Sustainability in HEP Computing

Rodney Walker

Ludwig-Maximillians-Universität, Munich, Germany

Abstract. The European energy crisis during 2022 prompted many computing facilities to take urgent electricity saving measures. In part voluntary, in response to EU and national appeals, but also to keep within a flat energy budget, as the electricity price rose. We review some of these measures and, as the situation normalizes, take a longer view on how the flexibility of high throughput computing can best contribute to the net-zero energy transition. We discuss carbon intensity estimation of workloads, automatically powering down nodes during low-demand periods, CPU frequency modulation to track variable electricity tariffs or renewable generation throughout the day and potential spin-down disk storage solutions.

1 Introduction

Due to the reduction and eventual stop of Russian gas deliveries to Europe in 2022, there was genuine concern that the we might run out of gas for industry, heating and electricity generation. A range of measures were taken, at the European and national levels, to boost alternative supply and reduce consumption. This included a call for a 10% reduction in electricity consumption, and a 15% reduction during peak periods. We describe some of the measure taken in the scientific computing area.

The lasting challenge of achieving net zero CO^2 requires long-term measures regarding our electricity consumption. The interplay of intermittent renewable generation, lack of storage capacity and variable demand makes the time of consumption much more important than the overall volume. We discuss the motivation to modulate our power consumption, and the methods available to High Throughput Computing(HTC).

We briefly discuss the so-called embedded carbon of our hardware, and its importance as the proportion of green electricity increases.

2 Energy saving measures

The overall energy consumption can be reduced by optimized data center design and hardware choice, but these are not things we can change in the short term. In terms of operating existing equipment, the energy saving measures are limited.

2.1 Power down

The simplest way to reduce consumption is to power down servers, starting with the least energy efficient. WLCG[1] compute centers typically continue to run older, out-of-warranty hardware until such a time that it fails, or the space, power or cooling capacity is needed for new hardware.

A running compute node has many batch jobs of different duration, typically up to 4 days. This leads to a lengthy draining process , where the machine is partially loaded, before it can be powered down. The majority of workloads from all users are not preemptable, so jobs cannot be killed without loss of work. For this reason, draining and powering down of nodes is only useful for longer, and infrequent periods.

As a voluntary reaction to the energy crisis, and a proof of principal, DESY Hamburg powered down older worker nodes over the Christmas holiday 2022[2]. As can be seen from Fig.1[2], the older hardware has significantly less CPU performance per Watt of power. In particular there is a step improvement in going from Intel to AMD CPUs.

Powering down the out-of-warranty nodes, to the left of the red bar, had the effect of reducing the power consumption by 60% for only a 40% reduction in the compute throughput. A total of 35MWh of electricity was saved during the 20 day period.
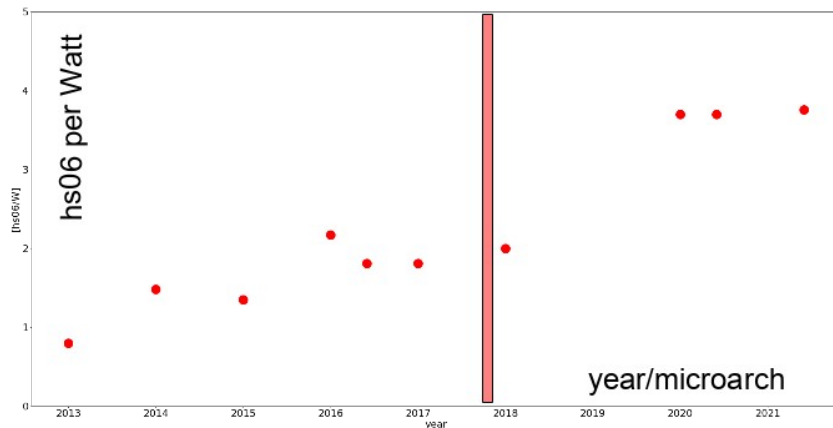

Figure 1: A benchmark of the compute performance, HEPSpec06, for various CPU models in use at DESY, versus the year purchased. The 3 data points on the right are for AMD CPU.

2.2 CPU frequency modulation

Modern CPUs are designed to significantly reduce energy consumption when not actively needed for processing. This ability is critical for battery-operated devices, to extend the on-battery operation time, but is also standard for server CPUs. In the case of a fully-loaded CPU, the frequency can still be forcibly reduced, to reduce energy consumption, at the expense of the payload running more slowly.

