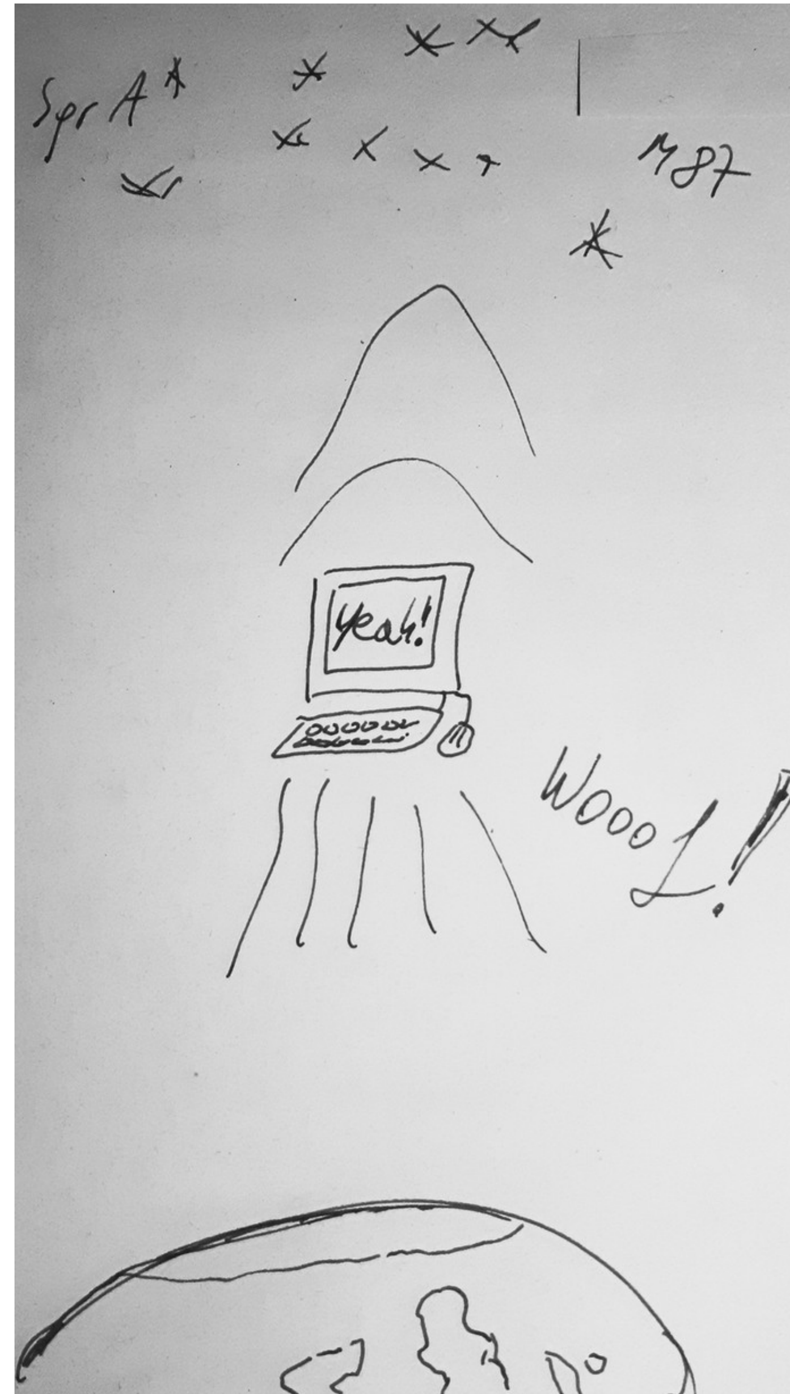

Optimal number of cores

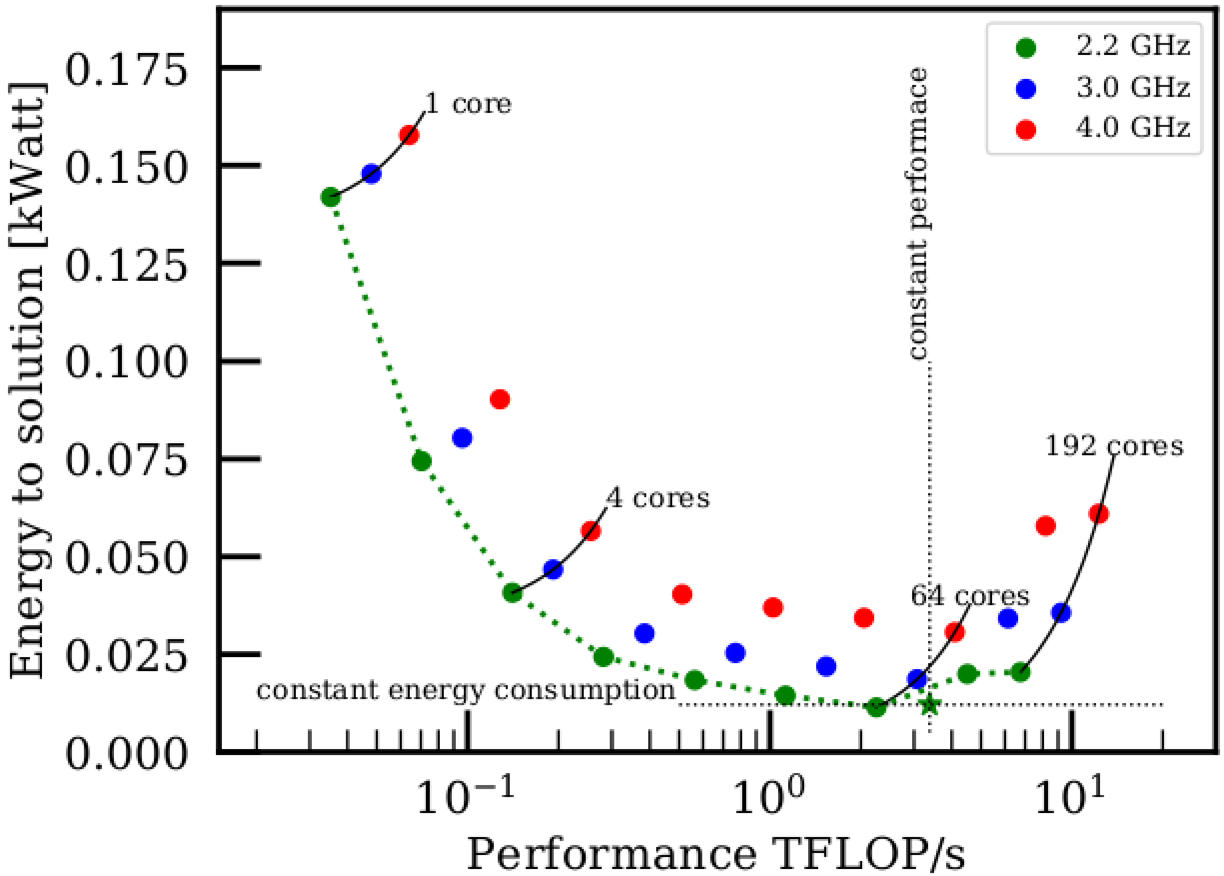

Figure 2: **Energy to solution as a function of code performance. The Z-plot gives for a number of processor (and processor frequencies) and the energy consumed (in kWatt) as a function of performance (in TFLOP/s) [9]. The runs (green dots) were performed using a quad CPU 24-core (48 hyperthreaded) Intel Xeon E7-8890 v4 at 2.20 GHz calculated with 1, 2, 4, ..., 192 cores. Curves of constant core-count are indicated for 1, 4, 64 and 192 cores (solid curves). The other colored points (blue and red) give the relation for overclocking the processor to 3 and 4 GHz, scaled from the measured points using over-clocking emission relations [17]. Dotted curves give constant energy-requirement-to-solution (horizontal) and sustained processor performance (vertical). The star at the cross of these two curves is measured using 96 cores. The calculations are performed Bulirsch-Stoer algorithm with a Leofrog integration [18] at a tolerance of** $\log(dE/E) = -8$ **using a wordlength of 64 bits.**

- Energy consumption as a function of the performance of 96 cores (192 hyperthreaded) workstation-in a private, not shared, use.
- Running single core @ workstation is inefficiently slow and produces more carbon than running multi-core. Performance continues to increase with core count, but optimal energy consumption is reached when running 64 and 96 physical cores (green star in Fig.). Running more cores will continue to reduce the time-to-solution, but at higher emission.
- Not shown here: reducing clock-speed slows down the computer while increasing the energy requirement.