Since the CPU power consumption is proportional to the frequency and the square of the voltage[3],

$$P \sim fV^2 \qquad (1)$$

it is advantageous to also reduce the voltage. This is possible because a lower voltage is sufficient at lower frequencies – roughly speaking, things don't need to be pushed around so quickly. Dynamic Voltage and Frequency Scaling(DVFS) ensures a faster, non-linear reduction in power consumption as the frequency is reduced. Fig.2 [3] shows the theoretical CPU power consumption versus frequency, assuming a linear dependence of the voltage with frequency. In practice, there are step-wise changes of voltages with frequency.
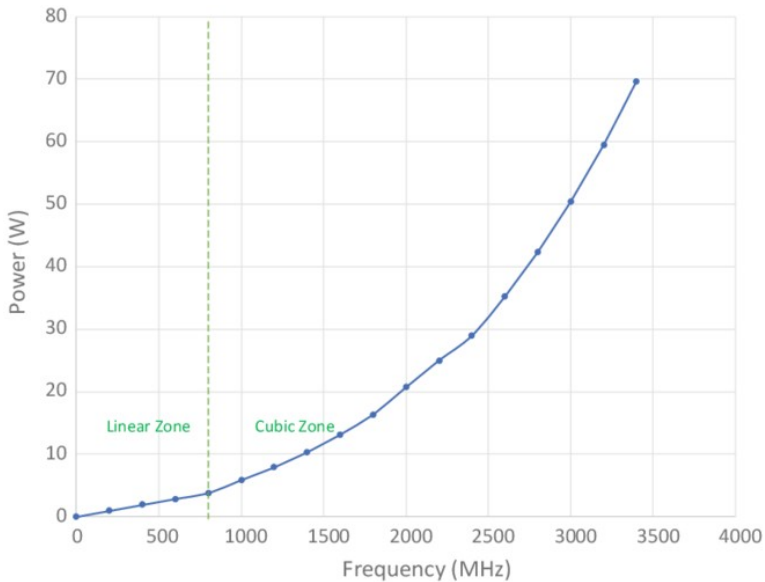

Figure 2: Theoretical CPU power versus frequency[3]

There is a base consumption of the server which does not reduce with CPU frequency. The power supply unit, disk, RAM and parts of the CPU all use power regardless. One concern might therefore be that throttling the frequency would result in a lower efficiency, in terms of processing work done per kWh of electricity. However, the non-linear reduction of power with frequency, due to the V^2 term, has the potential to offset the effect of the base consumption. Indeed, preliminary measurements[2,4] indicate that the optimal efficiency lies below the nominal frequency and that the work/kWh is not significantly lower even at the lowest frequency.

A financially motivated energy saving action on the Karolina HPC[5] reduced the compute node CPU frequency from 3300MHz to 2100MHz. The most arithmetically intensive loads run 16% longer using 30% less power.

2.3 Data Storage

It is estimated that some 40% of an WLCG data center electricity is consumed by disk storage. The typical configuration of PB scale storage is RAID6, where all disks are spinning continuously. Only a small part of the data is accessed by the local compute cluster, and only a fraction of the i/o bandwidth used. This raises the question whether some of the storage could operate with a lower quality of service, and thereby lower energy consumption. In such a scenario, most disk would be idle and spun-down. Access to data would be scheduled in a similar way to tape storage, where a spinning disk is analogous to a mounted tape.

Of course this has implications for the data placement, redundancy and operation, and would require more investigation.

3 Electricity grid carbon intensity

The amount of CO_2 emitted per kWh of electricity consumed depends on the energy mix in generation at that time. The intermittent nature of renewable power, solar and wind, and the variation in demand, leads to different amounts of generation from various fossil fuels being required. This is the case for Germany and most countries, although there are exceptions, such as the Nordics where hydroelectric power plays a dominant role.

Fig.3 [6] shows the generation mix in Germany over a week in May/June 2023. Generation from gas and coal is modulated to cover the load as the renewable component varies. The corresponding carbon intensity is estimated[7] by assigning a gCO_2/kWh to each generation type, and calculating the weighted average, as shown in Fig.4. There is a clear anti-correlation between carbon intensity and the proportion of renewable energy in the mix. It ranges from $200gCO_2/kWh$, when the sun is shining, to $400gCO_2/kWh$ overnight. Of course this will look quite different in the winter months where wind is the dominant renewable component.
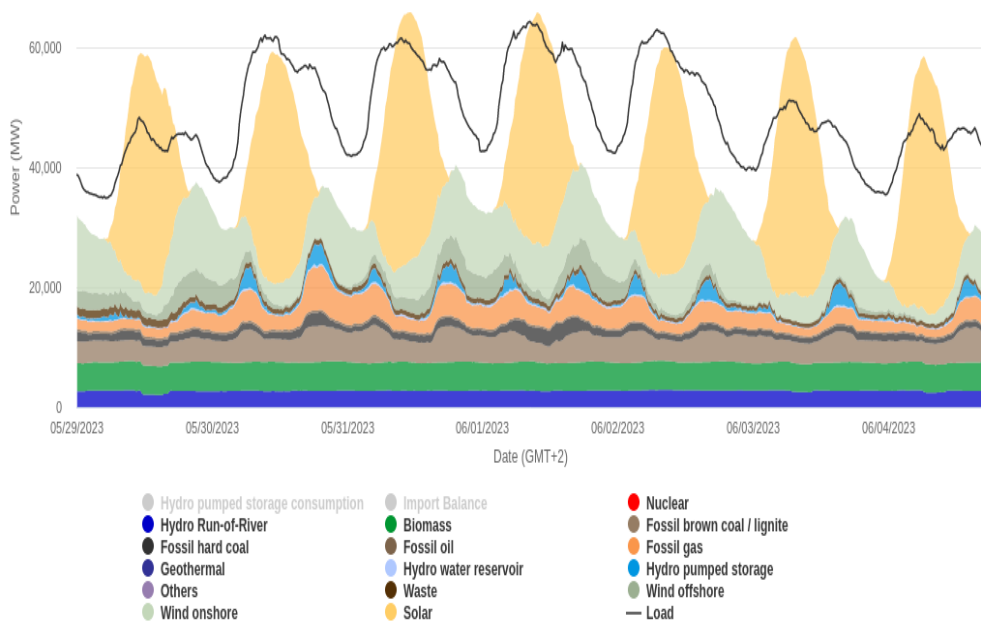

Figure 3: Electricity generation mix and load in Germany for Week 22 2023
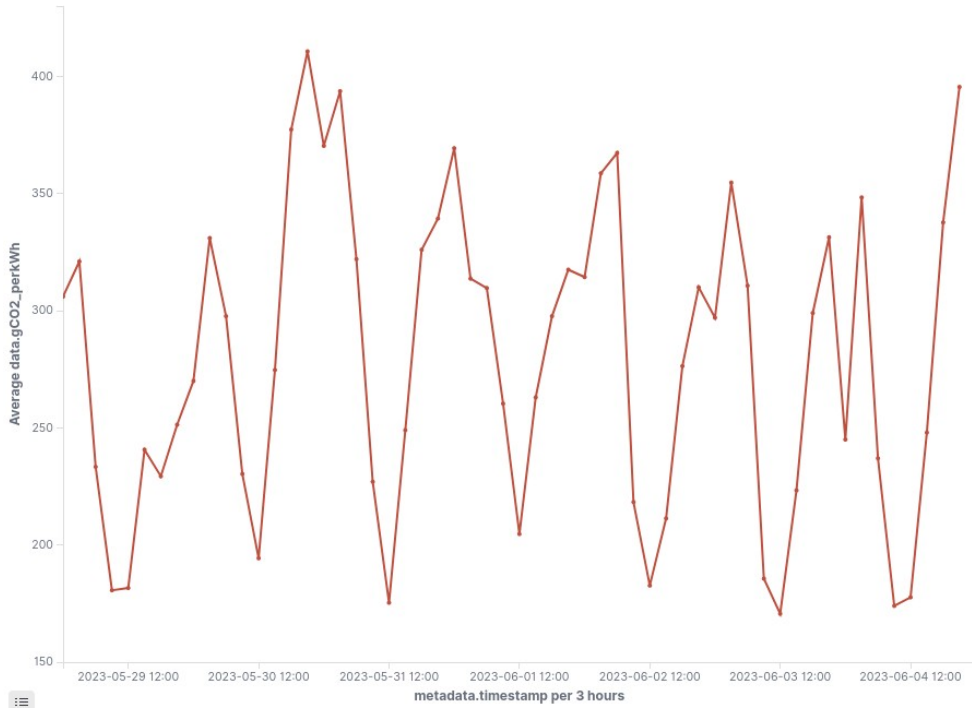
Figure 4: Estimated carbon intensity of German electricity for the week shown in Fig.3.