Ecological impact of computing language

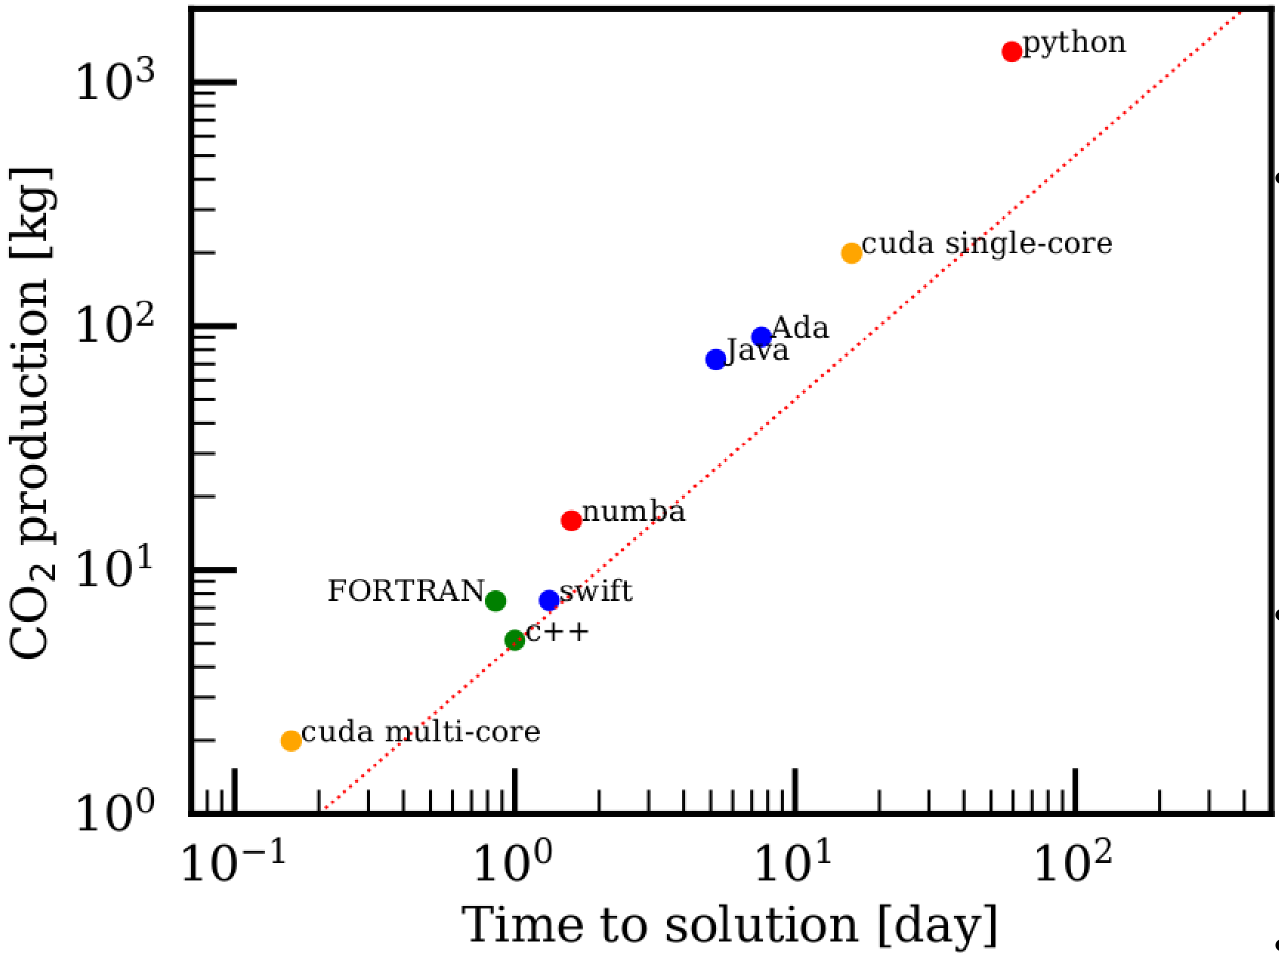

Figure 3: **Here we used the direct N-body code from [23] to measure execution speed and the relative energy efficiency for each programming language from table 3 of [22]. The dotted red curve gives a linear relation between the time-to-solution and carbon footprint ($\sim$ 5 kg CO_2/day). The calculations were performed on a 2.7GHz Intel Xeon E-2176M CPU and NVIDIA Tesla P100 GPU.**

- Results were obtained with the assumption that astrophysicists invest in full code optimization that uses the hardware optimally.
- In practice, most effort is generally invested into solving the research question; designing, writing, and running the code is not the primary concern, if the result is obtained reasonably fast. This is why **inefficient** (and **slow**) scripting languages as Python flourish.
- According to the Astronomical Source Code Library, ∼ 43% of the code is written in Python, 7 % Java, IDL and Mathematica. Only 18%, 17% and 16% of codes are written in Fortran, C and C++ respectively.
- Python and Java are also less efficient in terms of energy per operation than compiled languages, which explains the offset away from the dotted curve.
- Among 27 tested languages, only Perl and Lua are slower than Python- **popularity of Python should be confronted with the ecological consequences**.

How to improve?

- Runtime performance of Python can be improved **using numba or NumPy libraries**, which offer pre-compiled code for common operations-it leads to an enormous increase in speed and reduced carbon emission. However, these libraries are rarely adopted for reducing carbon emission or runtime with more than an order of magnitude.
- NumPy, for example, is mostly used for its advanced array handling and support functions. Using these will reduce runtime and, therefore, also carbon emission, but **optimization** is generally stopped as soon as the calculation runs within an unconsciously determined reasonable amount of time, such as the coffee-refill time-scale or a holiday weekend. We even teach Python to students, but without realizing the ecological impact.
- The carbon footprint of computational astrophysics can be reduced substantially by **running on GPUs,** but the development time of such code requires major investments in time and expertise.
- As an alternative, one could run concurrently **using multiple cores, rather than a single thread**. It is even better to **port the code to a supercomputer and share the resources**.
- Best for the environment is to **abandon Python** for a more environmentally friendly (compiled) programming language.
- Even better is to **use other interesting strongly-typed languages** with characteristics similar to Python, such as Alice, Julia, Rust, and Swift. They offer the flexibility of Python but with the performance of compiled C++.
- Educators may want to **reconsider teaching Python** to University students. There are plenty environmentally friendly alternatives.

An astronomical institute's perspective on meeting the challenges of the climate crisis

Knud Jahnke[1,*,+], Christian Fendt[1,+], Morgan Fouesneau[1,+], Iskren Georgiev[1,+], Tom Herbst[1,+], Melanie Kaasinen[1,+], Diana Kossakowski[1,+], Jan Rybizki[1,+], Martin Schlecker[1,+], Gregor Seidel[1,+], Thomas Henning[1], Laura Kreidberg[1], and Hans-Walter Rix[1]

[1]Max Planck Institute for Astronomy, Heidelberg, Germany
[*]jahnke@mpia.de
[+]these authors contributed equally to this work

ABSTRACT

Analysing greenhouse gas emissions of an astronomical institute is a first step in reducing its environmental impact. Here, we break down the emissions of the Max Planck Institute for Astronomy in Heidelberg and propose measures for reductions.

Energy and emissions cost of MPIA in 2018

Source	Amount	CO_2eq	CO_2eq/researcher	Percentage (%)
Travel (air)	1030 flights	1280 t	8.5 t	47
Electricity (on/off campus)	3,400,000 kWh	779 t	5.2 t	29
Heating (oil)	150,000 l	446 t	3.0 t	16
Commuting (car)	792,000 km	139 t	0.9 t	5
Paper / cardboard	0.15 / 7 t	35 t	0.2 t	1
Computer (desk-/laptops)	57 purchased	29 t	0.2 t	1
Meat (canteen)	1000 kg	16 t	0.1 t	<1
Total		**~2720**	**18.1 t**	100 %

Table 1. Summary of the MPIA's GHG emissions in 2018. Note that electricity includes both consumption at the MPIA campus, as well as in external supercomputing centers used by MPIA.

- GHG=greenhouse gas
- MPIA Heidelberg: ~150 researchers and ~320 employees in total
- emissions in seven categories; business flights, commuting, electricity, heating, computer purchases, paper use, and cafeteria meat consumption.
- Not all contributions could be measured easily, but authors estimate the major contributors to emissions, that is, flying and electricity, to be accurate to within 20%.
- The MPIA's total GHG emissions for 2018 amount to 18.1 tCO_2 eq per researcher. Alternatively, the contribution per refereed science publication, of which there were 583 either authored or co-authored by MPIA astronomers in 2018, is 4.6 tCO_2 eq.
- This is an alarming ~3 times higher than the German target for 2030 with the Paris Agreement.

Comparison with Australian (!) researchers in 2018

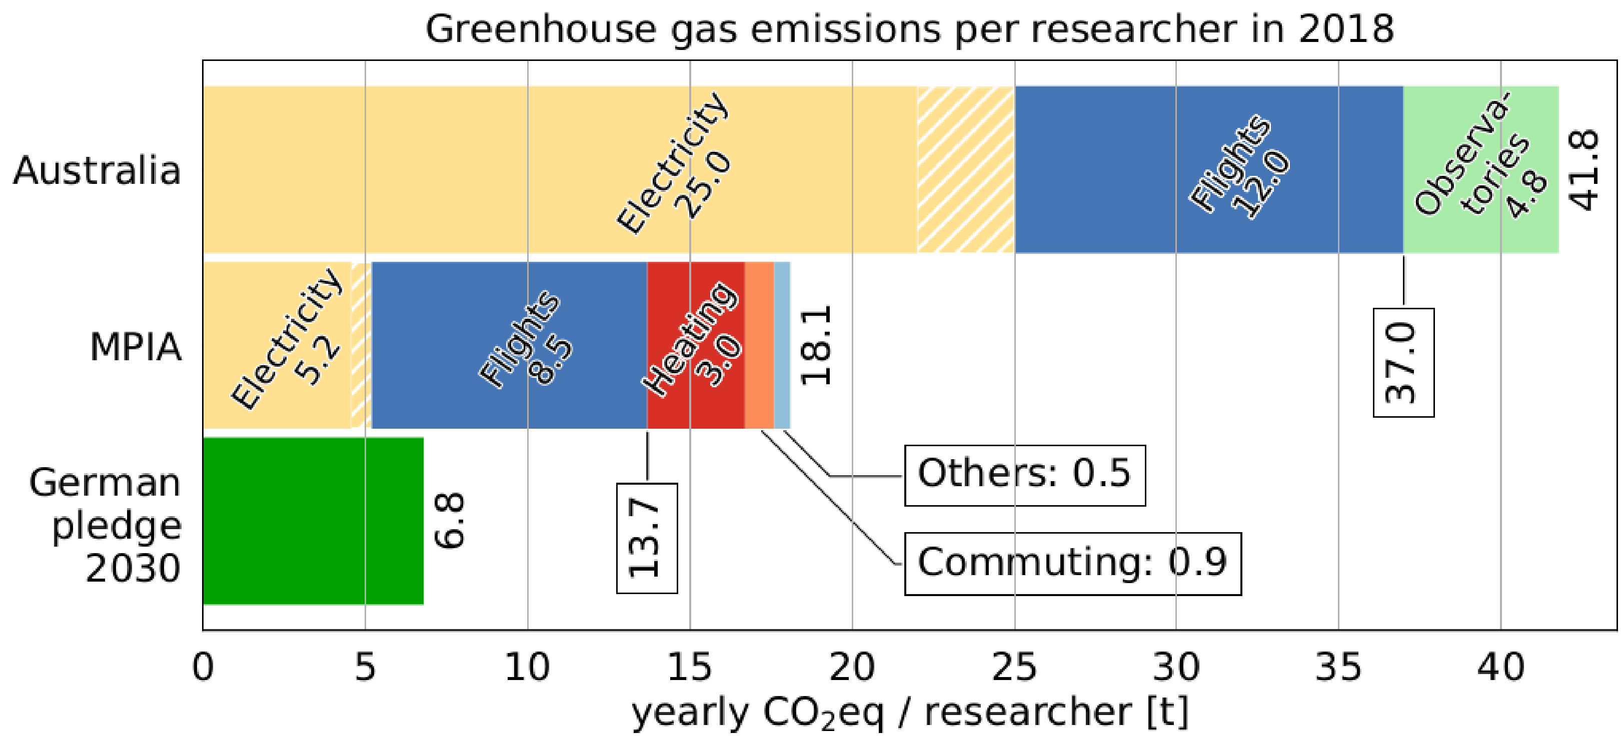

Figure 1. Average annual emissions in 2018 for an Australian and MPIA researcher in tCO_2eq/yr, broken down by sources. The sources include electricity, flights (converted to the same emission model, see text), observatory operation, office heating, commuting, and 'others', a category that combines office desktop and laptop hardware, paper and cardboard use, and meat consumption. Electricity related emissions include both computing and non-computing consumption, where for Australia computing is accounting for 88% of electricity emissions; we estimate a similar fraction for MPIA. In the plot, the smaller hatched part of the 'Electricity' bar indicates non-computing electrical power. Observatory operation is only given for Australia, while heating, commuting, and sources captured by the 'others' category are only given for MPIA. Therefore, emissions can only be compared between Australia and MPIA for electric power consumption and flights, which amount to 37.0 and 13.7 tCO_2eq/yr for Australian and MPIA researchers respectively. The major difference lies in the amount of GHG emissions per kWh electricity, which differs by a factor of ∼4 between the Australian astronomy and the MPIA. These values do not account for all emissions per capita. In particular, emissions not related to work are excluded. The combined MPIA emissions of 18.1 tCO_2eq/yr and researcher are also compared to the German pledge of a 55% reduction of the 1990 emissions by 2030, plotted per capita in dark green, which is close to 6.8 tCO_2eq/capita per year[9–11].

How to improve?

- The difference to Australia in electricity-related emissions is almost completely due to the different carbon intensity for electricity production: Whereas fossil fuel sources contributed 83% to Australia's generation of electricity in 2018, the contribution in Germany was ∼47% , and MPIA's delivery contracts have a carbon intensity even substantially below that.
- The high carbon intensity of MPIA's heating; needs to be addressed at the institute level.
- Flight-related GHG emissions dominate the MPIA's total emissions:

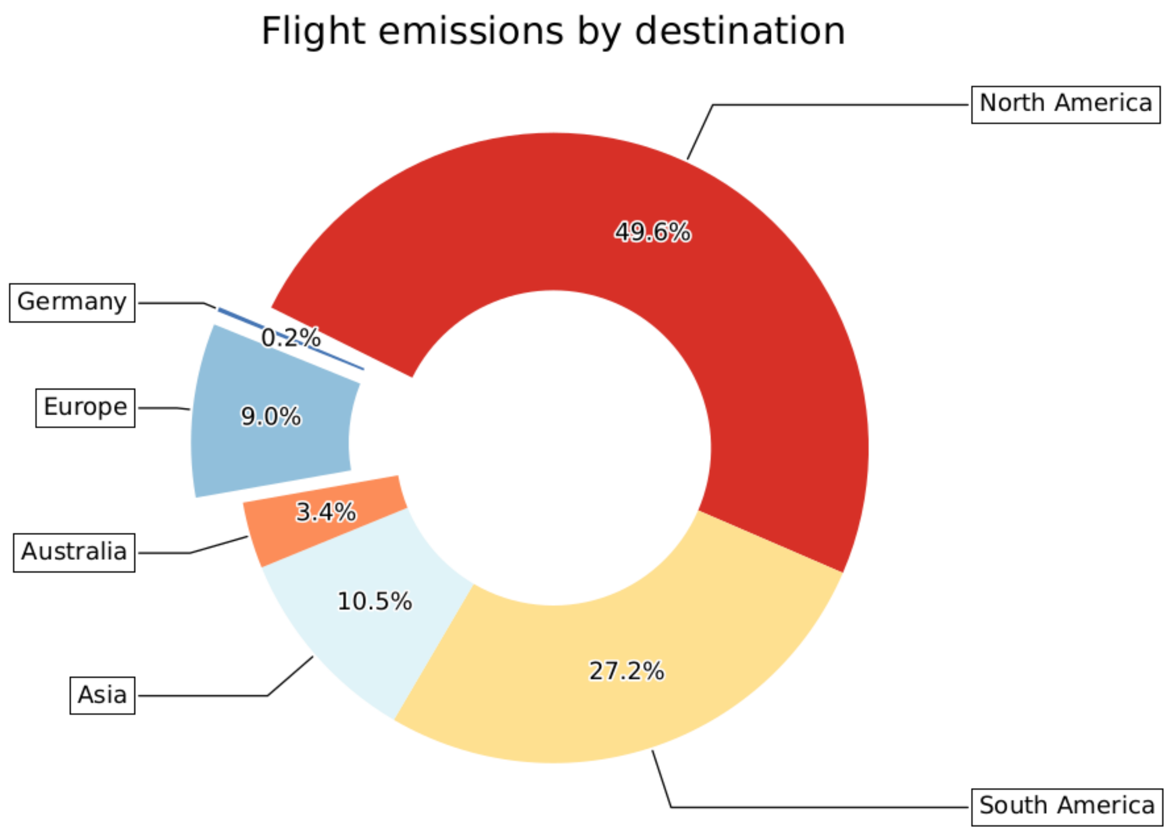

Figure 2. Relative GHG emissions broken down by flight destination for MPIA employees. Intercontinental flights that cannot be easily replaced by alternative means of transport make up about 91% of flying emissions. This is due to the number of flights, and the high climate impact of each intercontinental flight, primarily due to distance traversed, but also due to greater time averaged emission altitude, for example for nitrogen oxides.

How to improve?

- Changes to the German public servant's travel law in early 2020 ensure that **train trips** to well-connected European destinations are now **reimbursed, even if they are more expensive than a flight**. Moreover, at the individual level, many German researchers have pledged **not to fly distances under 1000 km**.
- To reduce our carbon footprint to anything approaching net zero by 2050, the expertise in **hosting virtual events** that was so rapidly developed during the 2020, should continue to be applied and expanded.
- Computing resources make 75–90% of our electricity consumption: The **sources of national/regional energy production** are decided at a political level, but the astronomical community, and indeed individual citizens, can collectively campaign for this change.
- Super-computing facilities may be moved to locations where renewables are available and **less electrical energy is needed for cooling**, for example, to Iceland, which has an average of 0.028 kgCO 2 eq/kWh emission for produced electricity in August 2020.
- Potential idle times, and hence the required amount of hardware, could be reduced by **switching to more cloud computing**, because there, capacity utilization is generally higher than for local computers. As a community, we should guarantee an **efficient use of super-computing** resources. This applies both to code efficiency, as well as regarding the computing architecture that we build up or rent.
- **Heating of** MPIA's **buildings** is the third largest contributor. Instead of oil, use of ground heat, in combination with an electrically-operated heat pump should be done.
- We proposed a **photovoltaic installation on** MPIA's **roof**, also currently under review, which would initially produce ∼10% of MPIA's on-campus electricity consumption at zero additional cost.

Conclusions

-Pay attention to **more efficient ways to compute**. Think non-Python.

-**Less flights**. More local collaborations, with train connections.

-**Participate** in political decisions about energy production.

-Think about energetically **more efficient ways to power/heat** the institute.

Shanghai sky on a good or bad day-AQI 30 or 200

Thank you.

Stay safe.
Stay sane.
Do Science!

Reminder: Nature is like this