To demonstrate how the future electricity generation might look, in the same week of 2030, we scaled the solar, offshore wind and onshore wind by factors 3, 4 and 2 respectively. These represent the German government goals for 2030. The demand is expected to rise due to battery electric vehicle(BEV) charging and heat pumps, but only by 11% according to a study[8] commissioned by the German government.

The resulting prognosis, excluding fossil fuel generation, is shown in Fig.5. One can see extended periods where the demand is entirely satisfied by renewables alone. However there are still times with a deficit of power where extra generation, potentially from fossil fuels, would be needed. The striking feature is the large excess of power due to solar generation. The ability to redistribute the demand from the periods with a deficit to those with an excess will be crucial in reducing the fossil fuel dependency. For example, if 7 Million BEVs were to charge during the solar peak, then the 70GW excess would be used, reducing demand when less renewable energy is available.
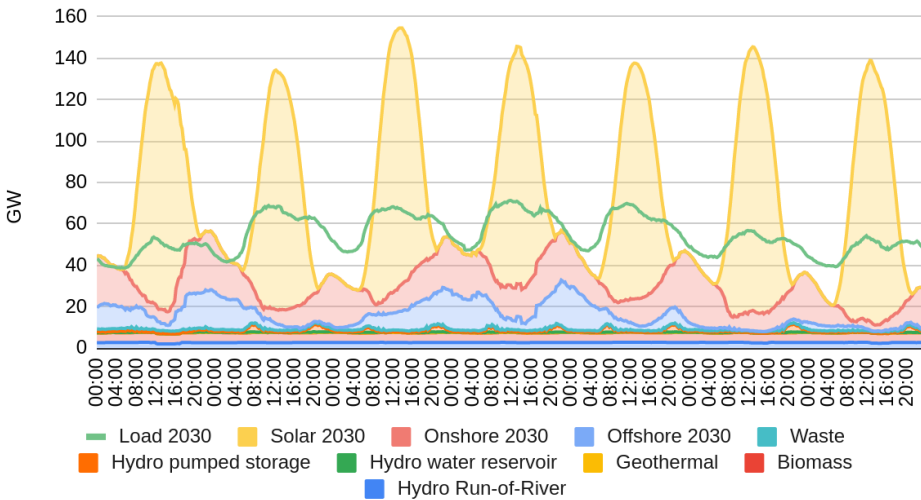
DE generation mix 2030 Week 22

Figure 5: Week 22 2023 scaled by German 2030 renewable targets for wind and solar generation capacities. The demand is also increased by 11%, the prognosis allowing for BEV charging, domestic heat pumps and efficiency improvements.

4 Price

The previous section shows that the CO_2 emissions can be reduced by reducing consumption at the times when the electricity grid carbon intensity is high. As publicly funded scientists we cannot altruistically reduce our carbon footprint in this way, in the long term. Fortunately electricity from renewable generation is also the cheapest form, and this is reflected in the market price. The market price per kWh for one sunny day in the aforementioned week is shown in Fig.6[9]. One can see the range from 0 to 14ct/kWh with the lowest price corresponding to the lowest carbon intensity. Although there are many other factors influencing the market price, the correlation to carbon intensity is noticeable and is likely to become stronger in future. An increase in the EU Carbon price, together with more renewables, should make the high CO_2 intensity periods less frequent but more expensive.
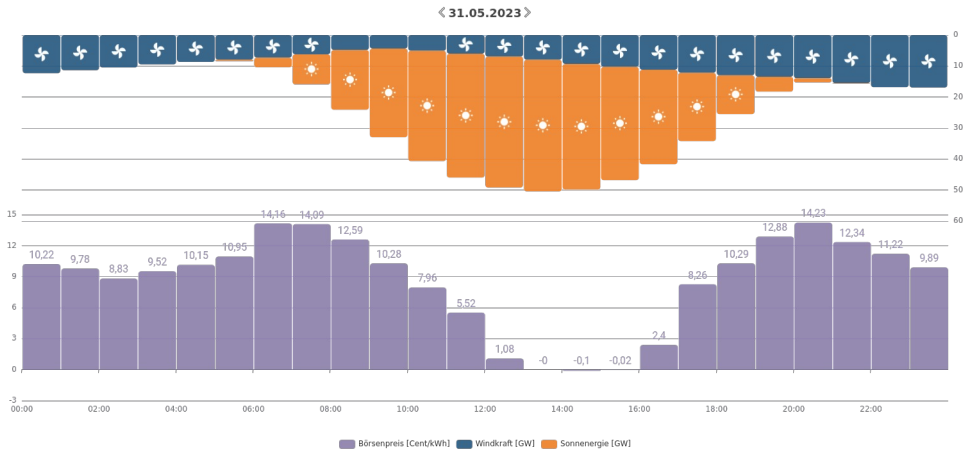
Figure 6: Market electricity price in Germany on 31.05.23. Graphic from Awattar[9]

Due to limited storage capacity, the ability to vary demand is crucial as we learn to live without fossil fuel generation. This means smoothing the demand peaks, and having consumption follow the renewable generation. Consumers unable to react to the strong price signal will pay more. Even though scientific computing is not a large consumer, we will anyway need to modulate consumption for purely financial reasons. In this sense, we may as well lead the way and demonstrate the kind of flexible demand which must become widespread.

5 Embedded carbon

Various studies into the manufacture of computer servers show this to be a significant part of the cradle-to-grave carbon footprint. A study from Dell[10] suggests that up to 50% of the carbon footprint is embedded, i.e. not emitted during operation. Of course there are many assumptions regarding the manufacturing processes, server cpu load and electricity grid carbon intensity. However it is clear that as the electricity for operation becomes cleaner, the proportion of the footprint due to manufacture will grow.

The natural consequence is that we should operate hardware longer, and this despite the poorer energy efficiency of older hardware. Since currently the grid carbon intensity is not zero but varies in time, we should only operate the inefficient hardware during periods of low carbon intensity.

6 Conclusion

The major part of WLCG scientific computing is HTC, and not time critical. It has the flexibility to reduce capacity for hours or days. We discussed two methods to reduce cpu power consumption: powering down compute nodes, and reducing the CPU frequency. Powering down nodes requires draining, and is therefore suitable for longer periods, such as several dark, windless days in winter. Modulating the CPU frequency can be done instantly and is therefore useful to track hourly changes in renewables and demand.

The planned capacities of renewable generation from solar and wind are high, in order to cover as much of the demand as possible. This inevitably leads to periods where production exceeds demand by a large amount. All current forms of storage are either limited, expensive, or both. The cheapest and most efficient 'storage' is that which can be avoided. Shifting useful loads from periods of high to low net demand, taking the renewables into account, reduces the need for storage.

The price structure will strongly encourage this load shifting so we should modulate consumption for financial reasons. Retaining and operating older compute nodes reduces the lifetime carbon footprint, and increases the capacity for running when there is an excess of electricity. This offsets the reduced capacity when electricity is expensive, and nodes are powered down or frequency throttled.

The next step is for a data center to move to a variable tariff, in order to motivate the development of automated power modulation.

References

1. [Computing for the Large Hadron Collider](), Ian Bird. "Annual Review of Nuclear and Particle Science", Vol. 61: 99-118, November 2011

2. Private communication with T.Hartmann, DESY

3. Energy Efficient High Performance Processors Recent Approaches for Designing Green High Performance Computing, Haj-Yahya, Jawad & Mendelson, Avi & Ben-asher, Yosi & Chattopadhyay, Anupam(2018) 10.1007/978-981-10-8554-3.

4. Energy Efficiency in H.E.P. (ARM vs. x86 and beyond) , E.Simili these proceedings.

5. https://docs.it4i.cz/general/energy/

6. https://www.energy-charts.info

7. https://www.co2signal.com/

8. https://www.prognos.com/de/projekt/entwicklung-des-bruttostromverbrauches-bis-2030

9. https://www.awattar.de/

10. https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf