


Computer Networks

Sustainable Networking



Romain Jacob

nsg.ee.ethz.ch

ETH Zürich

May 17, 2024

Partly based on material by Chris Adams

What does **sustainability** mean for you?

What does **sustainability** mean for you?

Meeting the needs of the present
without compromising the ability of
future generations to meet their own needs.

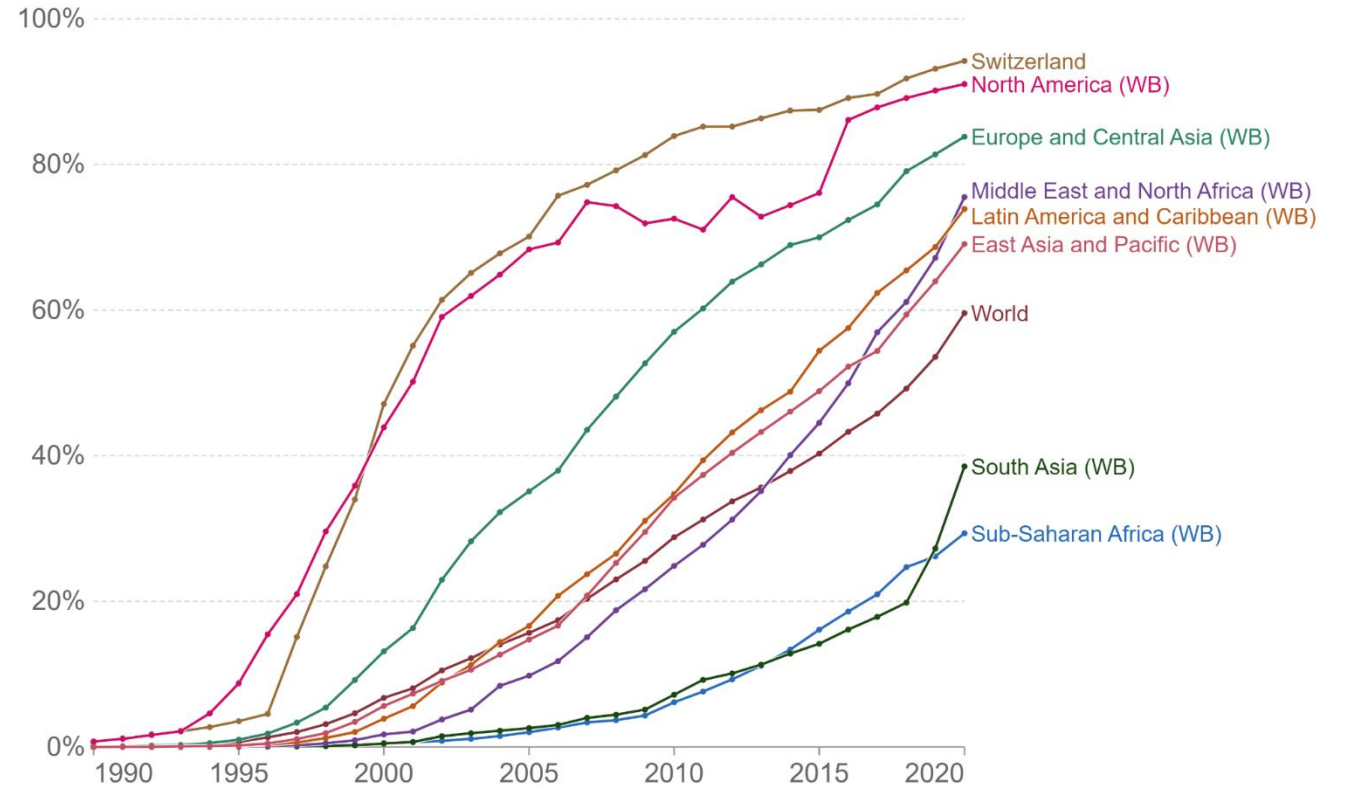
[UN Sustainability definition]

Internet access is still far from universal.

Share of the population using the Internet

Share of the population who used the Internet¹ in the last three months.

Our World in Data



Data source: International Telecommunication Union (via World Bank)

OurWorldInData.org/internet | CC BY

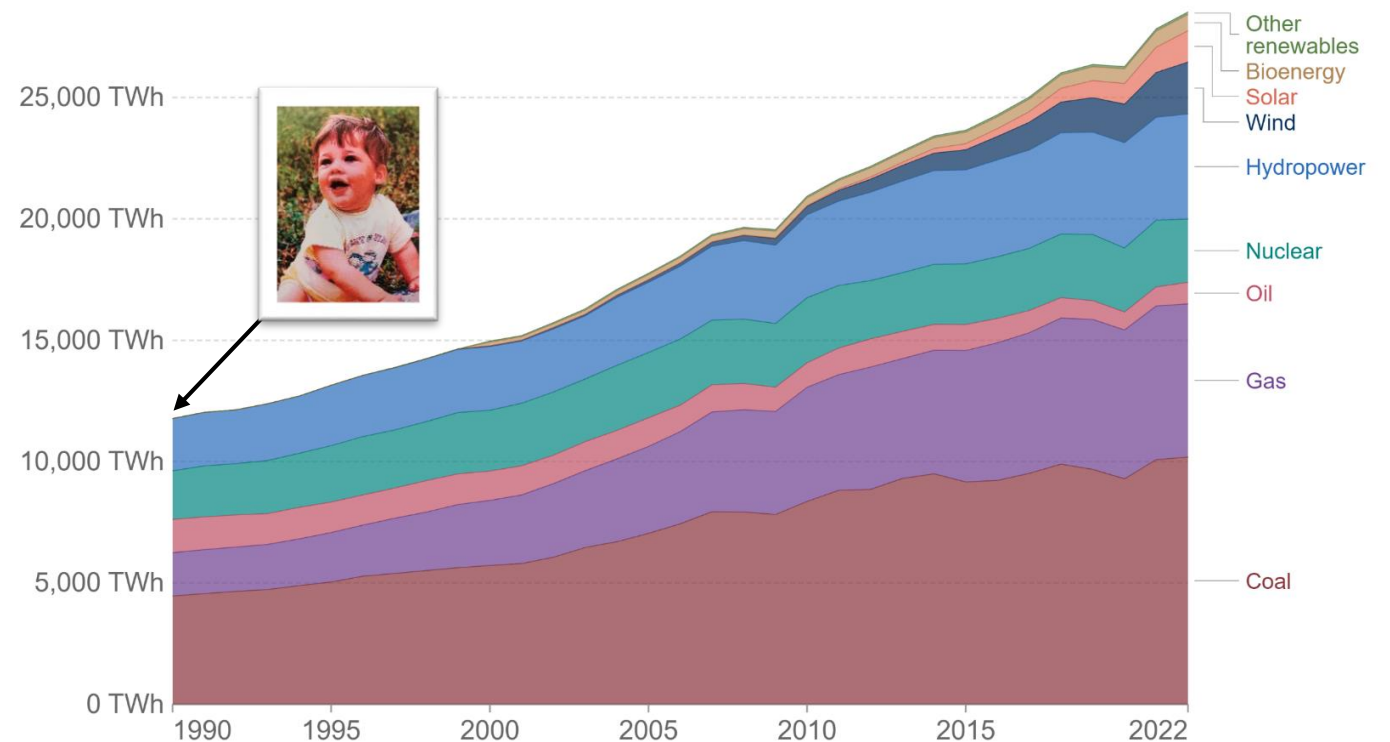
1. Internet user: An internet user is defined by the International Telecommunication Union as anyone who has accessed the internet from any location in the last three months. This can be from any type of device, including a computer, mobile phone, personal digital assistant, games machine, digital TV, and other technological devices.

[Internet users]

Electricity demand grows rapidly.

It **doubled** in my lifetime.

Electricity production by source, World



Data source: Ember's Yearly Electricity Data; Ember's European Electricity Review; Energy Institute Statistical Review of World Energy

Note: 'Other renewables' includes waste, geothermal, wave and tidal.

OurWorldInData.org/energy/ | [CC BY](https://creativecommons.org/licenses/by/4.0/)

[World energy production]

ICT is taking up an increasingly larger share.

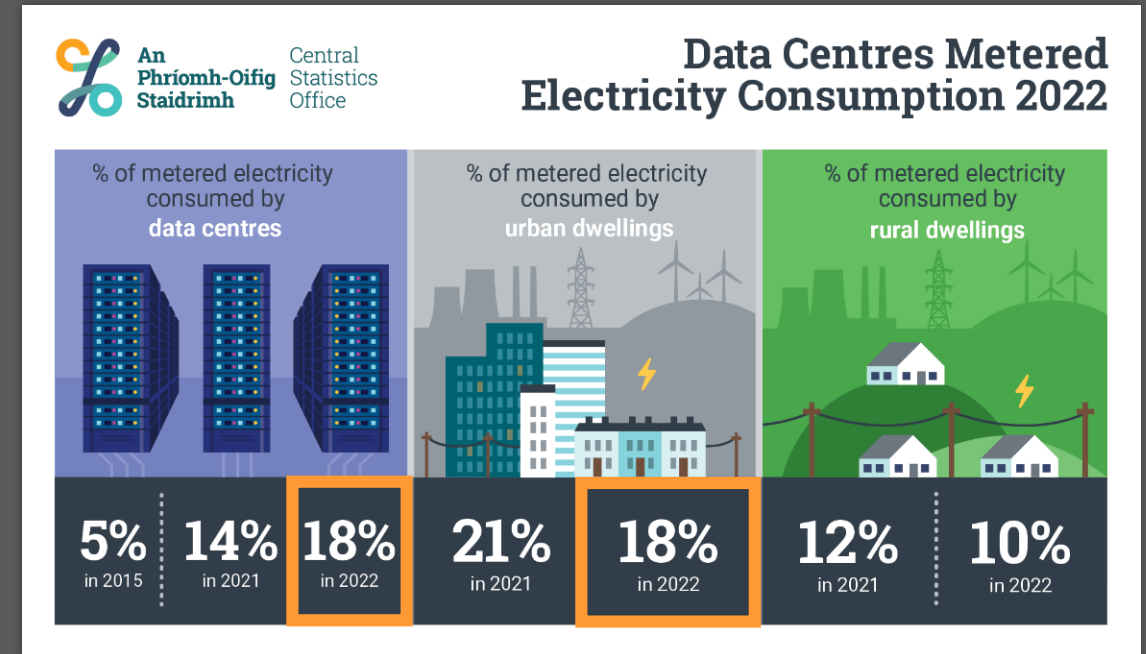
Estimated electricity consumption of

Data centres 240-340 TWh

Comm. networks 260-360 TWh

1-1.3% of global electricity demand, each

[IEA | DC and Telco]



DC energy demand equals all of urban housing Ireland, 2022


[Ireland DC consumption]

ICT is taking up an increasingly larger share. So companies get “creative.”

AWS buys 105 back-up diesel generators for new data center in Dublin

Move in response to concerns over impact on electricity networks in the region

October 25, 2022 By: Georgia Butler [Have your say](#)



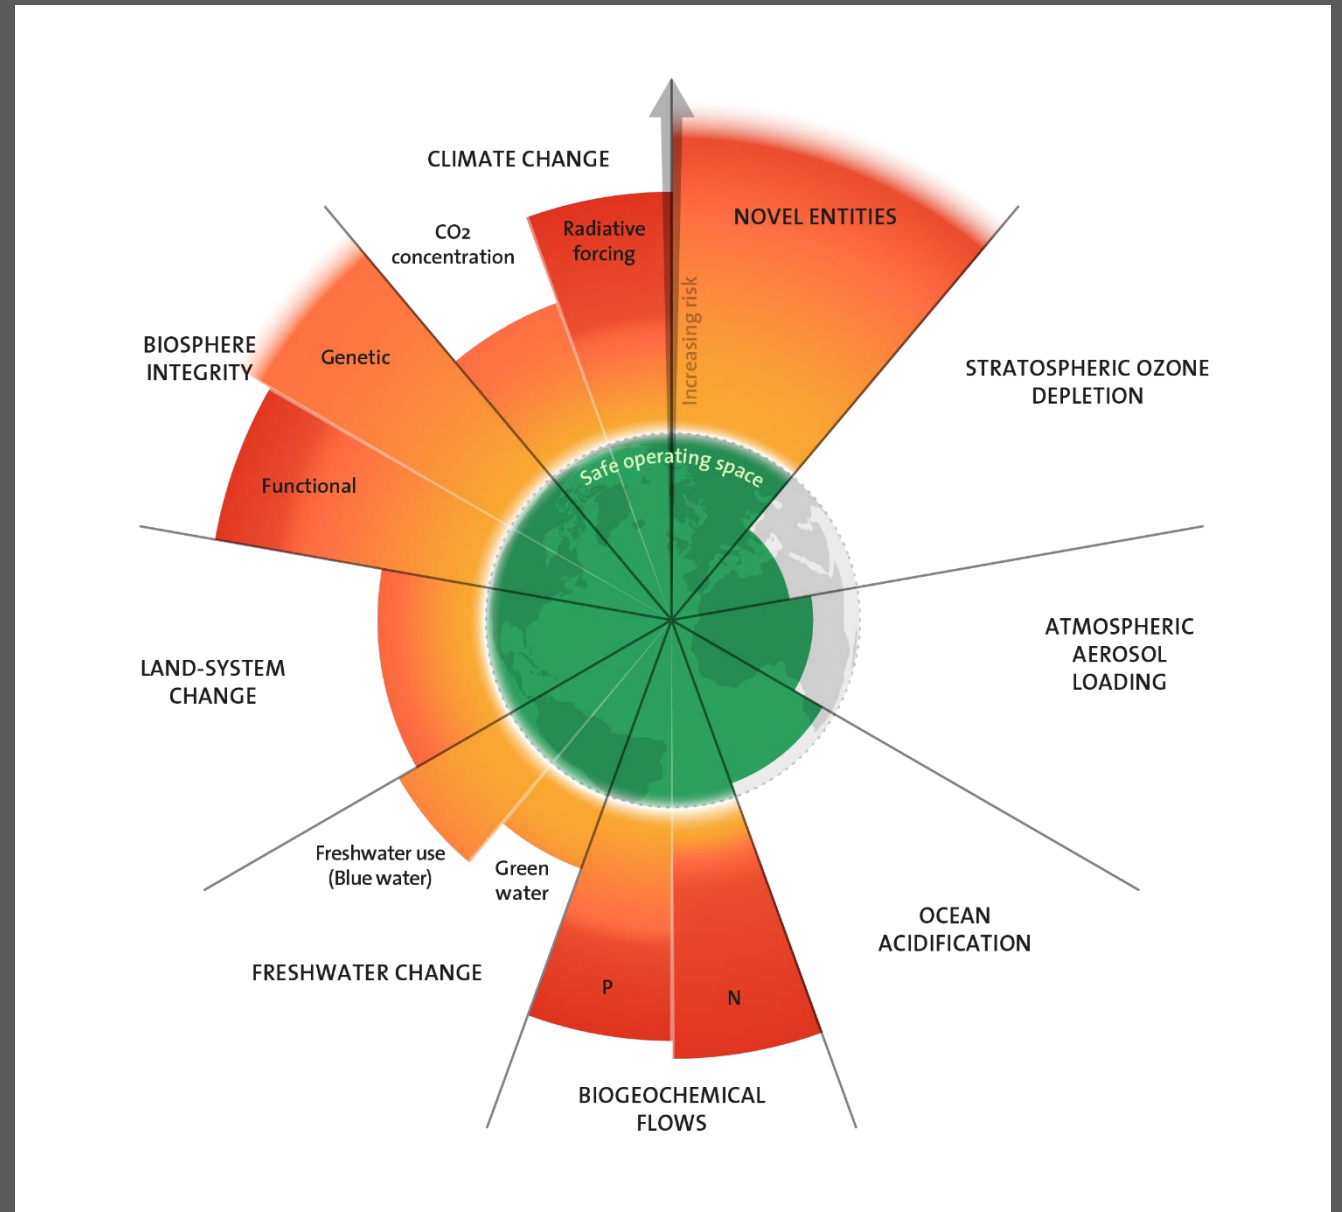
DC energy demand equals all of urban housing Ireland, 2022

[AWS buys diesel gen]

[Ireland DC consumption]

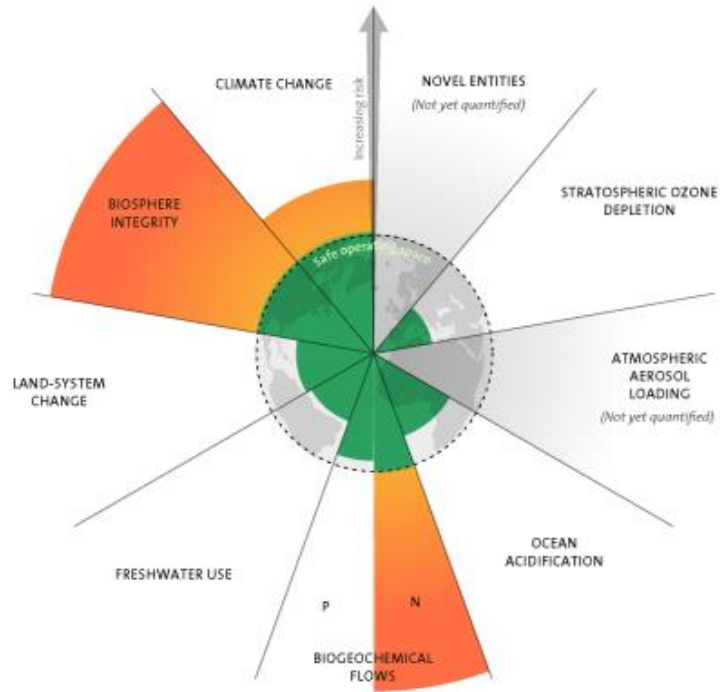
Sustainability is **not only** about energy!

Even if this lecture focuses on energy



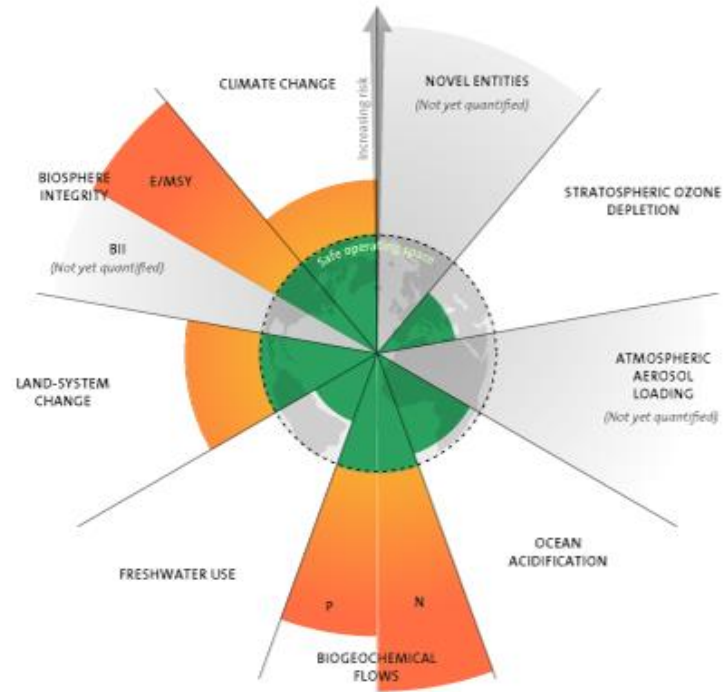
[Planetary boundaries]

2009



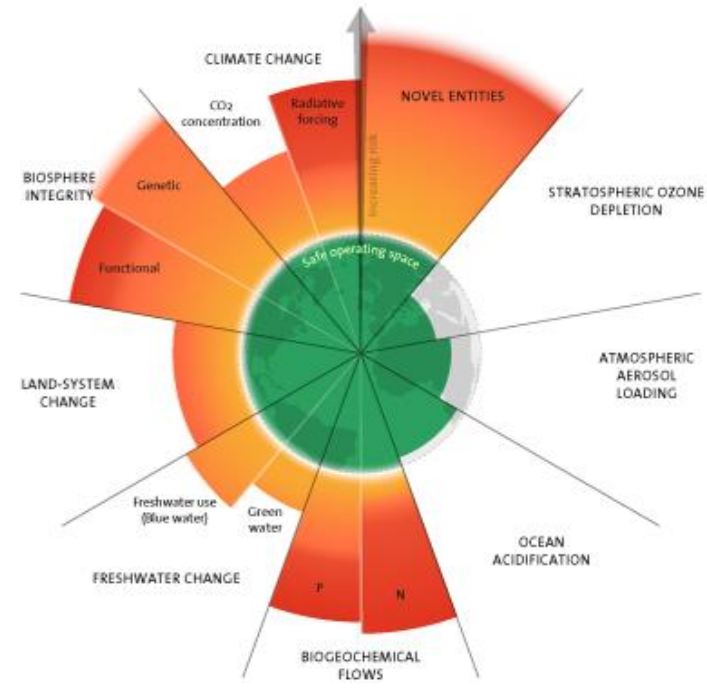
3 boundaries crossed

2015



4 boundaries crossed

2023



6 boundaries crossed

What does **sustainability** mean for you?

Meeting the needs of the present
without compromising the ability of
future generations to meet their own needs.



We've got work to do!



Disclaimer

Take all numbers with a grain of salt!

All estimates
largely depend on

- Hypotheses
often unclear
- Data sources

I've done my best to use only
reasonably trustworthy sources.

I expect I've got correct
orders of magnitude.

I may be wrong...



Be critical!




Read [Does not compute]

What's the carbon footprint of
one hour streaming Netflix?


What's the carbon footprint of one hour streaming Netflix?

Point this way



for "less than X"

Point that way



for "more than X"

Let's count in **number of boils** of an electric kettle.



Zero?

Four?

One?

Ten?

Two?

Fifty?

Three?

Hundred?

What's the carbon footprint of one hour streaming Netflix?

Let's count in **number of boils** of an electric kettle.

▶ 0.077 kWh in 2021 that is ~ three boils
55 gCO₂eq (*)
225m driving an ave. gasoline car

* : considering the 2021 average carbon intensity in the UK
[Netflix 2022 ESG report] [GHG equivalencies calculator]

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

Power is the rate of energy consumption.



Energy

vs.



Power

Dim.
Units

~ Distance (m)

~ Speed (m/s)

SI

Joules (J)

Watt (W)

Common

Kilowatt-hour (kWh)

Kilowatt (kW)

one kilowatt of power
delivered for one hour

rate corresponding to
1000 joules per second

Power is the rate of energy consumption.



Power is not “consumed.”
Power is “drawn.”



Energy

~ Distance (m)

Joules (J)

Kilowatt-hour (kWh)

one kilowatt of power delivered for one hour



Power

~ Speed (m/s)

Watt (W)

Kilowatt (kW)

rate corresponding to 1000 joules per second

Rate of energy usage

$1W = 1J/s$

Dim.
Units

SI

Common

“Carbon” is often used as metric for all emissions of greenhouse gas. —GHG



Dim.
Units

Carbon

Common

gCO₂eq or gCO₂e
or gCO₂-eq

- **Carbon** is often used as a broad term to refer to the impact of all types of emissions and activities on global warming.
- Carbon **equivalence** is a measurement term used to measure this impact.
E.g., 1 ton of methane has the same warming effect as about 84 tons of CO₂
▶ 84 tons CO₂eq
- We often shorten further to just “carbon,” which is then used to refer to all GHGs.

[Methane emissions]

There are different types of low-carbon energy sources.



Classification is unformal and somewhat subjective

Clean energy

comes from a source that

does not produce carbon emissions

E.g.,

Nuclear

Green energy

comes from nature

Hydropower

Renewable energy

do not expire

Wind, solar




There is some overlap between those categories

Producing electricity emits carbon.

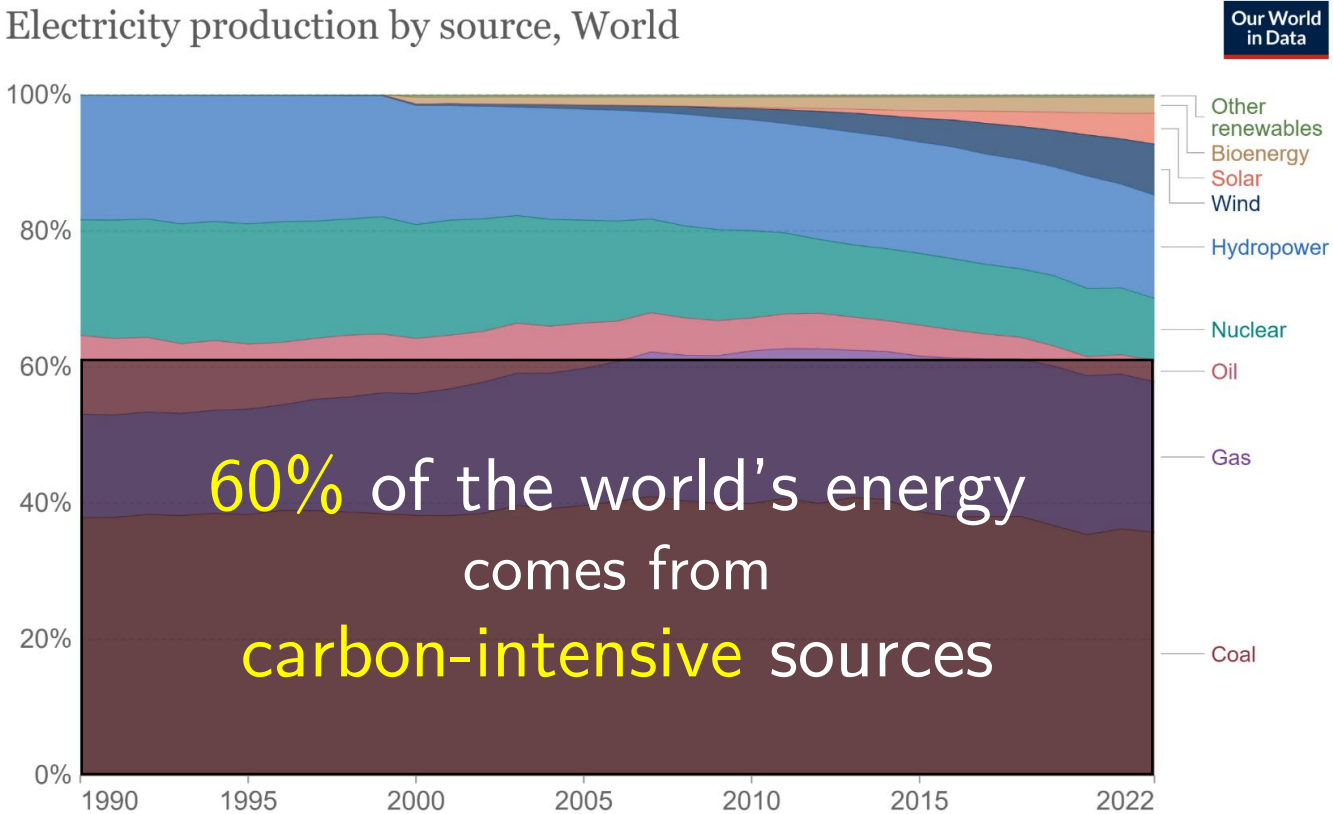


... Wait, what about solar, hydro, etc?



Producing electricity emits carbon.

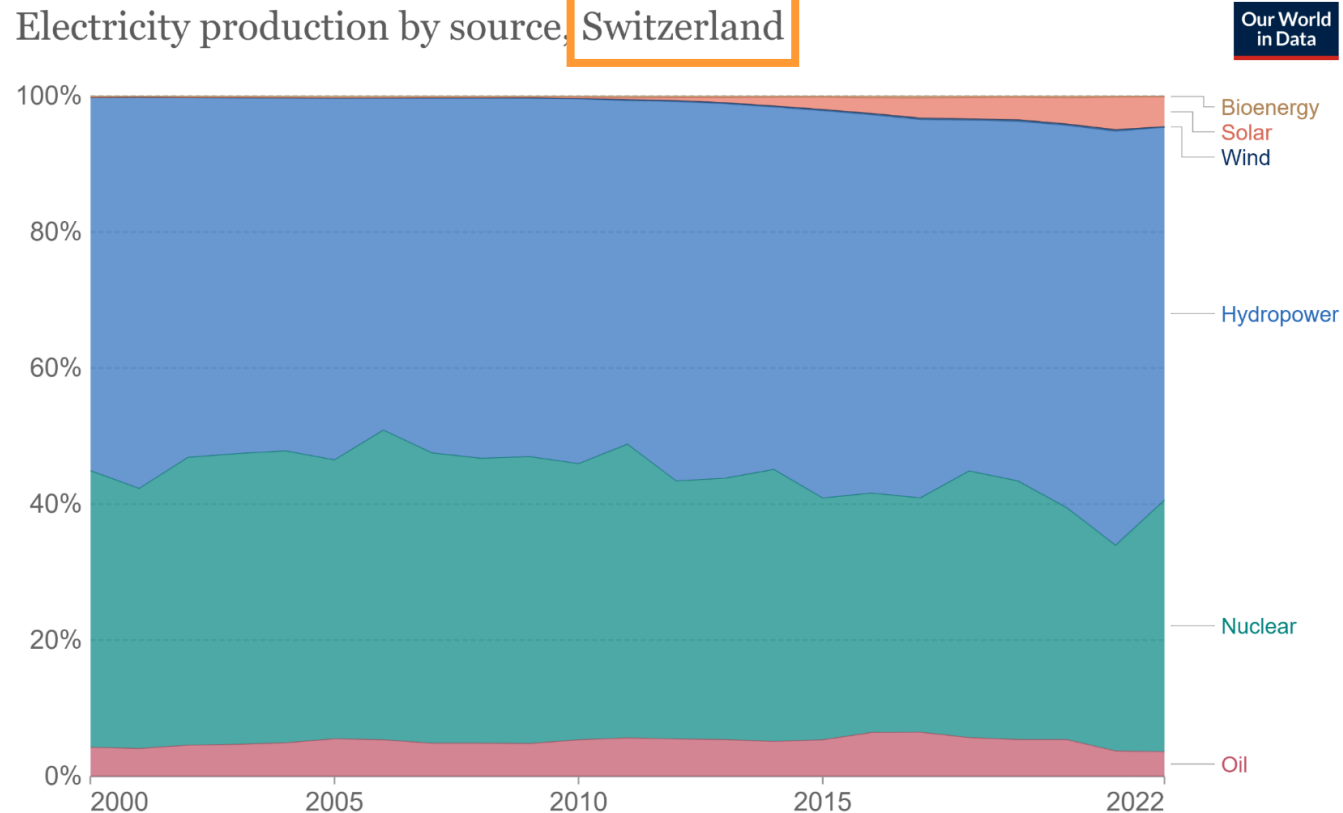
Electricity production by source, World



Despite the uptake of low-carbon energy, the majority of the world's energy still comes from carbon-intensive sources.

Data source: Ember's Yearly Electricity Data; Ember's European Electricity Review; Energy Institute Statistical Review of World Energy
Note: 'Other renewables' includes waste, geothermal, wave and tidal.
OurWorldInData.org/energy | CC BY

Producing electricity emits carbon. Switzerland is an outlier.



Despite the uptake of low-carbon energy, the majority of the world's energy still comes from carbon-intensive sources.

⚠ Electricity is quite clean in Switzerland, but hydropower cannot be scaled much further!

Data source: Ember's Yearly Electricity Data; Ember's European Electricity Review; Energy Institute Statistical Review of World Energy
Note: 'Other renewables' includes waste, geothermal, wave and tidal.
OurWorldInData.org/energy | CC BY

[Rel. energy prod. | CH]

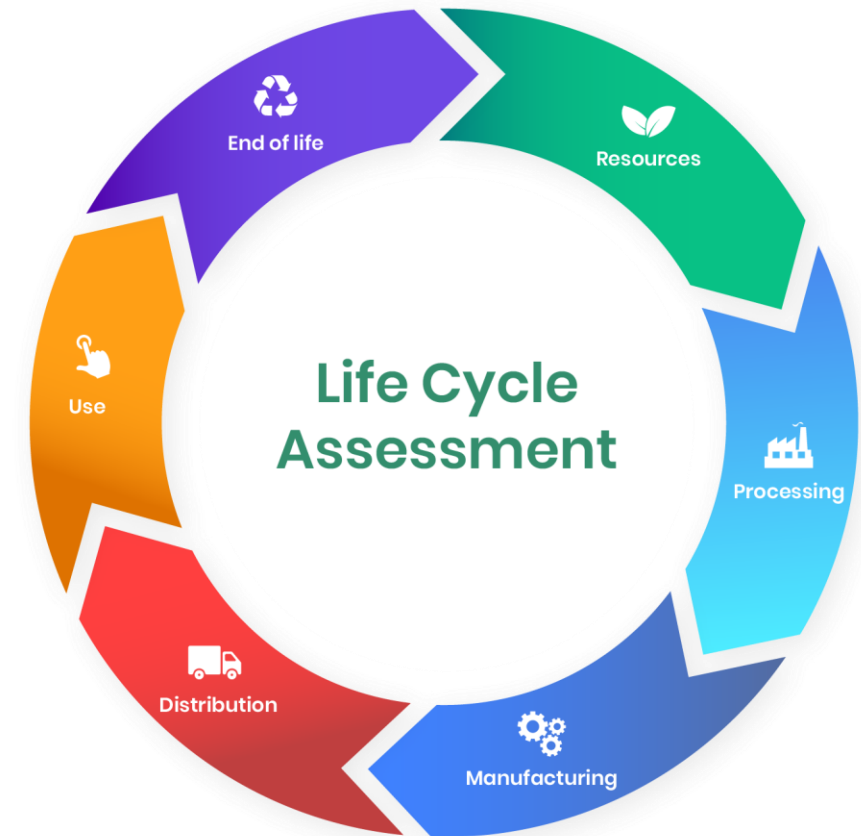
What consumes energy in computer networks?

Like for all (physical) products

- All steps of the product life cycle from components extraction to recycling.

During the use phase of the product

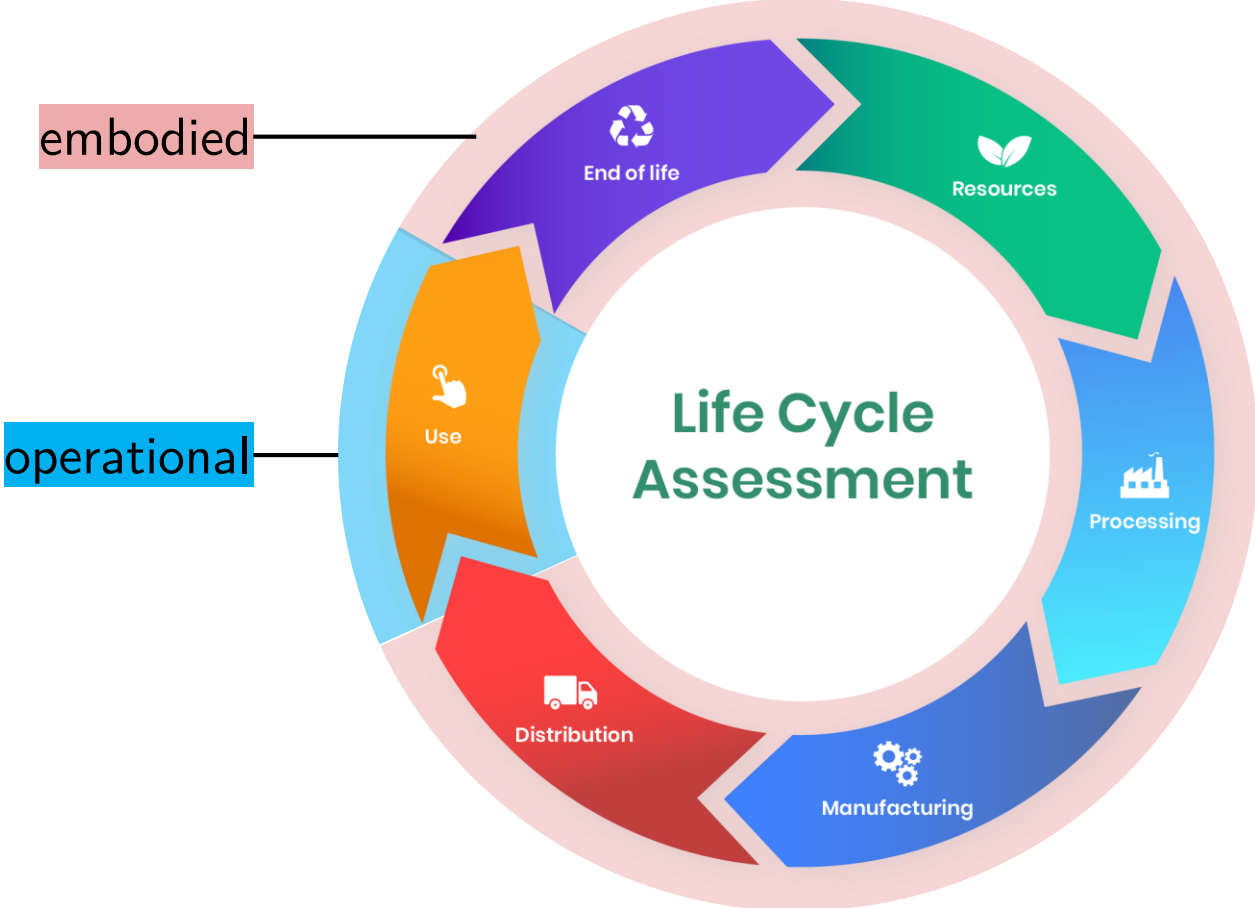
- Routers, transceivers, optical amplifiers
- The infrastructure needed to cool them!



[LCA]

It is important to distinguish operational and embodied carbon footprints

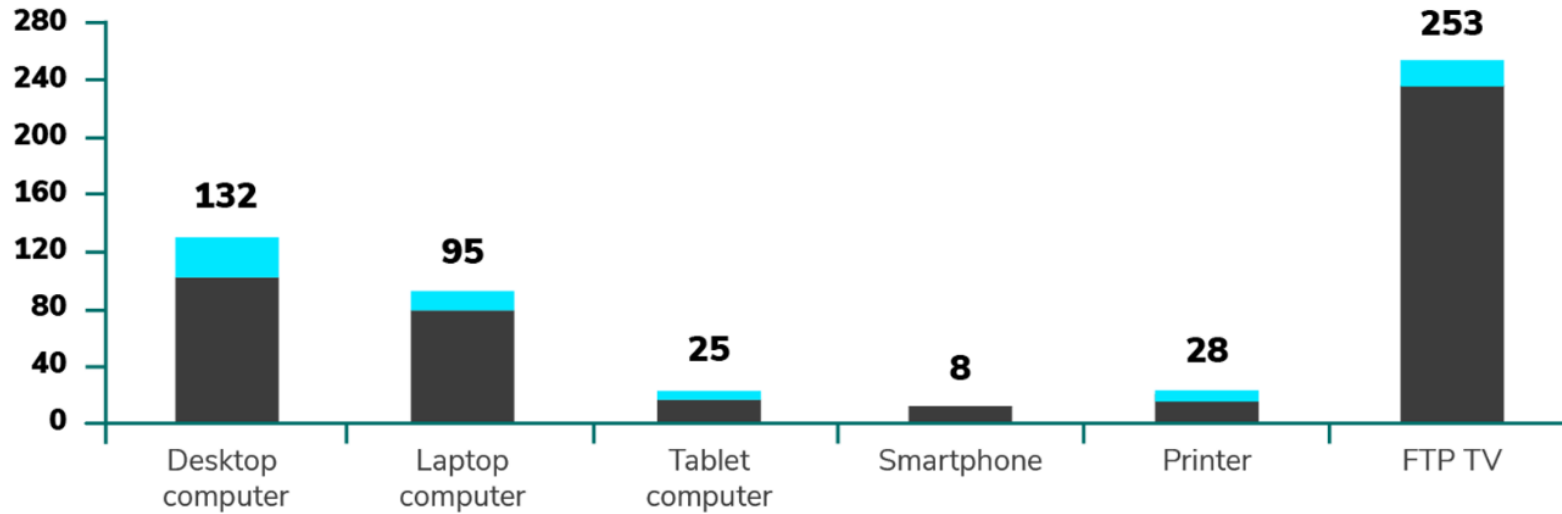
Embodied carbon or “embedded carbon” refers to the carbon pollution resulting from the creation and disposal of a product.



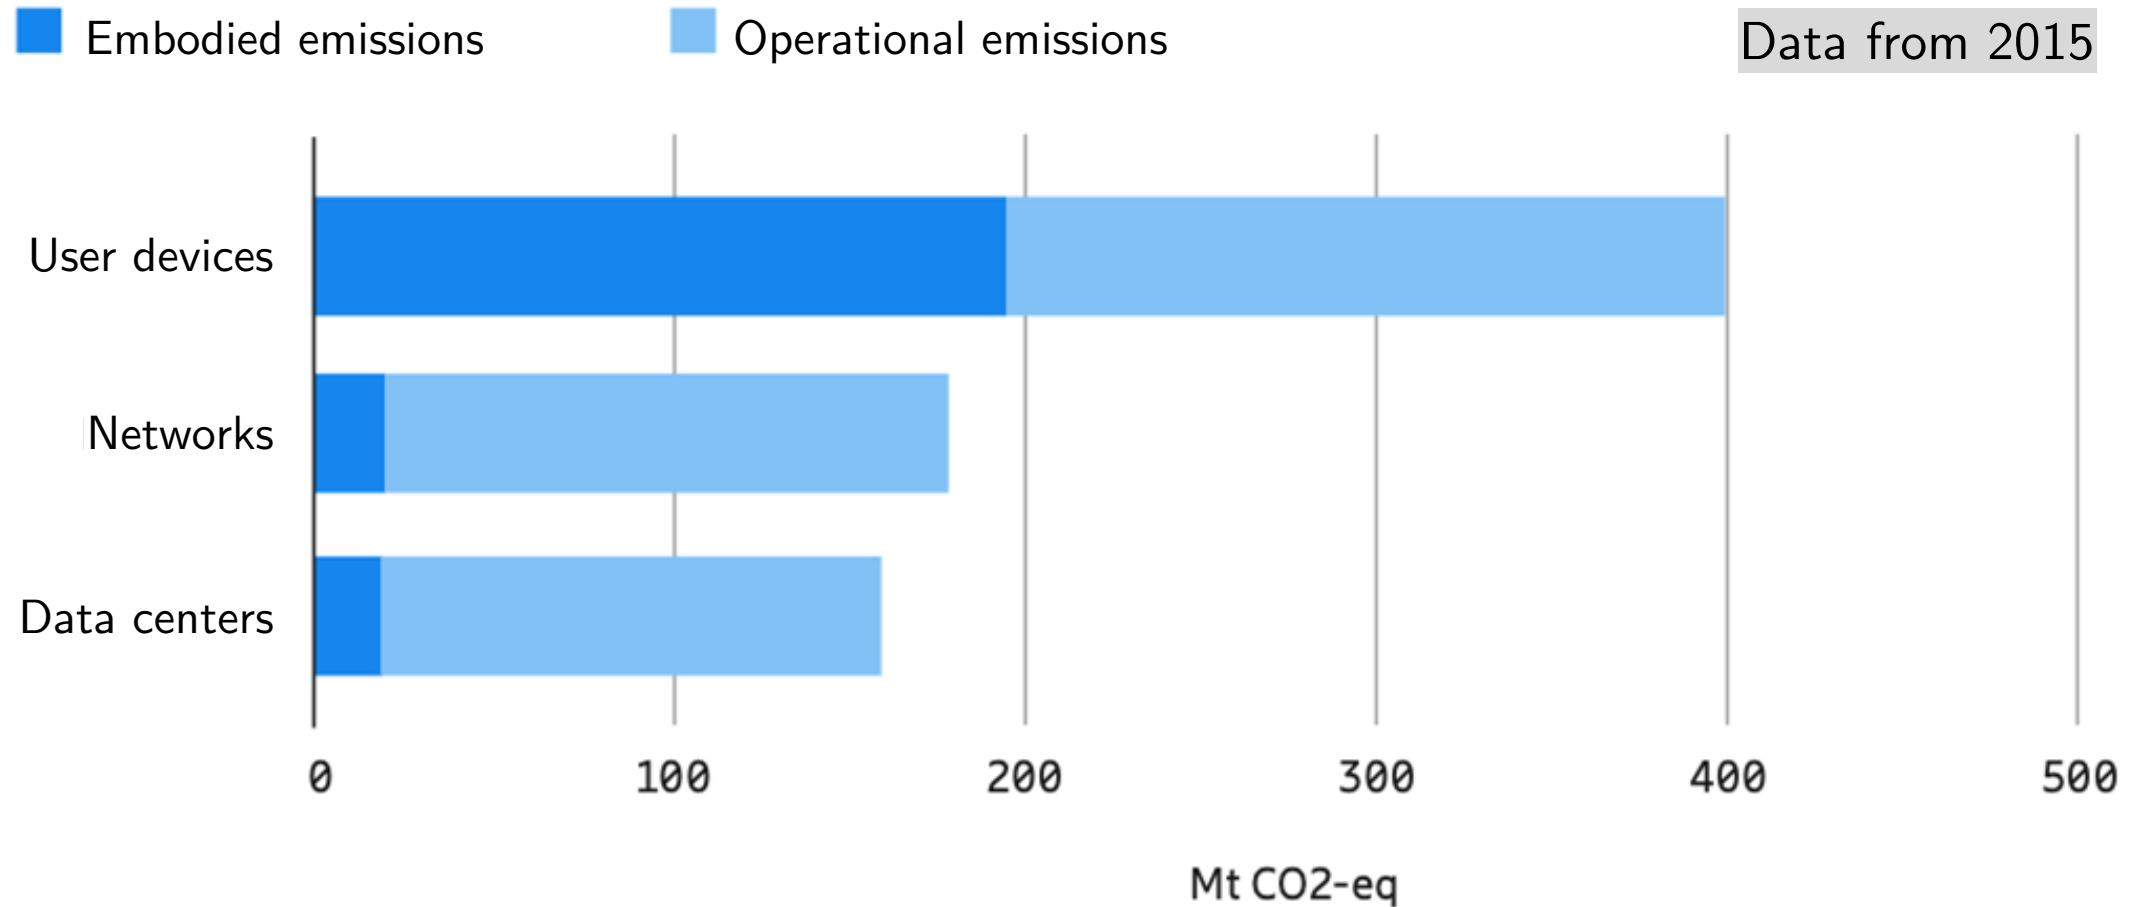
For consumer devices, the embodied footprint dominates.

CO_{2e} emission per ICT end user device


kg CO_{2e} / year



For networked devices, it tends to be the opposite



For networked devices, it tends to be the opposite because **their operational footprint is huge!**

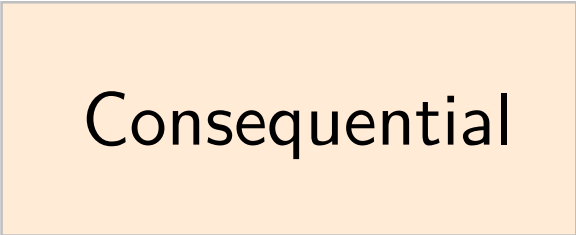


It is also important to distinguish attributional and consequential reasoning.



Attributional

Are these my carbon emissions?



Consequential

What are the consequences of this activity in terms of carbon emissions?

Attributional

Are these my carbon emissions?

Consequential


What are the consequences of this activity
in terms of carbon emissions?

Attributional reasoning is about the allocation of the footprint to different parties.

Policy

We divide the emissions equally between among parties.

Imagine 4 people driving on a trip to a meeting.



Attributional reasoning is about the allocation of the footprint to different parties.


Policy

We divide the emissions equally between among parties.

Imagine 4 people driving on a trip to a meeting.

How much is allocated to each party?

- ▶ $\frac{1}{4}$ of the cost for driving one car.




Attributional reasoning is about the allocation of the footprint to different parties.

Policy

We divide the emissions equally between among parties.

Now imagine a 5th person joins requiring a second car



Attributional reasoning is about the allocation of the footprint to different parties.


Policy

We divide the emissions equally between among parties.

Now imagine a 5th person joins requiring a second car

How much is allocated to each party?

▶ $2/5$ of the cost for driving one car.



Attributional

Are these my carbon emissions?



Consequential

What are the consequences of this activity
in terms of carbon emissions?

Consequential reasoning is about weighting the impact of alternatives.

Imagine five people driving to meeting, requiring a two cars.

We need the second car as a consequence of the fifth person coming.





Consequential reasoning is about weighting the impact of alternatives.

Imagine five people driving to meeting, requiring a two cars.

We need the second car as a consequence of the fifth person coming.

We can use this info to consider **alternative options**.



What if the fifth person joins the meeting remotely?

Consequential reasoning is about weighting the impact of alternatives.

Consequential reasoning weighs the pros and cons of decisions often in terms of total carbon emissions rather than focusing on how to allocate the responsibility to each party.

[Consequential LCA]

Are these my carbon emissions?

Attributional

- Accounting purposes
- “Post-mortem” analysis without predicting power

What are the consequences of this activity in terms of carbon emissions?

Consequential

- Informing decision making
- Alternative-based analysis aiming to predict the effect of a change

Accounting needs guiding principles.


Attributional

- Accounting purposes
- “Post-mortem” analysis without predicting power


Enters the GHG protocol.

[GHG protocol]


The GHG protocol divides emissions into three scopes.



Direct emissions from the fossil fuels you burn.



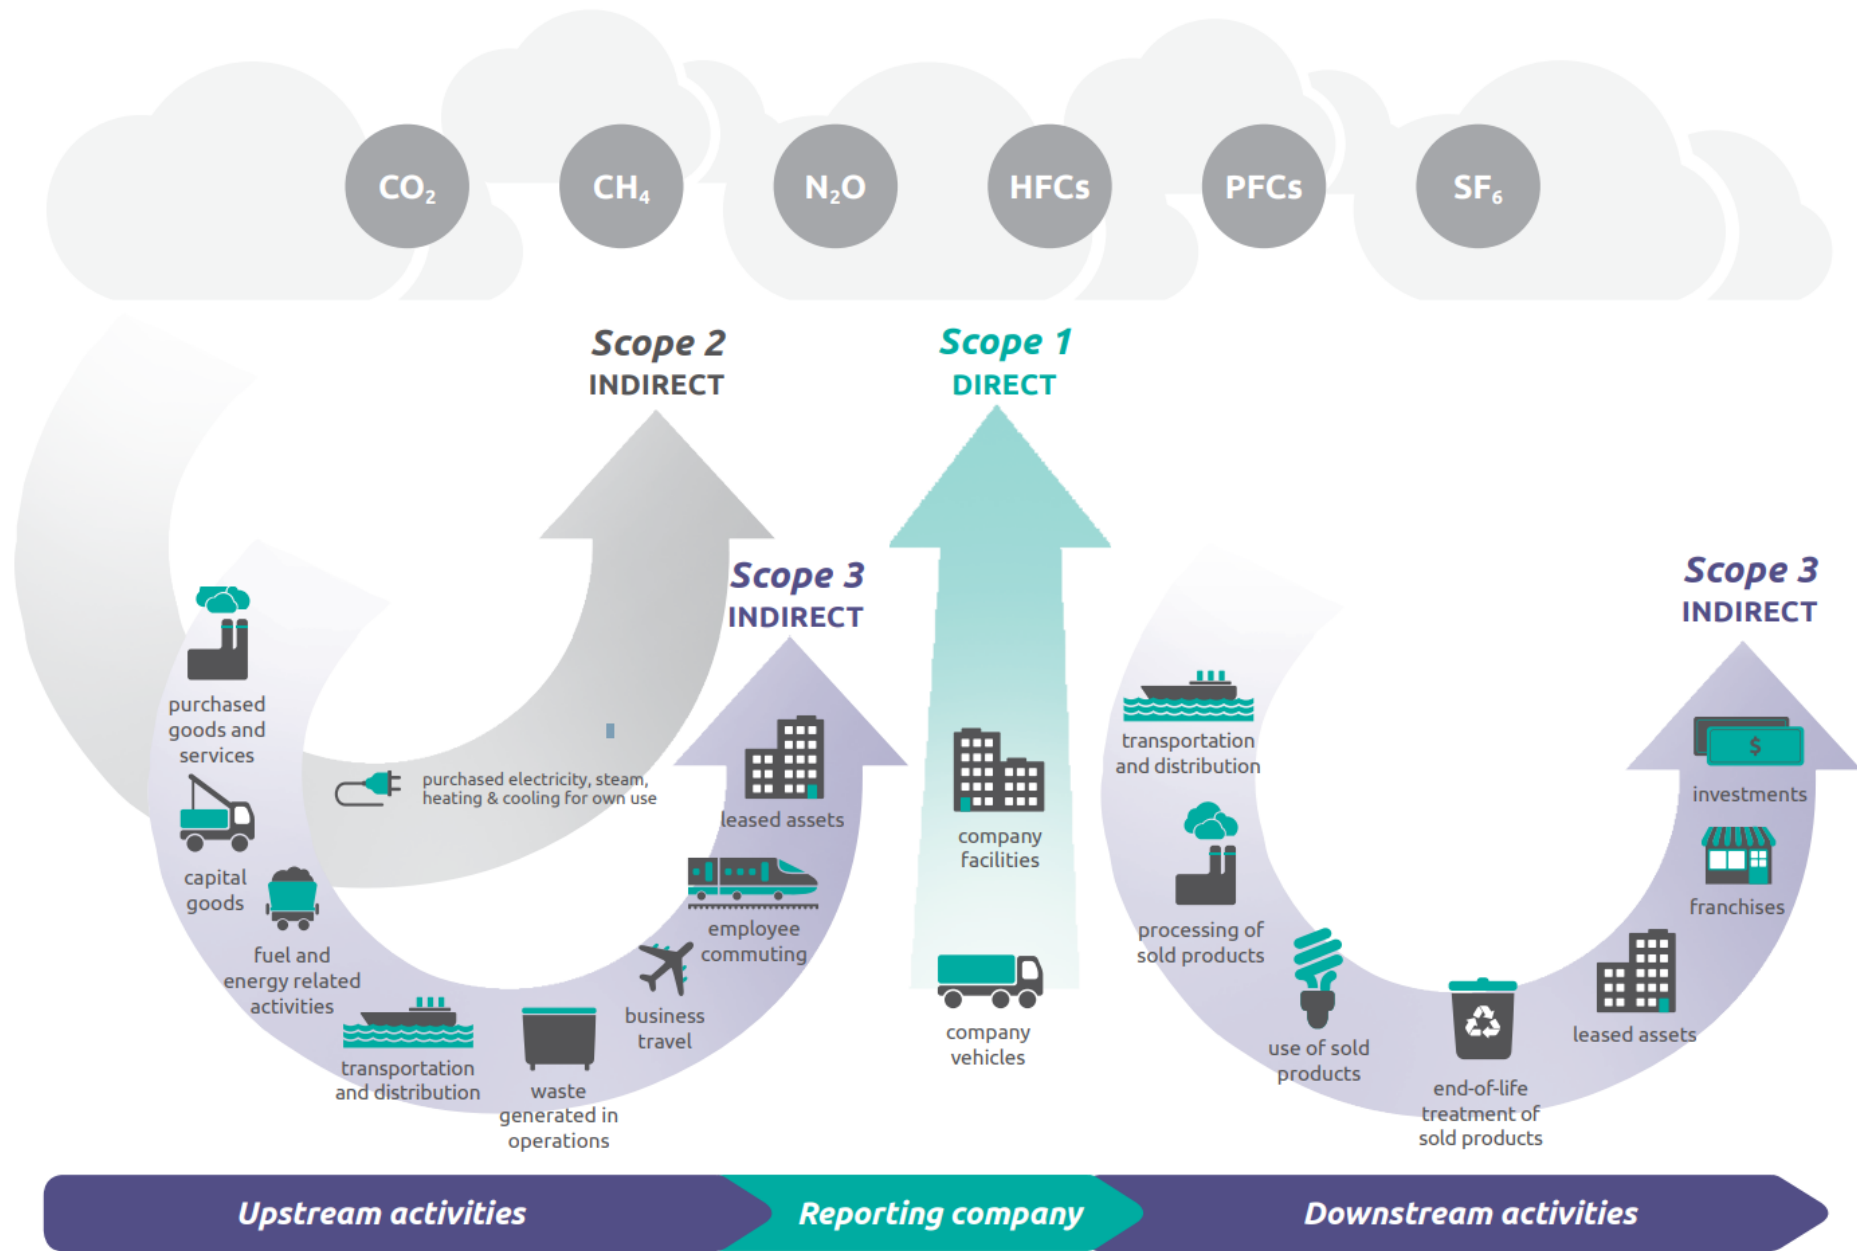
Indirect emissions from the electricity you use.



Indirect emissions from your supply and value chains.

That's "all the rest"





[GHG protocol]

The emissions' scope for a given product depends on who made it, and how.

Let's consider the preparation of your favorite hot beverage.

Scope 1



Direct emissions from the fossil fuels you burn.

Scope 2



Indirect emissions from the electricity you use.

Scope 3



Indirect emissions from your supply and value chains.

Setting the boundaries of Scope 3 is not trivial and somewhat subjective.



“Your scope 3 is someone else’s scopes 1&2”

What do you count in the footprint of...

[Hotcarbon 2023, A]

- a web search?
- a video call?
- an email?

On the Promise and Pitfalls of Optimizing Embodied Carbon

Noman Bashir, David Irwin, Prashant Shenoy
University of Massachusetts Amherst

ABSTRACT

To halt further climate change, computing, along with the rest of society, must reduce, and eventually eliminate, its carbon emissions. Recently, many researchers have focused on estimating and optimizing computing’s *embodied carbon*, i.e., from manufacturing computing infrastructure, in addition to its *operational carbon*, i.e., from executing computations, primarily because the former is much larger than the latter but has received less research attention. Focusing attention on embodied carbon is important because it can incentivize i) operators to increase their infrastructure’s efficiency and lifetime and ii) downstream suppliers to reduce their own operational carbon, which represents upstream companies’ embodied carbon. Yet, as we discuss, focusing attention on embodied carbon may also introduce harmful incentives, e.g., by significantly overstating real carbon reductions and complicating the incentives for directly optimizing operational carbon. This position paper’s purpose is to mitigate such harmful incentives by highlighting both the promise and potential pitfalls of optimizing embodied carbon.

CCS CONCEPTS

• **Hardware** → **Impact on the environment**; • **General and reference** → **Metrics**; • **Social and professional topics** → **Sustainability**.

KEYWORDS

Embodied and operational carbon emissions, metrics, sustainability.

ACM Reference Format:

Noman Bashir, David Irwin, Prashant Shenoy. 2023. On the Promise and Pitfalls of Optimizing Embodied Carbon. In Proceedings of the ACM Conference on Environmental and Sustainable Computing. ACM, New York, NY, 1–11.

tax or cap-and-trade system, to provide a direct financial incentive for businesses to adopt low-carbon energy. Such an incentive would be configurable based on the magnitude of carbon’s price. Many governments have adopted carbon taxes and cap-and-trade systems [12]. Of course, since carbon pricing raises energy costs, it can hurt legacy carbon-based energy businesses. As a result, many countries including the U.S. are unlikely to ever introduce a direct carbon pricing policy, and instead are using more indirect means. For example, the recent U.S. Inflation Reduction Act takes an indirect approach to financially incentivizing lower carbon energy by providing various tax subsidies for actions that promote its use [19].

Since the financial incentives to adopt low-carbon energy introduced by the policies above are complex and likely not strong enough to reduce carbon emissions fast enough to avoid the worst outcomes of climate change, there has also been growing social pressure for companies to reduce their carbon footprint from their investors, customers, and employees, i.e., as part of Environmental, Social, and Governance (ESG) investing initiatives [14]. As a result, many companies now publicly report their annual estimated carbon emissions based on the Greenhouse Gas (GHG) protocol [7], which is required in some countries and may soon be in the U.S. [20]. The GHG protocol divides carbon and other emissions into Scopes 1, 2, and 3: Scope 1 emissions derive from directly burning fuels and other chemicals, e.g., by company vehicles, generators, industrial processes, etc.; Scope 2 emissions derive from purchasing energy, e.g., from the electric grid; and Scope 3 emissions derive from all other aspects of a company’s value chain, including carbon emissions from manufacturing the products and services a company uses. Scopes 1 and 2 roughly represent a company’s *operational carbon*, while Scope 3 represents a company’s *embodied carbon*.

The GHG protocol provides a general framework, to be translated into domain-specific guidelines.

One example

Created to measure, understand and ultimately reduce the emissions of serving **digital media** and **entertainment products**.

The logo for DIMPACT features the word "DIMPACT" in a bold, dark blue, sans-serif font. The letter "D" is significantly larger than the other letters. A green horizontal bar with a white bar chart pattern is positioned below the "D".

DIMPACT

Insight to action on digital carbon impacts

[DIMPACT]

Quick recap

There are several useful ways of looking at carbon.

- Operational vs. Embodied emissions
- Attributional vs. Consequential reasoning
- Scoped emissions
 - 1 my direct emissions
 - 2 my indirect emissions
 - 3 other indirect emissions

- Operational vs. Embodied emissions
- Attributional vs. Consequential reasoning
- Scoped emissions
 - 1 my direct emissions
 - 2 my indirect emissions
 - 3 other indirect emissions

How do we improve?

It is useful to normalize the footprint by considering efficiency metrics.

$$\begin{array}{l} \text{Operational} \\ \text{Carbon efficiency} \end{array} = \boxed{\frac{\text{J used}}{\text{Task}}} \times \boxed{\frac{\text{J supplied}}{\text{J used}}} \times \boxed{\frac{\text{Carbon}}{\text{J supplied}}}$$

Device efficiency (HW + SW) Datacenter PUE Carbon intensity

It is useful to normalize the footprint by considering efficiency metrics.

Operational
Carbon efficiency =

$$\frac{J \text{ used}}{\text{Task}}$$

x

$$\frac{J \text{ supplied}}{J \text{ used}}$$

x

$$\frac{\text{Carbon}}{J \text{ supplied}}$$


Device efficiency
(HW + SW)

Datacenter
PUE

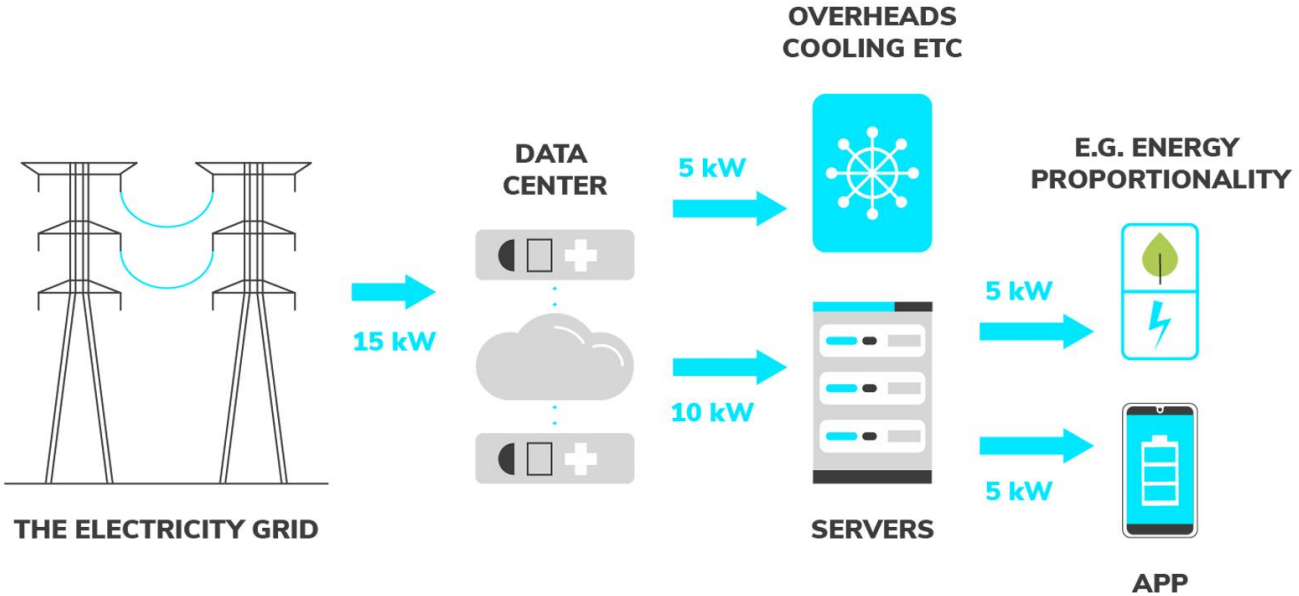
Carbon
intensity

= Energy efficiency
(Wh/b)

A device efficiency depends on its utilization and energy propotionality.



The Power Usage Effectiveness (PUE) quantifies the infrastructure overhead.



$$\text{PUE} = \frac{\text{Total energy use}}{\text{Energy used for compute and comm.}}$$


aka "Pointless Use of Energy"



PUE = 1.5

static power = 5kW

Since the introduction of PUE in 2006, progress have been made.




“What is the annual PUE for your largest data center?”

What is the average annual PUE for your largest data center?


Source: Uptime Institute Global Survey of IT and Data Center Managers 2019, n=624

Hyperscalers have already reached the plausible limits.

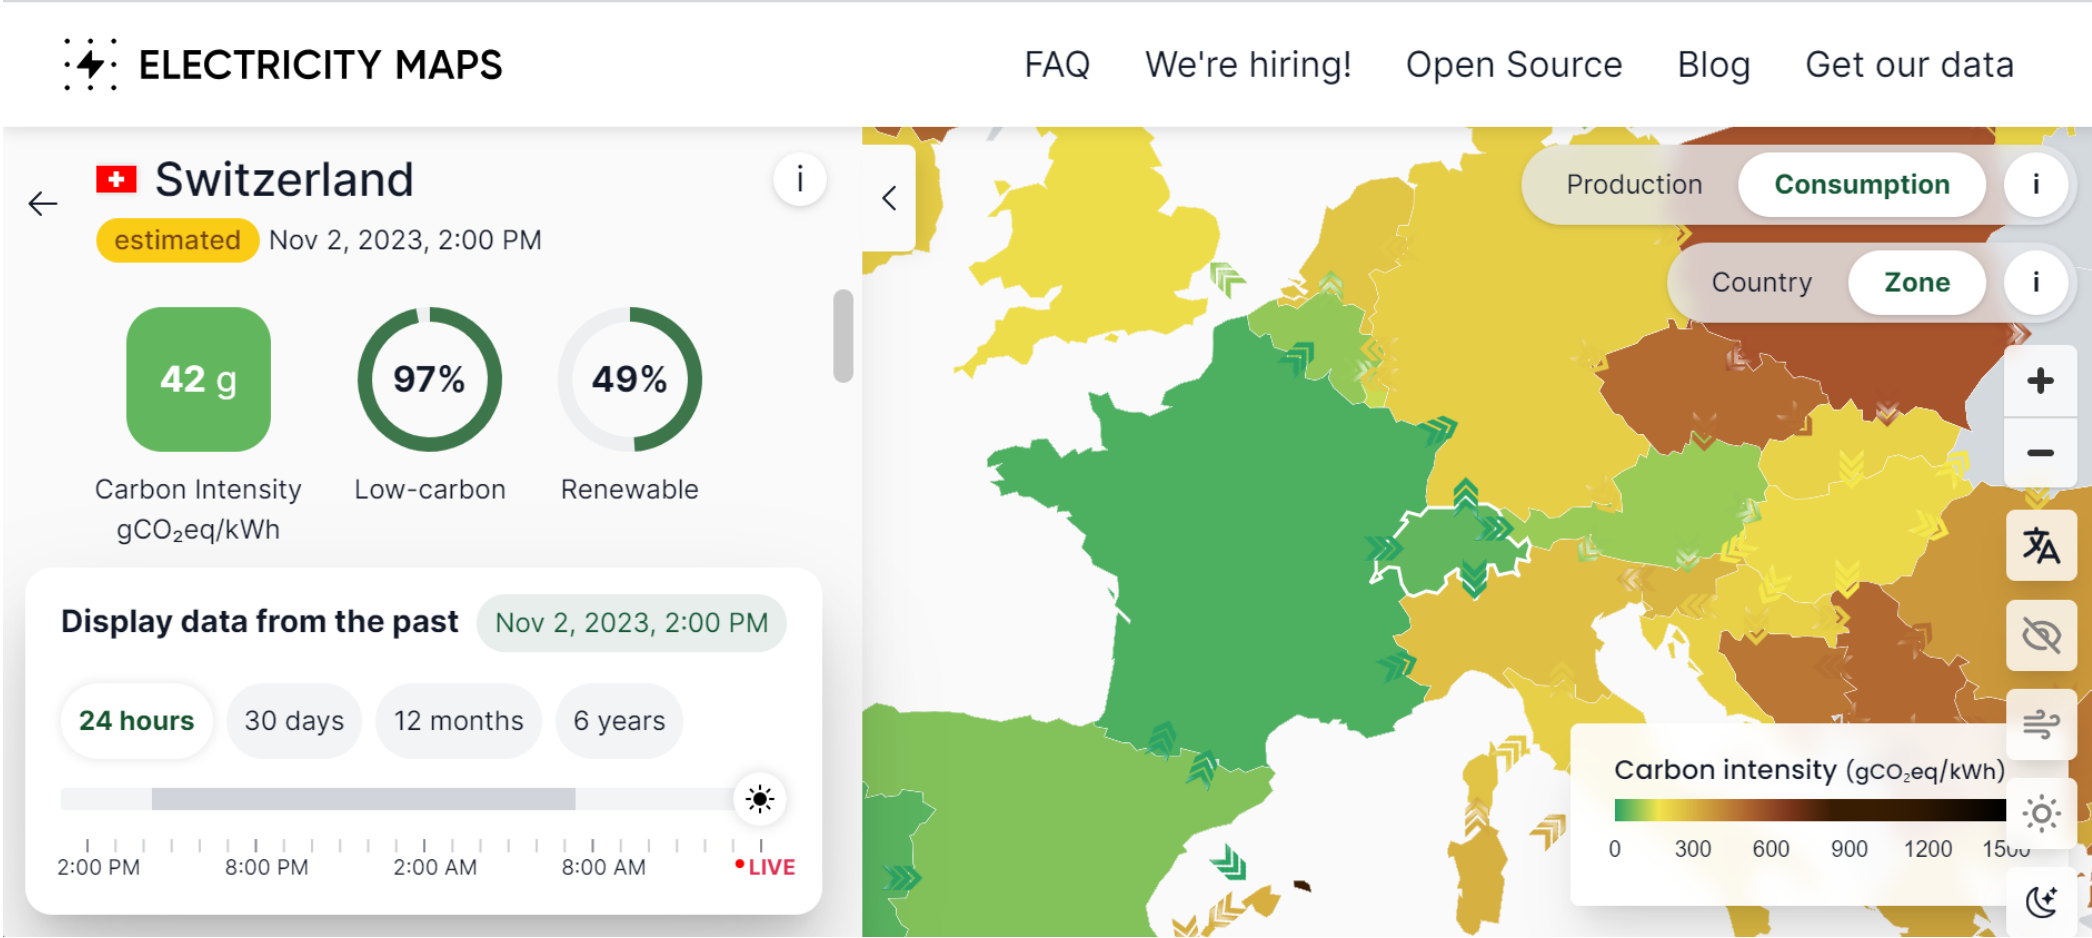


Continuous PUE Improvement

Average PUE for all data centers



The carbon intensity of the electricity grid fluctuates in time and space.



[Electricity Maps]

Breaking down the carbon efficiency allows identifying where there is room for improvements.

|
Focus of the next hour

$$\begin{array}{l} \text{Operational} \\ \text{Carbon efficiency} \end{array} = \boxed{\frac{\text{J used}}{\text{Task}}} \times \boxed{\frac{\text{J supplied}}{\text{J used}}} \times \boxed{\frac{\text{Carbon}}{\text{J supplied}}}$$

Device efficiency (HW + SW) Datacenter PUE Carbon intensity

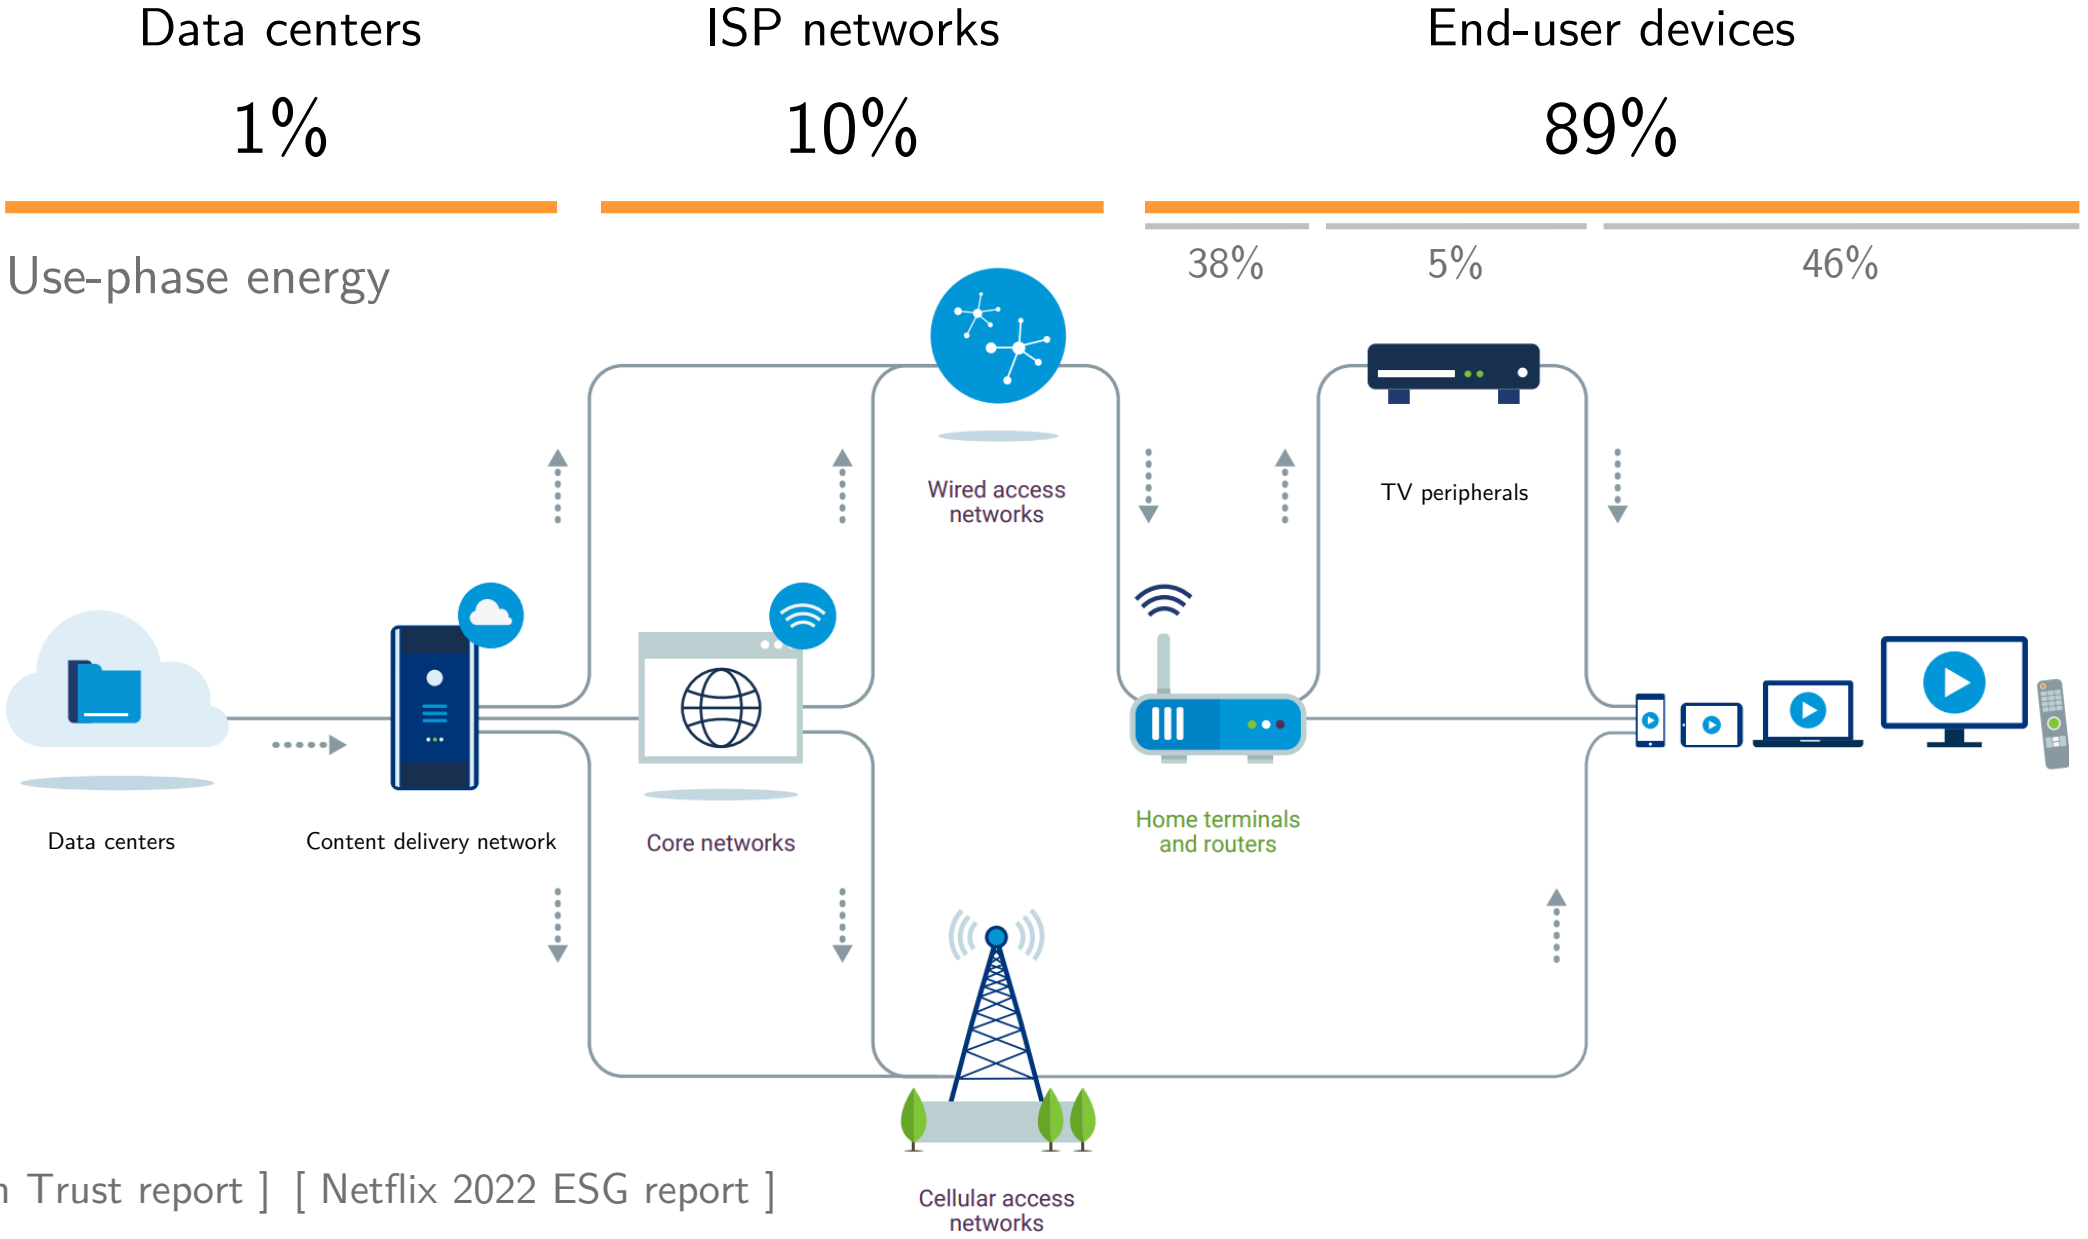
What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

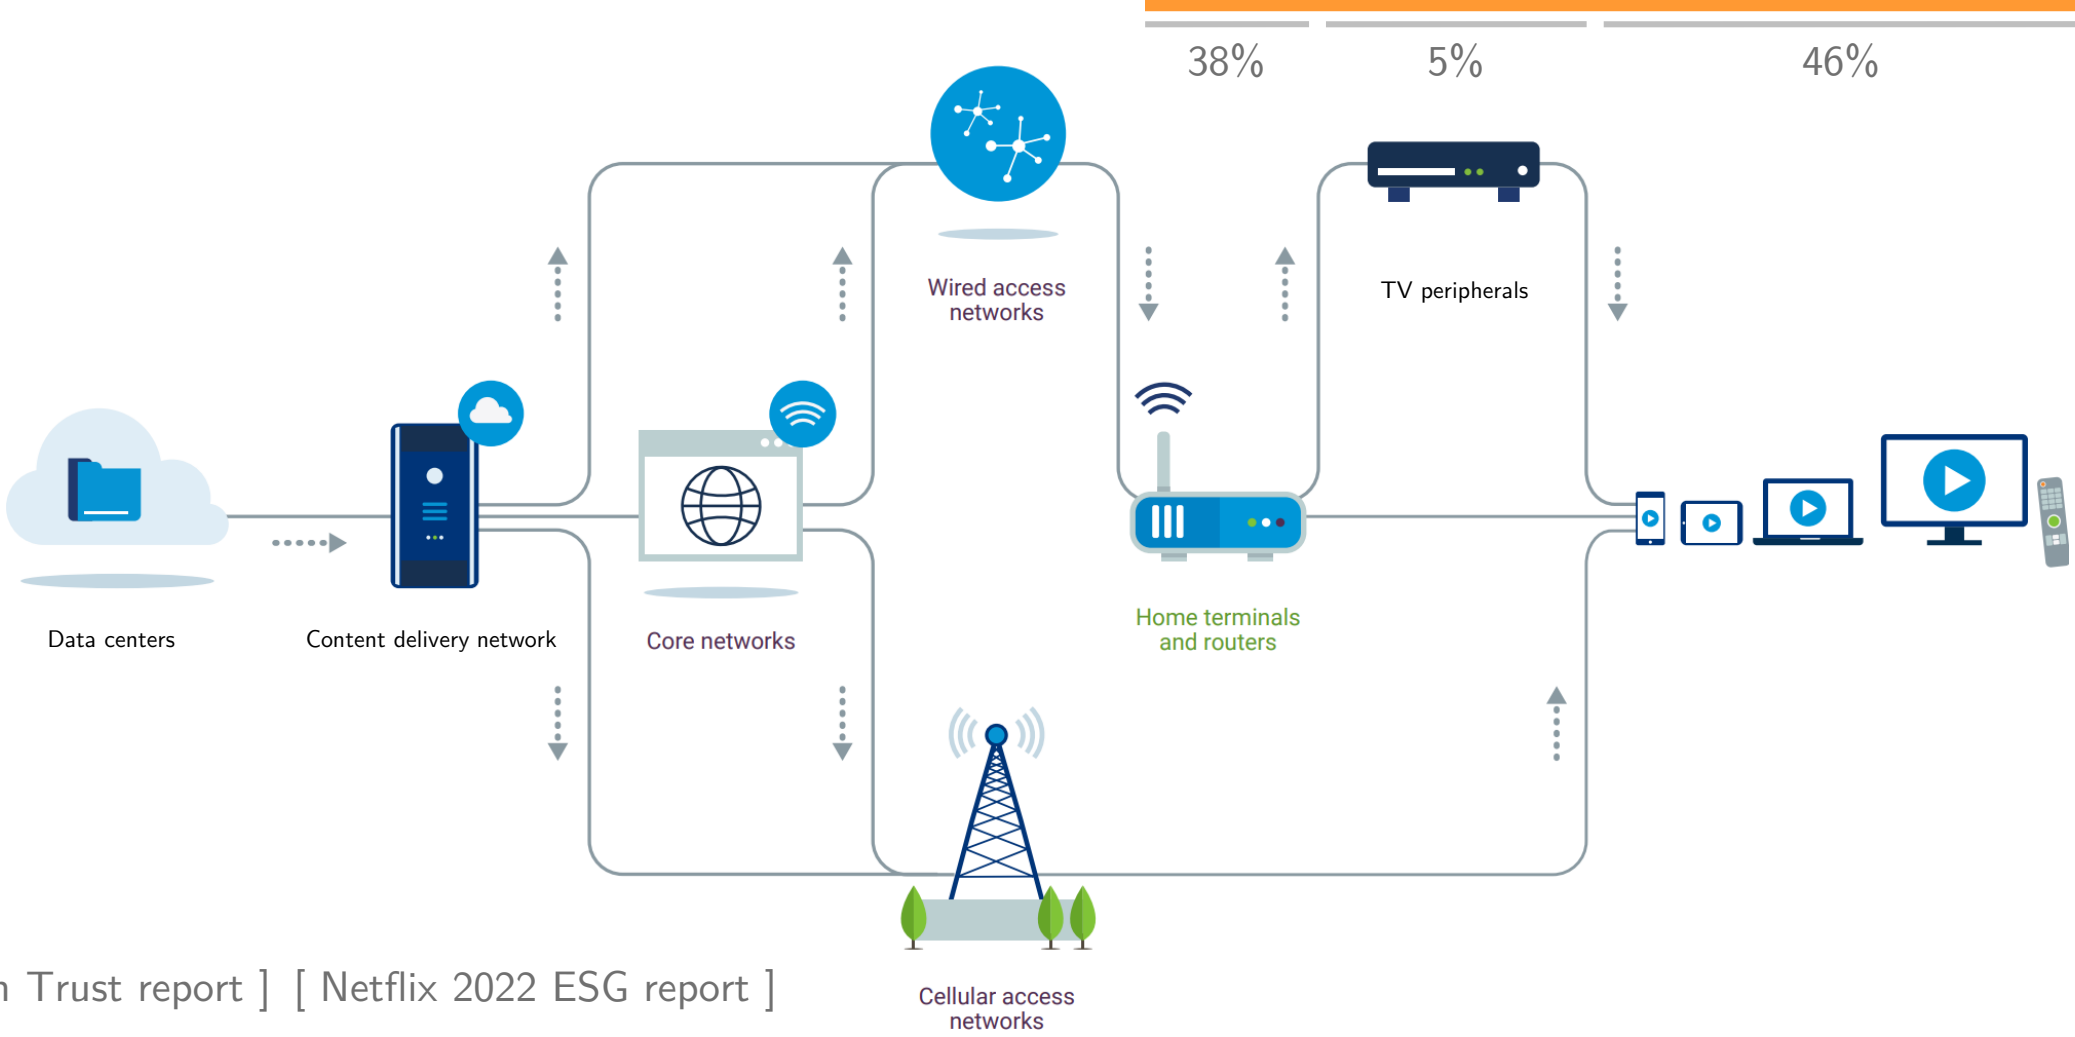
Can technology
save us?



[Carbon Trust report] [Netflix 2022 ESG report]

End-user devices dominate the use-phase energy usage.

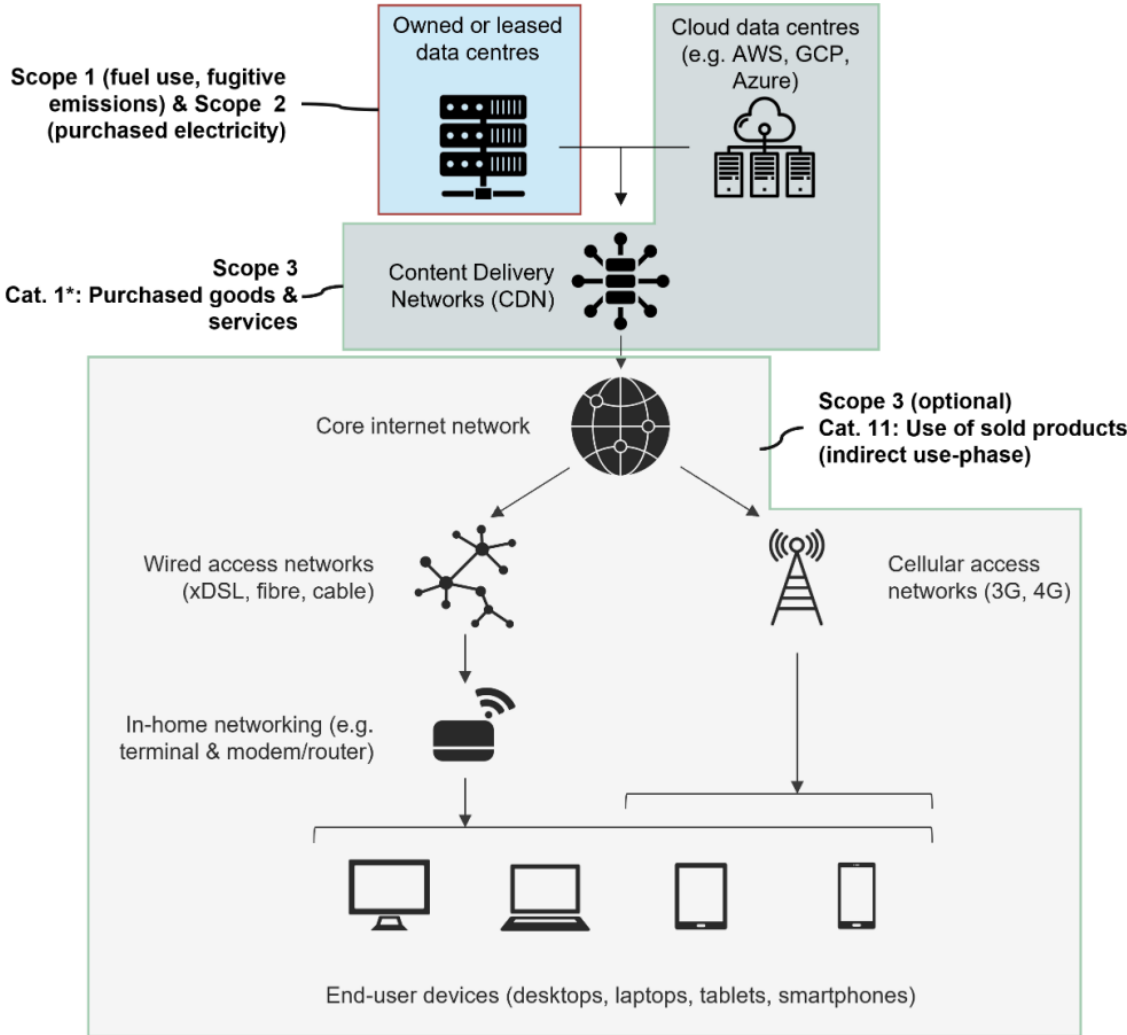
End-user devices
89%



[Carbon Trust report] [Netflix 2022 ESG report]

The “use of sold products” is beyond Netflix’s control and is not included in their reported scoped emissions.

Our responsibility



What Netflix commits to improve on

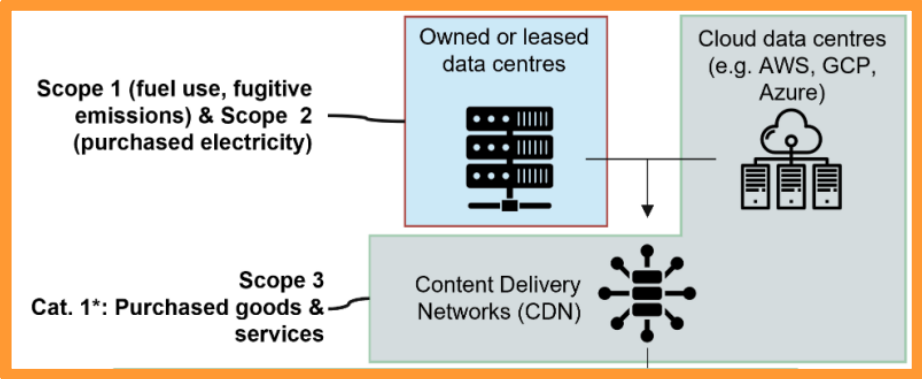
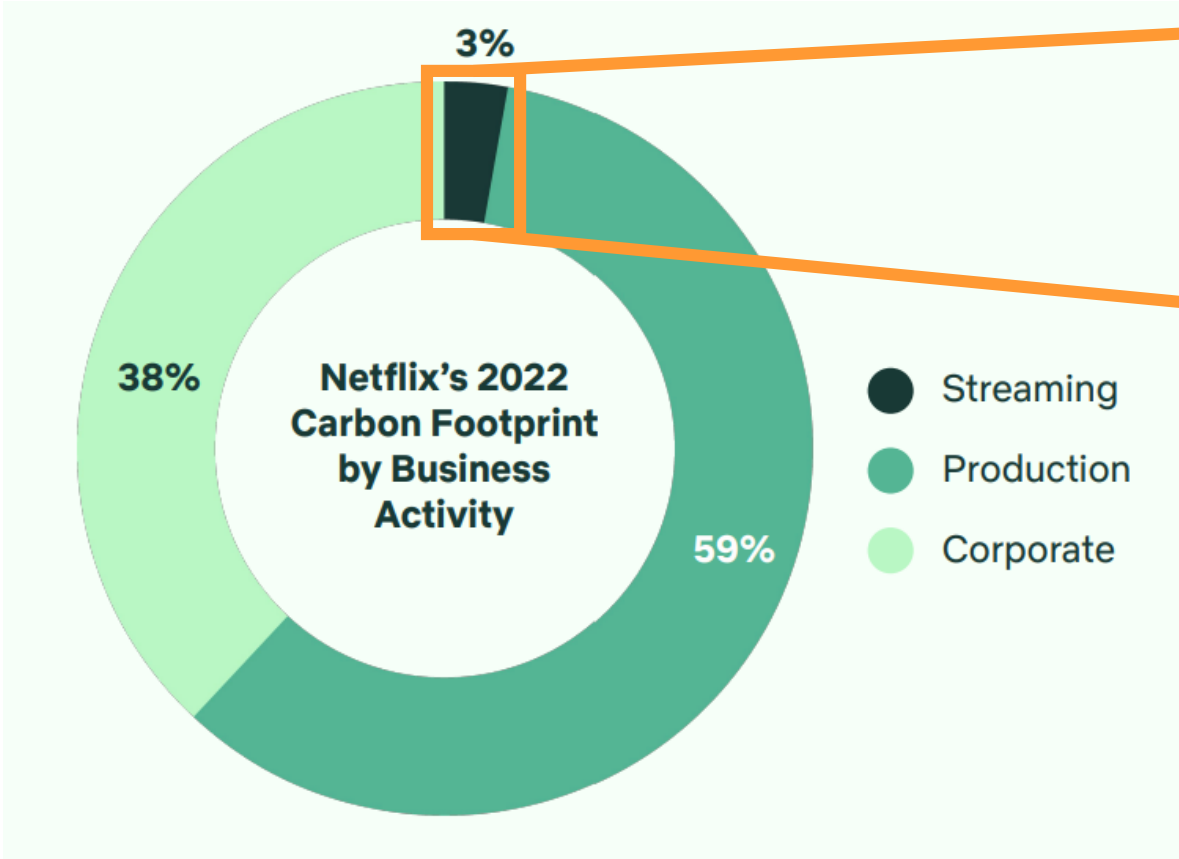
BUT the 55gCO₂e figure **does** include the energy used by end-user devices while streaming!

Netflix’ argument

This energy **is spent** and **matters**, even if it is not on us to reduce it.

* Categories based on GHG Protocol Corporate Value Chain (Scope 3) Accounting and Reporting Standard

Moreover, the actual streaming is by far the smallest part of Netflix's carbon footprint.



[Netflix 2022 ESG report]

Finally, these figures
are all **attributional**.

1h streaming ~ 55gCO₂e



- 55gCO₂e is emitted when you stream for 1h
- 55gCO₂e is 'saved' if you stream 1h less



This analysis assigns responsibilities.
It **does not predict** effects of changes.

Adding up the numbers is... scary.

IT WOULD TAKE 5.12 MILLION TREES GROWING FOR 100 YEARS TO REMOVE THIS CARBON.

LET'S THINK AND ACT MORE CREATIVELY, DRIVE MORE INNOVATION AND DRIVE ACTION TOGETHER



per year



Even if the network is “only” 10% it is relevant to try and improve it!

**CARBON.
CONTENT.
CONFUSION.**

**WHAT NETFLIX TAUGHT ME
ABOUT SUSTAINABILITY**



**204 Million
Global
Subscribers**

**Average 8.82
Hours
consumption pw**



**0.077 KWh ENERGY USED
PER STREAMING HOUR**



**7236989.76
MWh**

**5,128,728
MTCO2e**



POWER 617,618 HOMES



DRIVE 12.8 BILLION MILES



CHARGE 623 BILLION SMART PHONES

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

Breaking down the carbon efficiency allows identifying where there is room for improvements.

Operational
Carbon efficiency = $\frac{\text{J used}}{\text{Task}} \times \frac{\text{J supplied}}{\text{J used}} \times \frac{\text{Carbon}}{\text{J supplied}}$

... improves with better... Networks & compute Infrastructure & HW design Application elasticity

Operational
Carbon efficiency

=

$$\frac{\text{J used}}{\text{Task}}$$

×

$$\frac{\text{J supplied}}{\text{J used}}$$

×

$$\frac{\text{Carbon}}{\text{J supplied}}$$

... improves
with better...

Networks
& compute

Infrastructure
& HW design

Application
elasticity

Not really a
network matter



Operational
Carbon efficiency

=

$$\frac{\text{J used}}{\text{Task}}$$

×

$$\frac{\text{J supplied}}{\text{J used}}$$

×

$$\frac{\text{Carbon}}{\text{J supplied}}$$

... improves
with better...

Networks
& compute

Infrastructure
& HW design

Application
elasticity



Improves with
better networks

Operational
Carbon efficiency

=

$$\frac{\text{J used}}{\text{Task}}$$

×

$$\frac{\text{J supplied}}{\text{J used}}$$

×

$$\frac{\text{Carbon}}{\text{J supplied}}$$

... improves
with better...

Networks
& compute

Infrastructure
& HW design

Application
elasticity



Main focus
of networking

Operational
Carbon efficiency

=

$$\frac{\text{J used}}{\text{Task}}$$

×

$$\frac{\text{J supplied}}{\text{J used}}$$

×

$$\frac{\text{Carbon}}{\text{J supplied}}$$

... improves
with better...

Networks
& compute

Infrastructure
& HW design

Application
elasticity



Main focus
of networking

Let's consider two energy usage profiles for the same task.



Which option is more energy efficient?

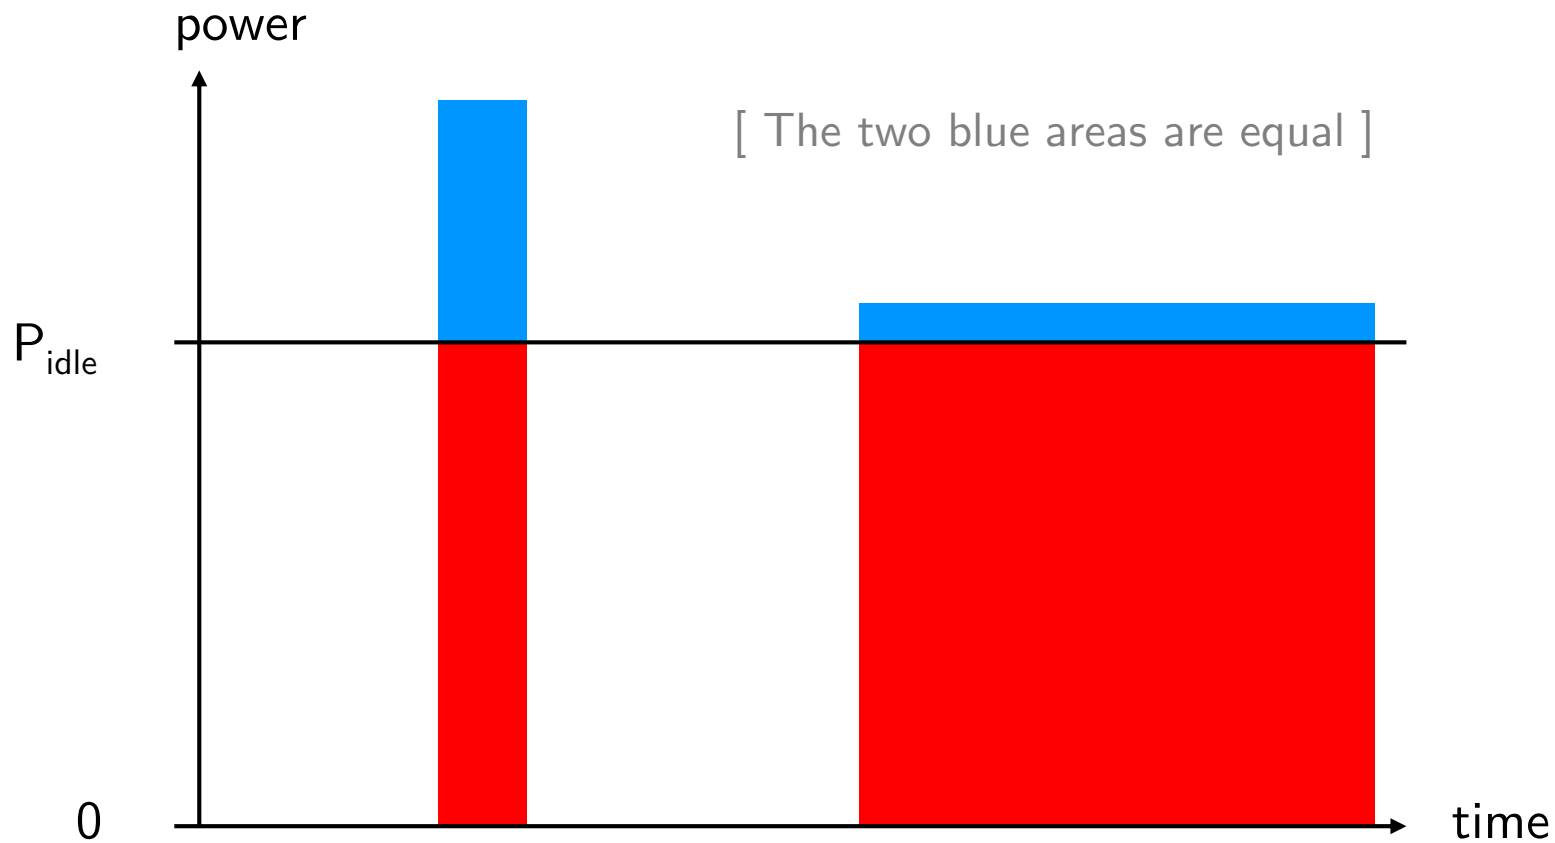
Option 1

- High power
- Short time

Option 2

- Low power
- Long time

Let's consider two energy usage profiles for the same task.



- ◀ Which option is more energy efficient?
- ◀ What about now?

Turning components off whenever possible
is the fundamental way of saving energy.

aka “sleeping”


Sleeping is implemented
in all consumer IT

- Screens Laptops, phones
- Radio duty-cycling IoT devices
- ...

Sleeping is implemented
in all consumer IT

- Screens Laptops, phones
- Radio duty-cycling IoT devices
- ...

What about
network devices?




How does such a plot look like for a switch?


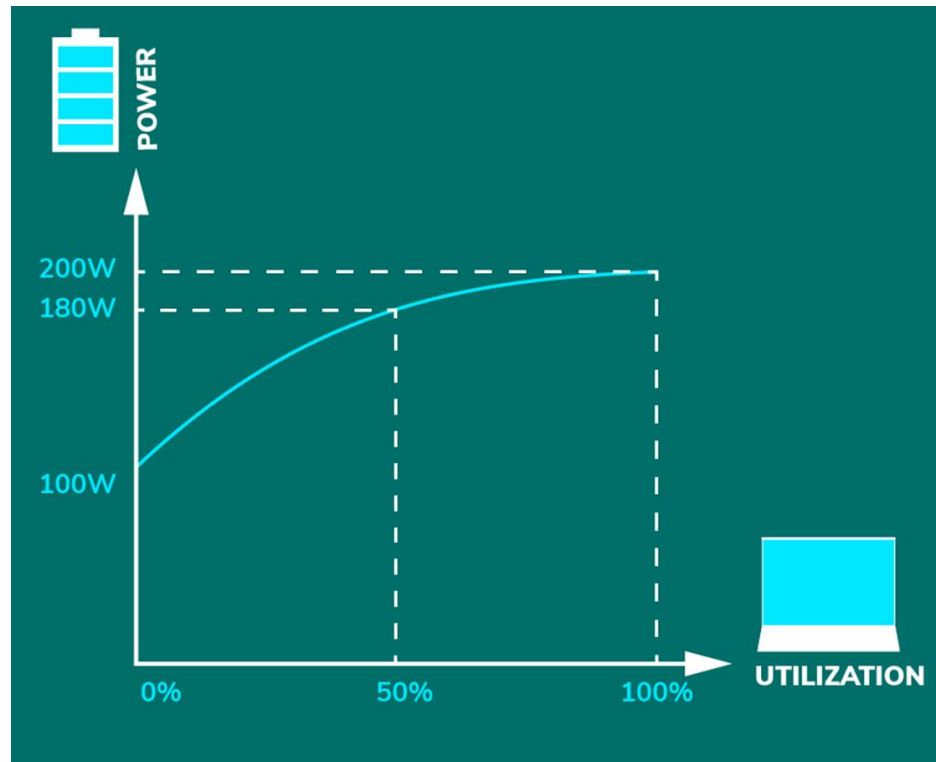
Vendors only provide information about the maximum power... So we measure ourselves!



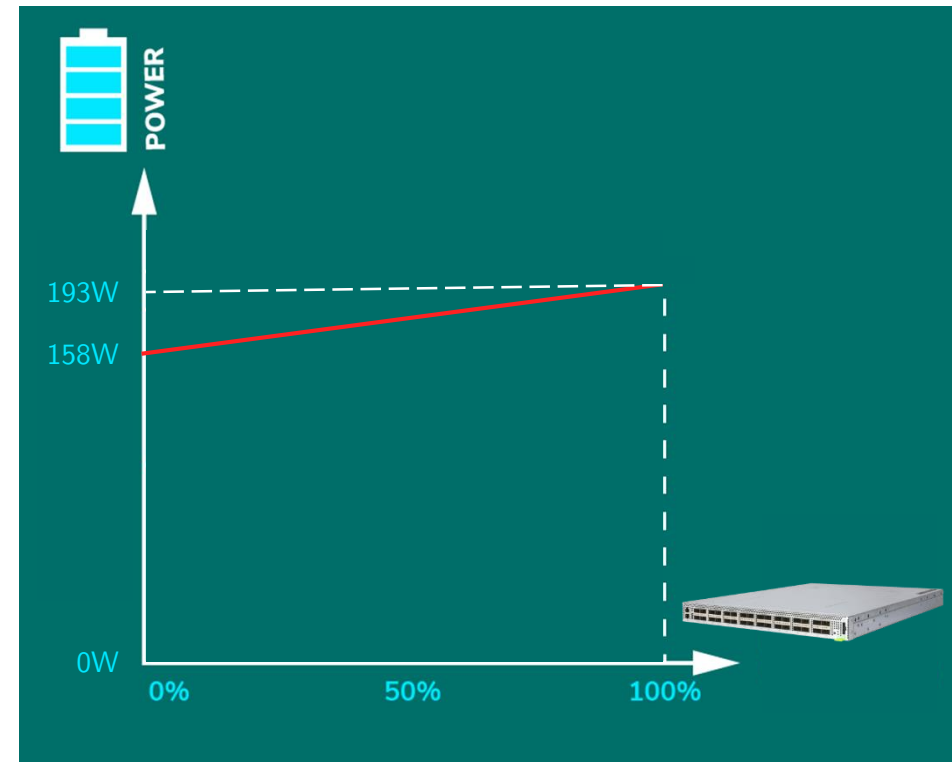

WEDGE 100BF-32X
32 x 100G QSFP28 ports with Tofino 32D



The idle power dominates
i.e., network power is **inelastic**.



The idle power dominates
i.e., network power is **inelastic**.



The traffic in the entire SWITCH network is less than

250 Tb/month

[Oct. 23]


Assuming 100% utilization

One switch could forward one year-worth of traffic in

7,8 min




Typical utilisation is really low.



6.4 Tbps

How “bad” is power inelasticity?

On the bright side, inelasticity means we can carry more traffic with the same power!



[Does not compute]

On the dark side, it results in very inefficient wired networks...

[SIGCOMM 2003]

Greening of the Internet

Maruti Gupta
Department of Computer Science
Portland State University
Portland, OR 97207
mgupta@cs.pdx.edu

Suresh Singh
Department of Computer Science
Portland State University
Portland, OR 97207
singh@cs.pdx.edu

ABSTRACT

In this paper we examine the somewhat controversial subject of energy consumption of networking devices in the Internet, motivated by data collected by the U.S. Department of Commerce. We discuss the impact on network protocols of saving energy by putting network interfaces and other router & switch components to sleep. Using sample packet traces, we first show that it is indeed reasonable to do this and then we discuss the changes that may need to be made to current Internet protocols to support a more aggressive strategy for sleeping. Since this is a position paper, we do not present results but rather suggest interesting directions for core networking research. The impact of saving energy is huge, particularly in the developing world where energy is a precious resource whose scarcity hinders widespread Internet deployment.

Categories and Subject Descriptors

C.2.1 [Network Architecture & Measurement]: [Network Topology]; C.2.2 [Network Protocols]: [Routing Protocols]; C.2.6 [Internetworking]: [Routers, Standards]

General Terms

Algorithms, Measurement, Economics

Keywords

Energy, Internet, Protocols

1. INTRODUCTION

Recently, an opinion has been expressed in various quarters (see [5, 12]) that the energy consumption of the Internet is "too high" and that since this energy consumption can only grow as the Internet expands, this is a cause for concern. One may disagree, as we do, with the qualitative statement that the energy consumption of the Internet is too high, because it is a small fraction of the overall energy

Device	Approximate Number Deployed	Total AEC TW-h
Hubs	93.5 Million	1.6 TW-h
LAN Switch	95,000	3.2 TW-h
WAN Switch	50,000	0.15 TW-h
Router	3,257	1.1 TW-h
Total		6.05 TW-h

Table 1: Breakdown of energy draw of various networking devices (TW-h refers to Tera-Watt hours and AEC to Annual Electricity Consumption).

consumption. However, the absolute numbers do indicate a need to be more energy efficient. We use the analysis presented by these observers as a starting point to discuss an exciting new direction for future core networking research. We believe that if energy can be conserved by careful engineering then there is no reason why we should not do so as this has implications not only for reducing energy needs in the U.S. but also on speeding up Internet deployment and access in the developing world where energy is very scarce.

Table 1 [14] summarizes the energy consumption by Internet devices in the U.S. as of the year 2000. These values are copied from Tables 5-59 (Hub), 5-61 (LAN switch), 5-62 (WAN switch), and 5-64 (Router) of [14]. The data is broken up based on network device type, which is useful in analyzing where and how energy savings can be garnered. In order to arrive at the various energy numbers in the table, the authors took into account the percentage of different types of devices deployed (e.g., number of CISCO 2500 type routers, number of 7505s, etc) and then used the average energy consumption values of these devices to arrive at the final numbers shown in the table¹. Two energy values missing from the table are the energy cost of cooling the equipment and that of UPS (Uninterruptable Power Supplies) equipment². The future expectation is that the energy consumption of networking devices will increase by 1 TW-h by 2005 [14].

Expressed as a percentage of total U.S. energy expenditure in the year 2000, the energy drawn by the devices in Table 1 accounts for approximately 0.07% of the total. Given


The Internet core consumes more Joules per Bytes than wireless LANs.

2x and 24x more...

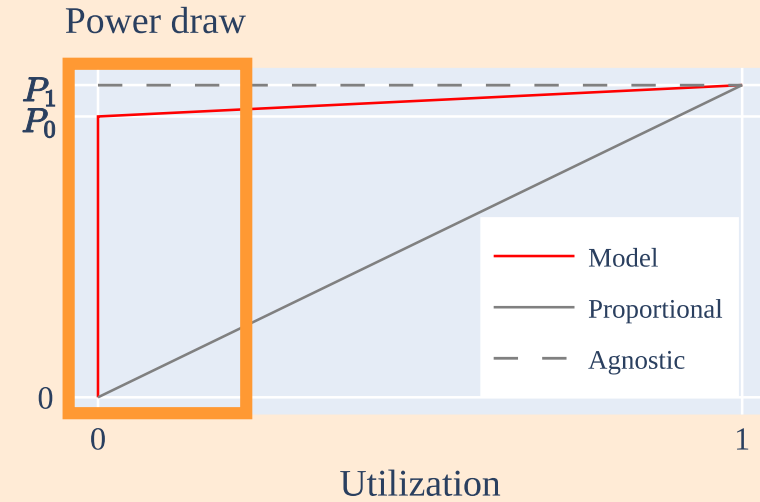
depending on your hypotheses

1. Network devices are always “on.”
2. Network devices’ energy consumption is mainly independent of traffic load.
3. Network devices are under-utilized.

1. Network devices are always “on.”
2. Network devices’ energy consumption is mainly independent of traffic load.
3. Network devices are under-utilized.




1. Network devices are always “on.”
2. Network devices’ energy consumption is mainly independent of traffic load.
3. Network devices are under-utilized.

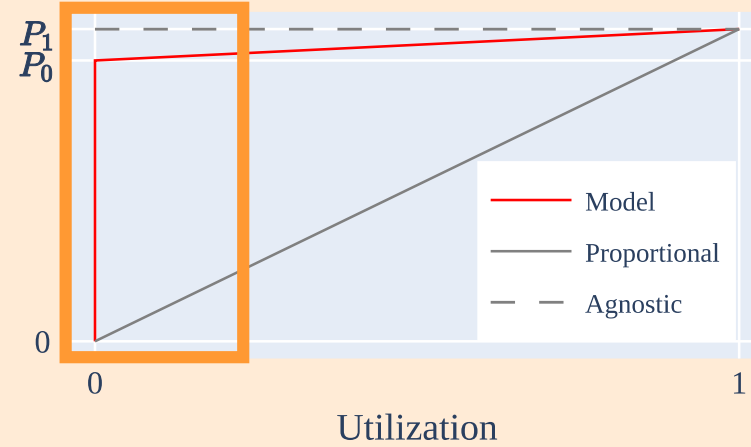


ISP overprovision networks to support

- Peak traffic
- Fault tolerance

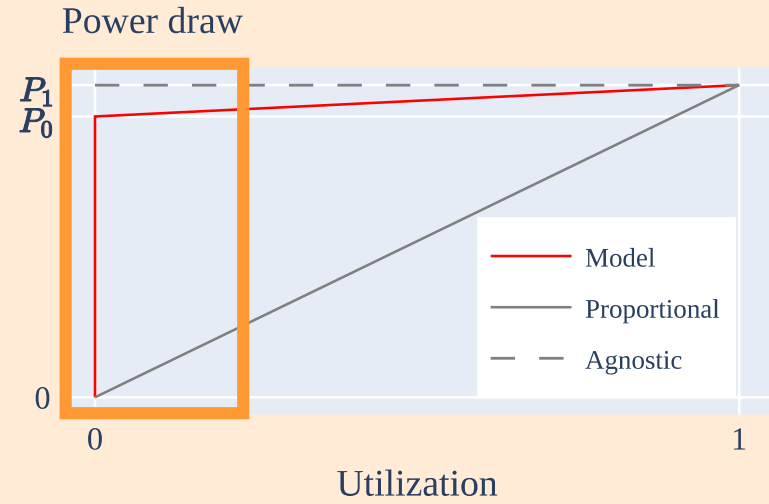



Power draw



ISP overprovision networks to support

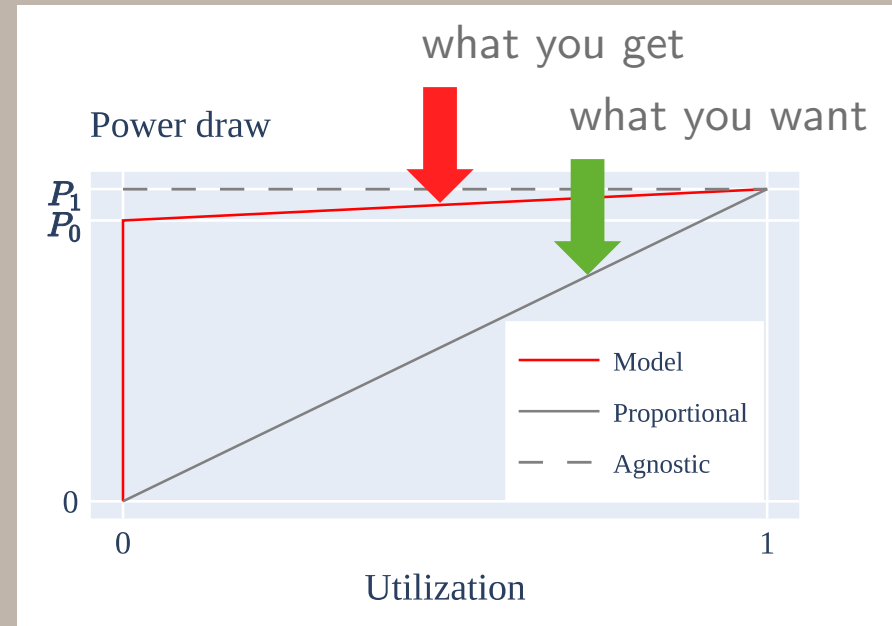
- Peak traffic
- Fault tolerance



ISP overprovision networks to support

- Peak traffic
- Fault tolerance


1. Network devices are always “on.”
2. Network devices’ energy consumption is mainly independent of traffic load.
3. Network devices are under-utilized.



ISP overprovision networks to support


- Peak traffic
- Fault tolerance

What does proportionality mean for our toy example?




- As idle power dominates, low utilization wastes a lot.

What does proportionality mean for our toy example?




- As idle power dominates, low utilization wastes a lot.
- Reducing idle power yields better proportionality.

What does proportionality mean for our toy example?



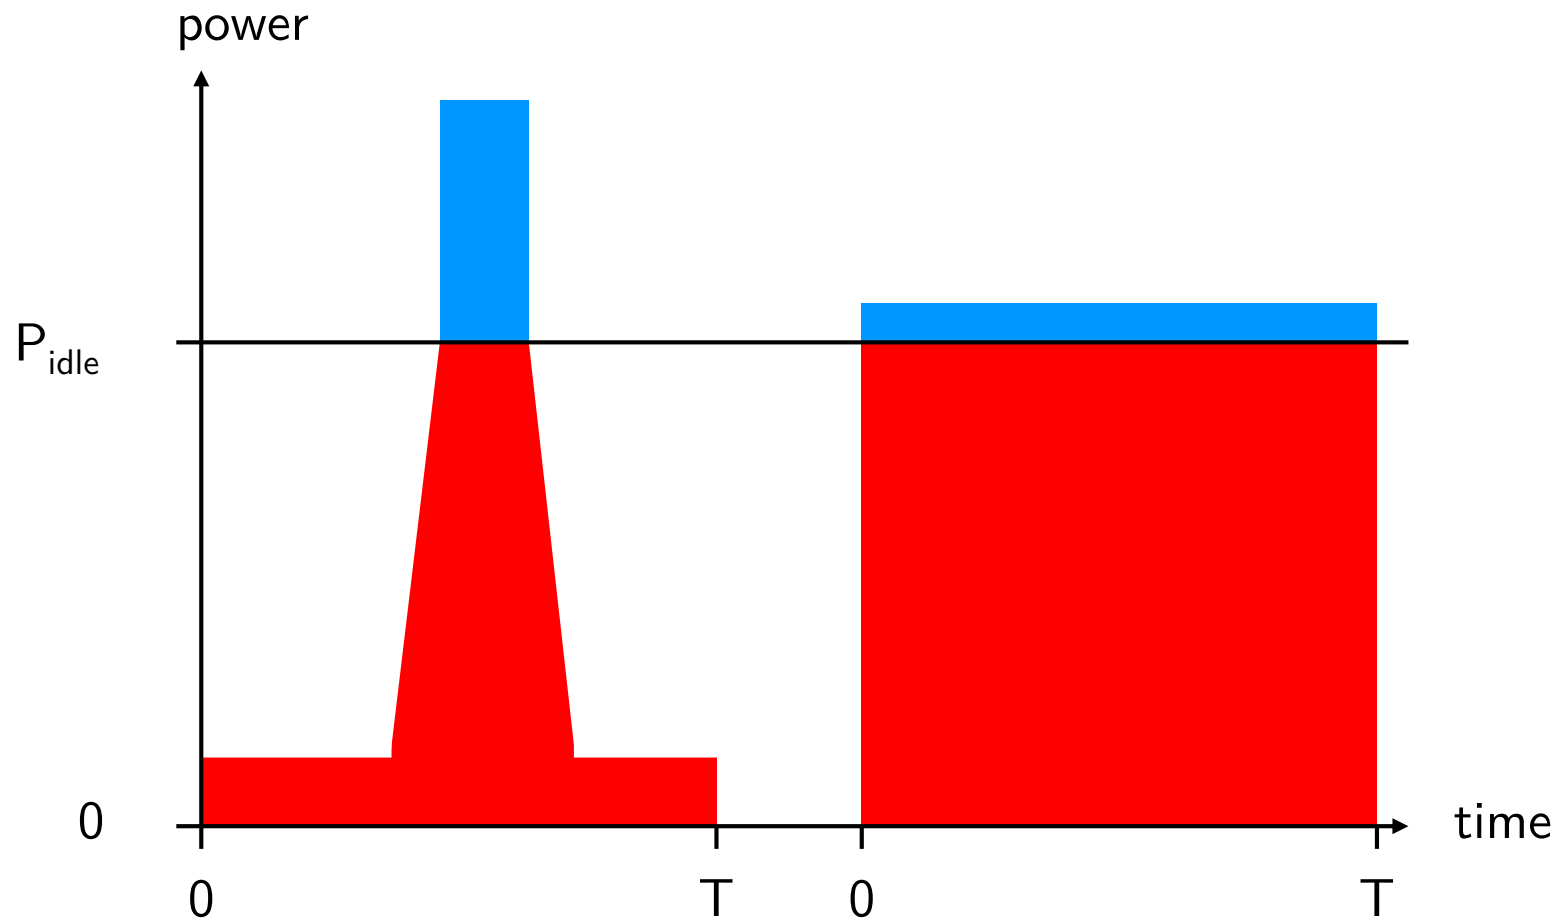
- As idle power dominates, low utilization wastes a lot.
- Reducing idle power yields better proportionality.
- Idle power is always there!

What does proportionality mean for our toy example?



- As idle power dominates, low utilization wastes a lot.
- Reducing idle power yields better proportionality.
- Idle power is always there!
- ... and it dominates.


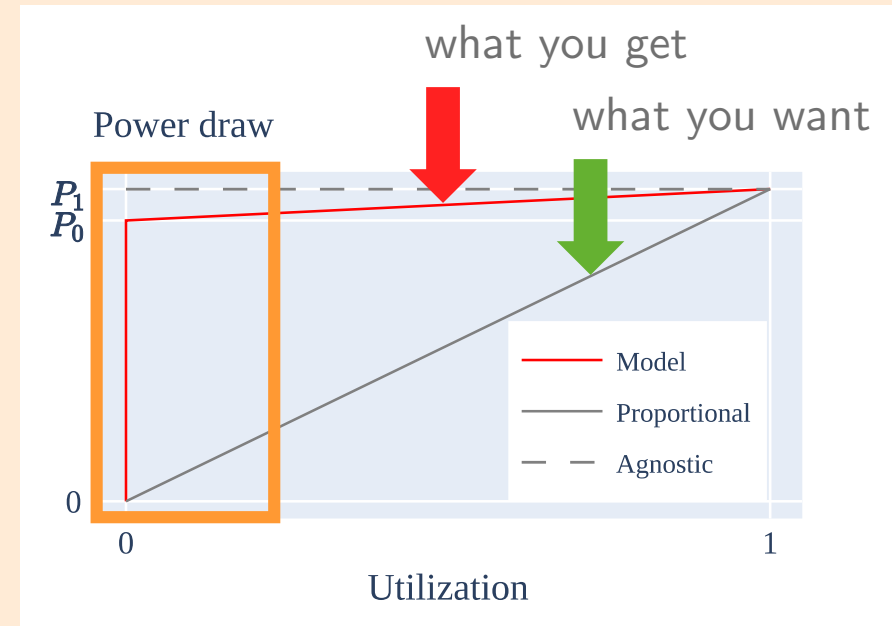
What does proportionality mean for our toy example?



- As idle power dominates, low utilization wastes a lot.
- Reducing idle power yields better proportionality.
- Idle power is always there!
- ... and it dominates.
- ▶ Improving proportionality is essentially about taking the “average idle power” down.

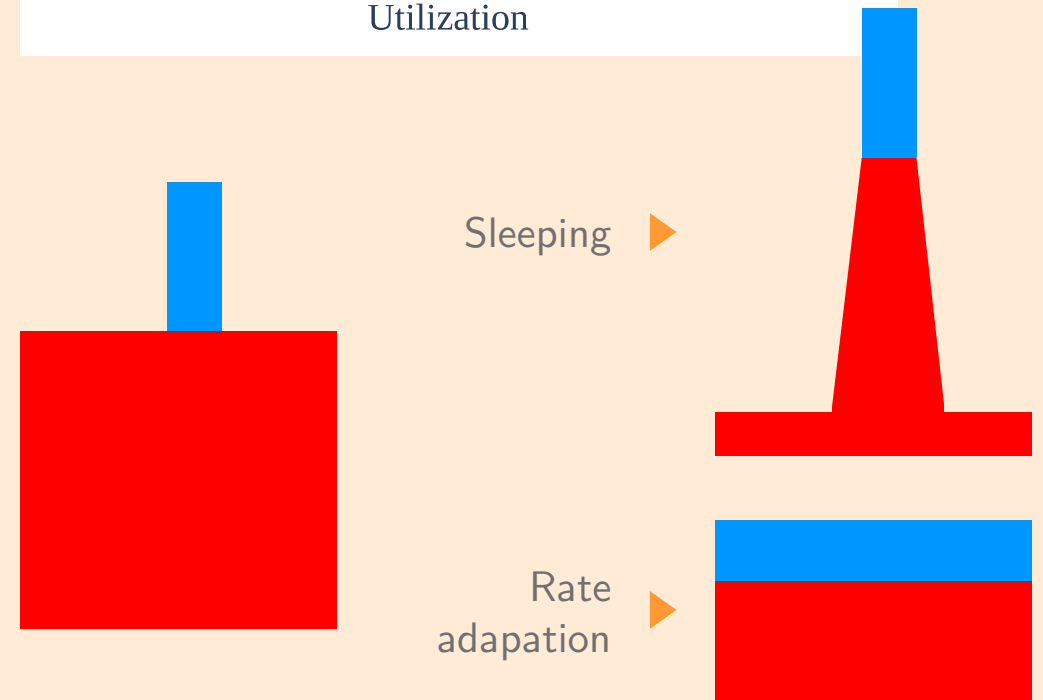
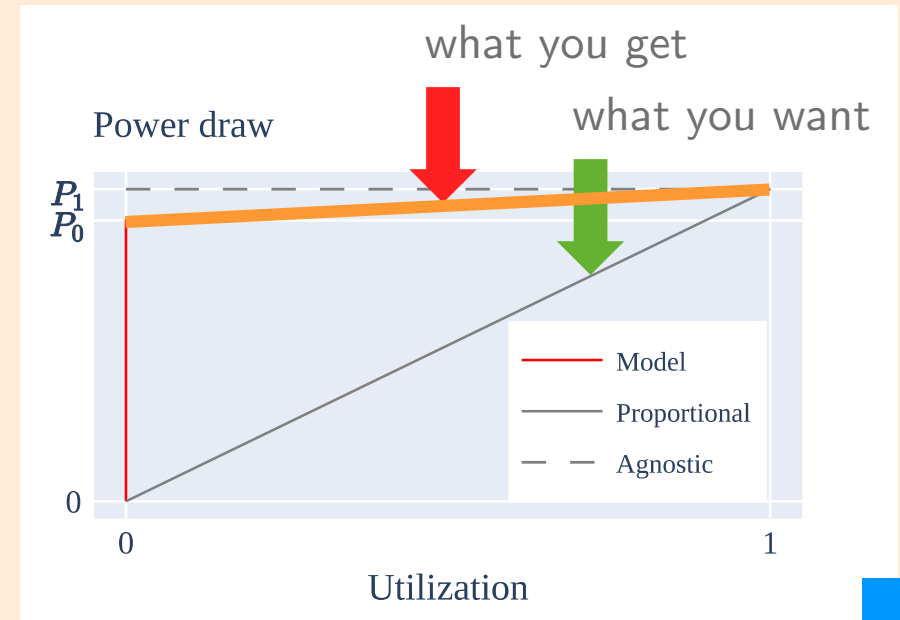
There two ways to improve energy efficiency.

- Run more often at high utilization
 - ▶ Better efficiency
 - Increase in total energy...



There two ways to improve energy efficiency.

- Run more often at high utilization
 - ▶ Better efficiency
 - ▶ Increase in total energy...
- Take low-utilization power down



What can be turned off?



- Ports
- Line cards [set of ports]
- Switch ...

- Memory banks
- Power supplies
- LEDs ... etc.

What can be turned off?

It can be more subtle than on/off.

- Change a port rate from 100G to 10G
- Down-clock the ASIC
- Cache frequently used FIB entries

- Ports
- Line cards [set of ports]
- Switch ...
- Memory banks
- Power supplies
- LEDs ... etc.

What can be turned off?

It can be more subtle than on/off.

- Change a port rate from 100G to 10G
- Down-clock CPUs
- Cache frequently used FIB entries

■ Ports

- Line cards [set of ports]
- Switch ...
- Memory banks
- Power supplies
- LEDs ... etc.

The theory says we can save tens of energy % in low-utilization networks.

[NSDI 2008]

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

Sergiu Nedeveschi[†] Lucian Popa^{*†} Gianluca Iannaccone[†]
Sylvia Ratnasamy[†] David Wetherall^{‡§}

Abstract

We present the design and evaluation of two forms of power management schemes that reduce the energy consumption of networks. The first is based on putting network components to sleep during idle times, reducing energy consumed in the absence of packets. The second is based on adapting the rate of network operation to the offered workload, reducing the energy consumed when actively processing packets.

For real-world traffic workloads and topologies and using power constants drawn from existing network equipment, we show that even simple schemes for sleeping or rate-adaptation can offer substantial savings. For instance, our practical algorithms stand to halve energy consumption for lightly utilized networks (10-20%). We show that these savings approach the maximum achievable by any algorithms using the same power management primitives. Moreover this energy can be saved without noticeably increasing loss and with a small and controlled increase in latency (<10ms). Finally, we show that both sleeping and rate adaptation are valuable depending (primarily) on the power profile of network equipment and the utilization of the network itself.

1 Introduction

In this paper, we consider power management for networks from a perspective that has recently begun to receive attention: the conservation of energy for operating and environmental reasons. Energy consumption in network exchanges is rising as higher capacity network equipment becomes more power-hungry and requires greater amounts of cooling. Combined with rising energy costs, this has made the cost of powering network exchanges a substantial and growing fraction of the total cost of ownership – up to half by some estimates[23]. Various studies now estimate the power usage of the US network infrastructure at between 5 and 24 TWh/year[25, 26], or \$0.5-2.4B/year at a rate of \$0.10/KWh, depending on what is included. Public


via standards such as EnergyStar. In fact, EnergyStar standard proposals for 2009 discuss slower operation of network links to conserve energy when idle. A new IEEE 802.3az Task Force was launched in early 2007 to focus on this issue for Ethernet [15].

Fortunately, there is an opportunity for substantial reductions in the energy consumption of existing networks due to two factors. First, networks are provisioned for worst-case or busy-hour load, and this load typically exceeds their long-term utilization by a wide margin. For example, measurements reveal backbone utilizations under 30% [16] and up to hour-long idle times at access points in enterprise wireless networks [17]. Second, the energy consumption of network equipment remains substantial even when the network is idle. The implication of these factors is that *most* of the energy consumed in networks is wasted.

Our work is an initial exploration of how overall network energy consumption might be reduced without adversely affecting network performance. This will require two steps. First, network equipment ranging from routers to switches and NICs will need power management primitives at the hardware level. By analogy, power management in computers has evolved around hardware support for *sleep* and *performance* states. The former (e.g., C-states in Intel processors) reduce idle consumption by powering off sub-components to different extents, while the latter (e.g., SpeedStep, P-states in Intel processors) tradeoff performance for power via operating frequency. Second, network protocols will need to make use of the hardware primitives to best effect. Again, by analogy with computers, power management preferences control how the system switches between the available states to save energy with minimal impact on users.

Of these two steps, our focus is on the network protocols. Admittedly, these protocols build on hardware support for power management that is in its infancy for networking equipment. Yet the necessary support will readily be deployed in networks where it makes

Energy Savings (%)



Let us consider both sleeping and rate adaptation.

Energy model

$$E = p_a T_a + p_i T_i$$

Effect of
... sleeping



- p_i reduces to $p_s = \gamma \cdot p_i$
- p_a is unchanged

... rate adaptation

- reduces both p_i and p_a
- T_a increases

Switching penalty

δ , in time [no energy penalty]



The sleeping approach is to “buffer-and-burst” packets.

Buffer-and-burst from the network edges


- Ingress router bundles same-destination packets together
- Buffers packets for some time then sends everything until the buffer is empty
- Turn off the link in-between bursts

Parameters


Defaults

- δ wake up delay $1ms$
- B buffering time $10ms$


Buffer-and-Burst



Classical forwarding

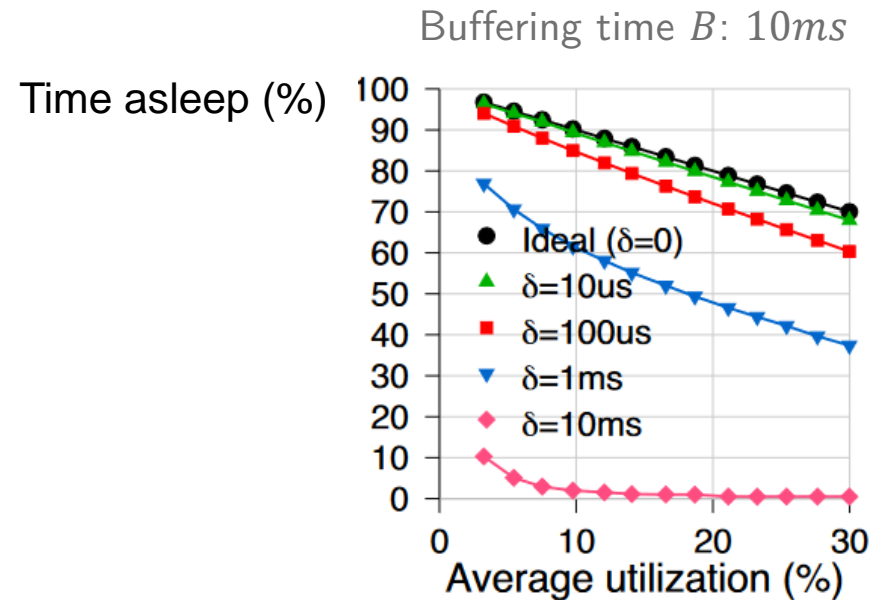


As expected, the wake-up and buffering times strongly affect the scheme's performance.

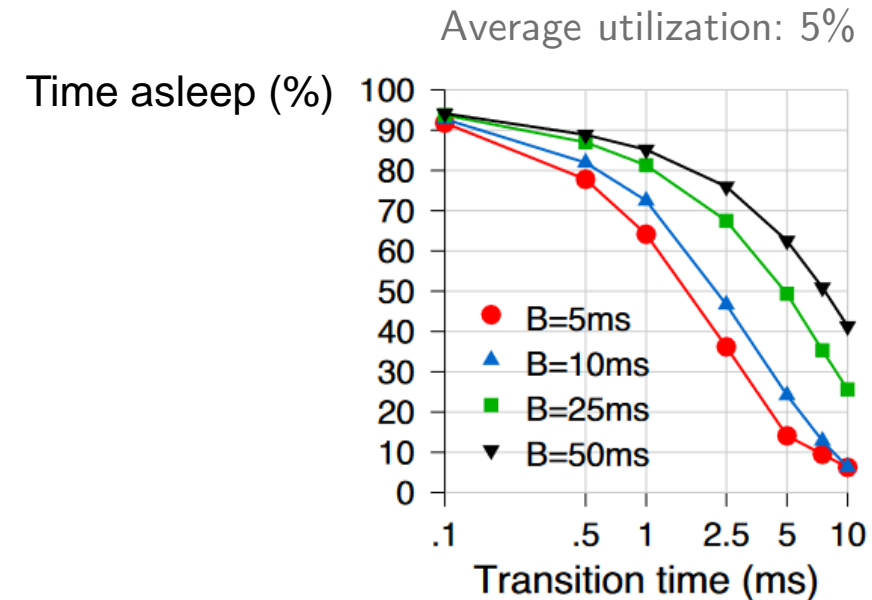


- ▶ Faster wake-up delay allows more time asleep
- ▶ Buffering must be longer than the wake-up delay...

As expected, the wake-up and buffering times strongly affect the scheme's performance.




- ▶ Faster wake-up delay allows more time asleep
- ▶ Buffering must be longer than the wake-up delay...



- ▶ Longer buffering allows more time asleep

Wake-up delay δ : 1ms

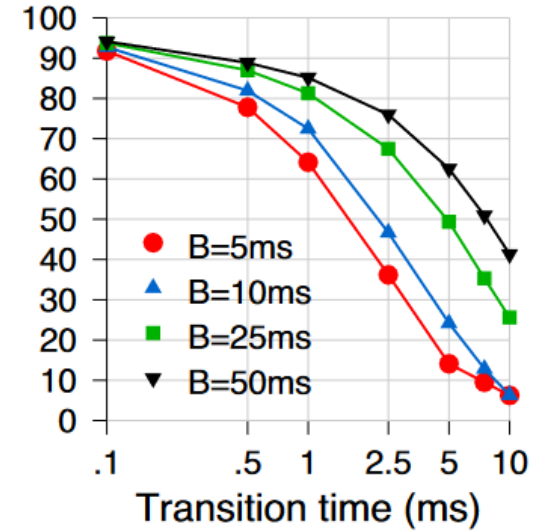
Ave. delay (ms)



▶ Longer buffering induces more packet delays!

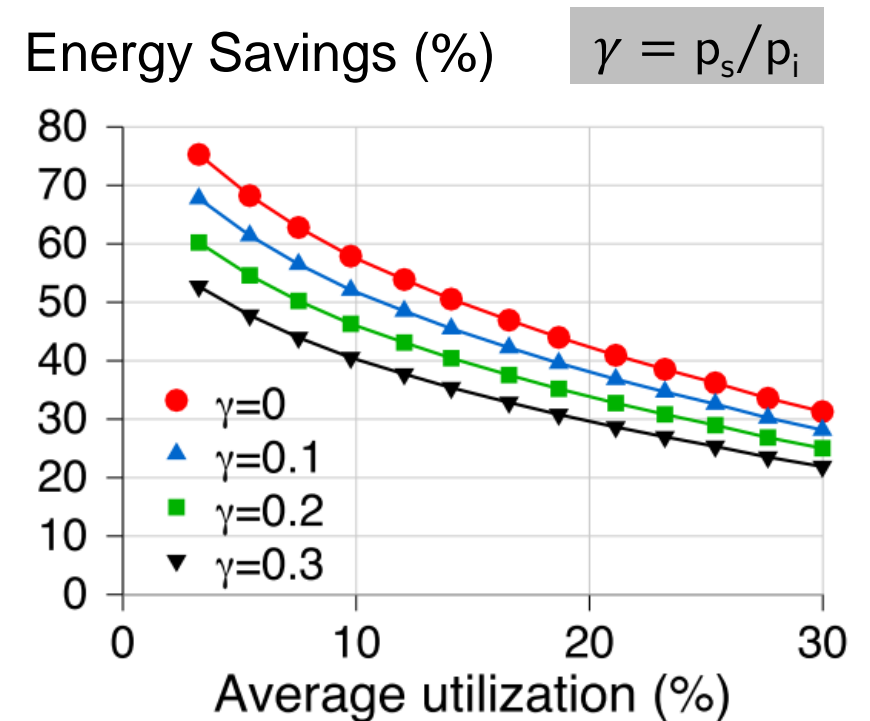
Average utilization: 5%

Time asleep (%)



▶ Longer buffering allows more time asleep but...

Assuming we can implement it,
buffer-and-burst promise sizable savings.




Does it work in practice?

How much does it save?

How fast can we wake-up?

Energy Savings (%)

$$\gamma = p_s / p_i$$




Does it work in practice?

How much does it save?



How fast can we wake-up?

Energy Savings (%)

$$\gamma = p_s / p_i$$



Turning ports off improves efficiency but we are still far from proportional.



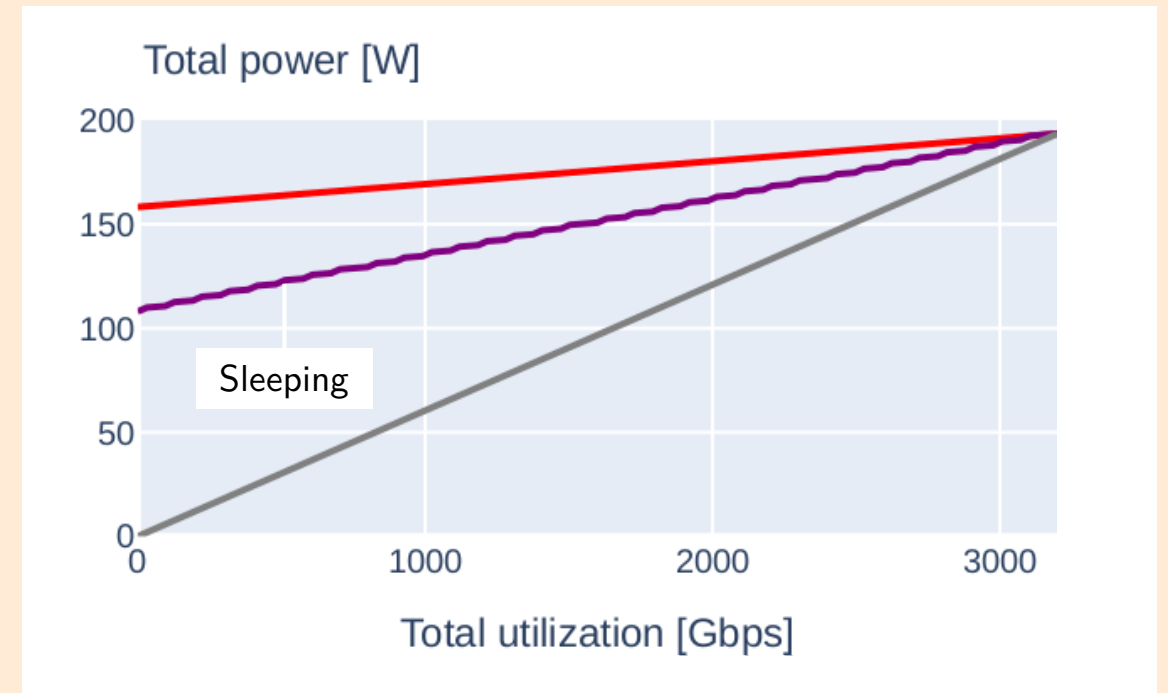
Still far from proportional but already much better!

However

This implicitly assumes that, either



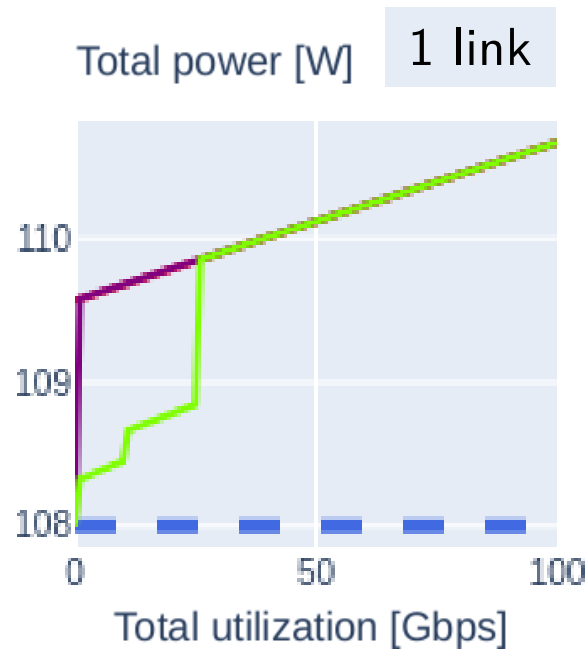

- All ports are connecting to the same endpoint,
- We normally keep useless ports on.

▶ Neither is very realistic.



▶ Still far from proportional but already much better!

Sleeping and rate adaptation save even with only a few parallel links.

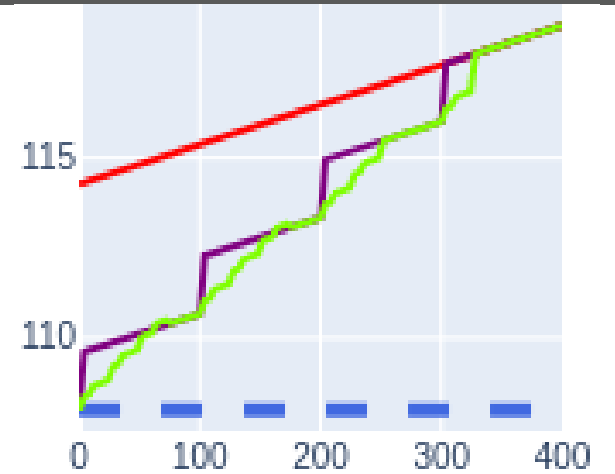
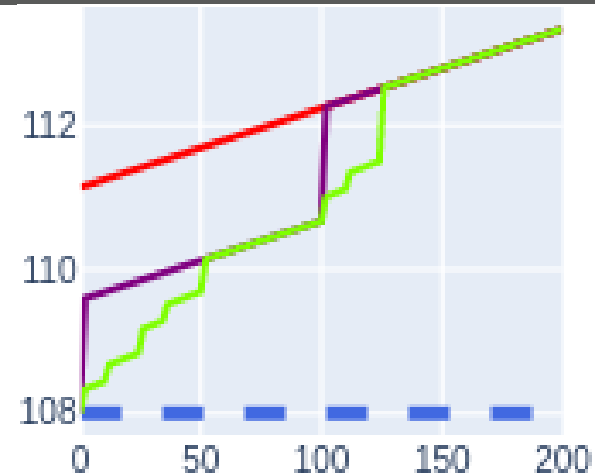



- Cannot sleep
- Down-rating helps

- Sleeping is simple and effective
- Down-rating helps further

- The more parallel links, the more possible savings

The power savings are **only a few watts** though it depends on the transceivers.



Does it work in practice?

How much does it save?

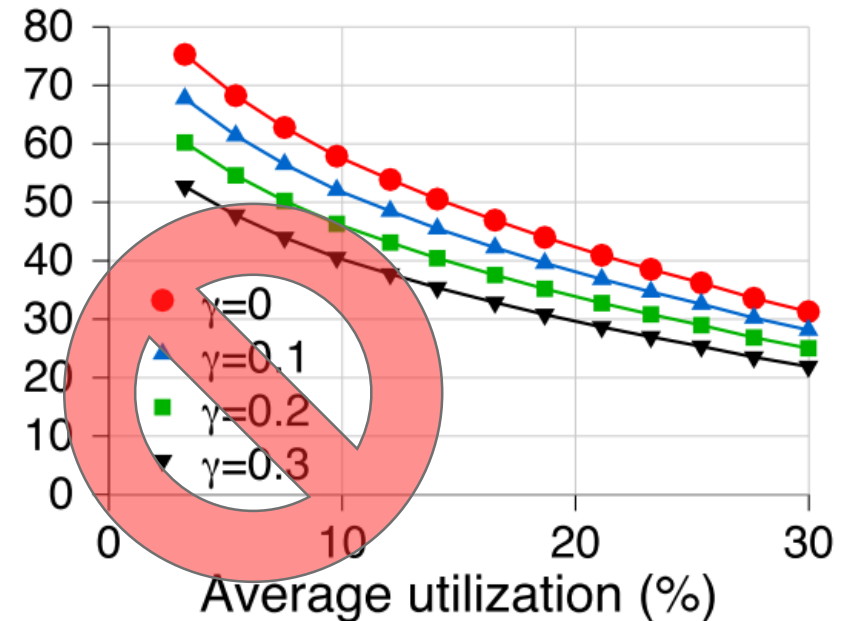
How fast can we wake-up?



If sleeping ports only, energy savings are small.

Energy Savings (%)

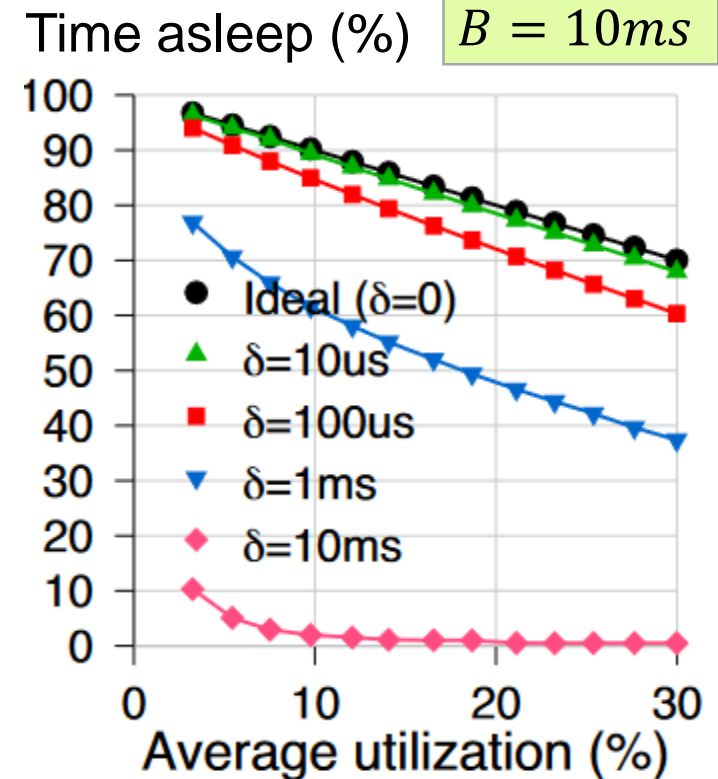
$$\gamma = p_s / p_i$$



Does it work in practice?

How much does it save?

How fast can we wake-up?



▶ Buffering must be longer than the wake-up delay

The real challenge is to turn things back on.

It's **easy** to turn off things that are never used.

▶ It's **harder** to turn off something that **you usually don't need**, but that you may, eventually.

[Cisco 8800 power post]

- Ports
- Line cards [set of ports]
- Switch ...

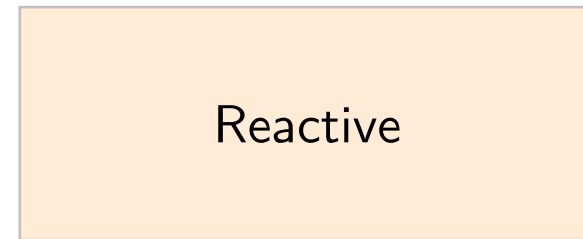
- Memory banks
- Power supplies
- LEDs ... etc.

To turn back on efficiently, we need either good reaction time or prediction. Ideally, both.

Objective


Stuff must be ON when you need it.

Two approaches



Be quick at turning on and switch on-demand

▶ Phone screens




Accurately predict future demand and switch on early enough

▶ Electricity grid

Theory

These findings reinforce our intuition that hardware support featuring low-power sleep states and quick transitions (preferably $< 1ms$) between these states are essential to effectively save energy.

Practice



We cannot turn on transceiver at traffic timescales (today).

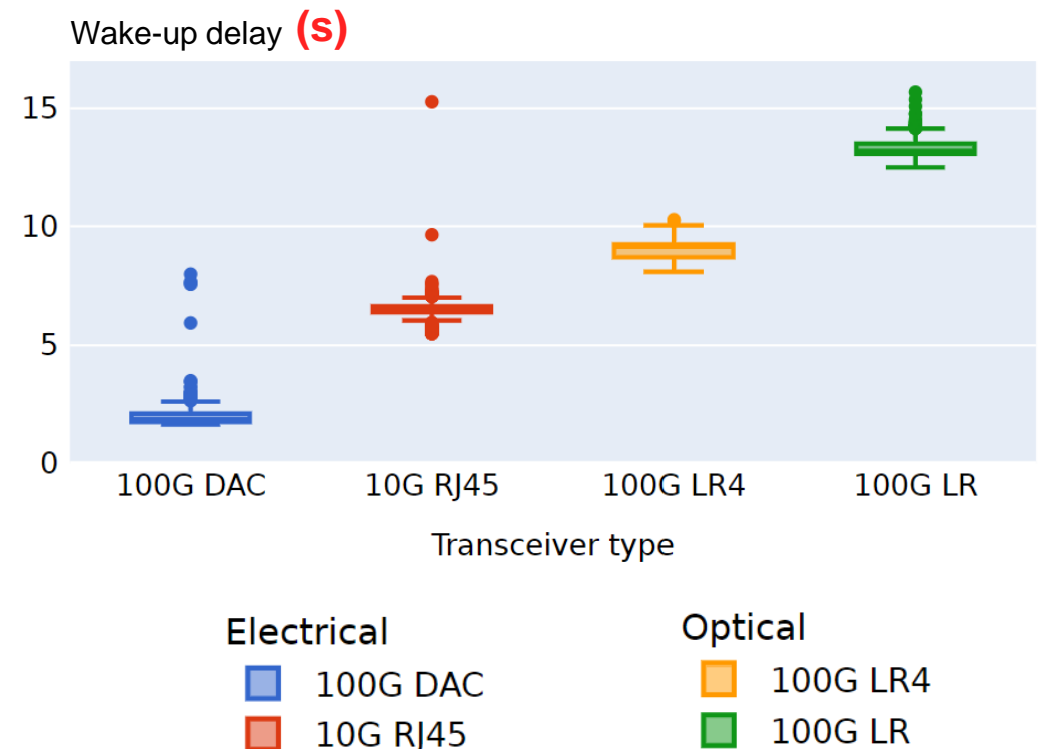
Theory

These findings reinforce our intuition that hardware support featuring low-power sleep states and quick transitions (preferably $< 1ms$) between these states are essential to effectively save energy.



Today's transceivers are **1000x slower to start** than required for "effective savings"

Practice



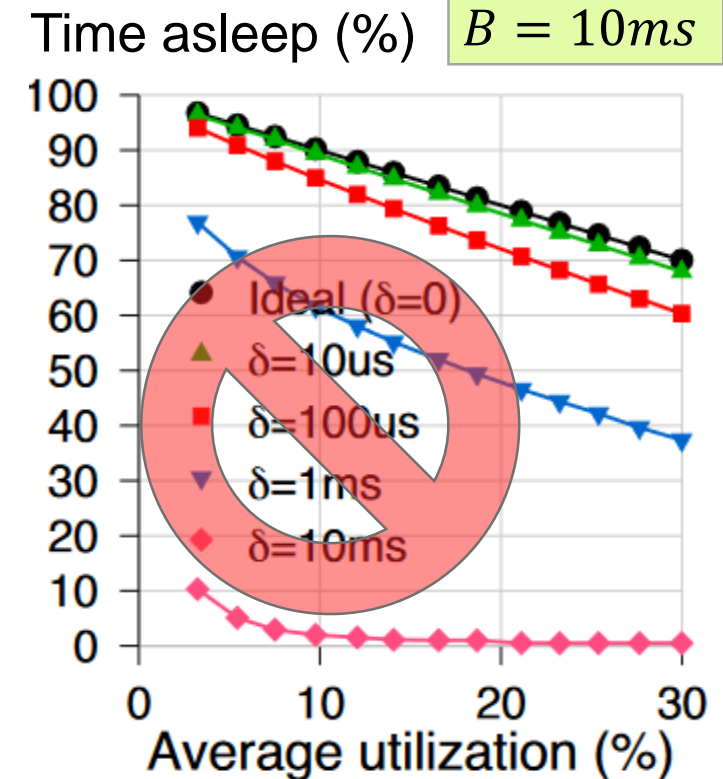
Does it work in practice?

How much does it save?

How fast can we wake-up?



The wake-up delay is too long to buffer.



Buffering must be longer than the wake-up delay

Turning back on efficiently
is **really hard** today. Ideally,

Reactive

Be quick at turning on
and switch on-demand

▶ 1000x too slow

Proactive

Turning back on efficiently
is **really hard** today.


Reactive

Proactive

Accurately predict the
future traffic demand

▶ Known to be
hard in networks

Turning back on efficiently is really hard today... unless we **change timescales!**




What keeps your network up at night?

Lukas Röllin, Romain Jacob, Laurent Vanbever

Observation

Network links are **underutilized, power-hungry** and **inefficient**



Theory


Save energy with sleeping and buffering

Assumption: Transceiver ready within milliseconds




Practice

Transceiver **wake-up** takes **seconds!**




Contribution

Turning links off still works when considering longer timeframes




No disruption to the network if the traffic doesn't change too fast




Result

TCP limits the impact of congestion if traffic changes too fast



Future

Faster wake-up **boosts energy savings** and **reduces performance impact**



Operational
Carbon efficiency

=

$$\frac{\text{J used}}{\text{Task}}$$

×

$$\frac{\text{J supplied}}{\text{J used}}$$

×

$$\frac{\text{Carbon}}{\text{J supplied}}$$

... improves
with better...


Networks
& compute


Infrastructure
& HW design

Application
elasticity

Not really a
network matter

Most of the energy inefficiency comes from cooling




 Numbers are old!

[Data Center Energy, 2016]

Most of the energy inefficiency comes from cooling ... so cool less

[SIGMETRICS 2012]

 **Temperature Management in Data Centers:
Why Some (Might) Like It Hot**

Nosayba El-Sayed Ioan Stefanovici George Amvrosiadis Andy A. Hwang
Bianca Schroeder
Department of Computer Science
University of Toronto
{nosayba, ioan, gamvrosi, hwang, bianca}@cs.toronto.edu

ABSTRACT
The energy consumed by data centers is starting to make up a significant fraction of the world's energy consumption and carbon emissions. A large fraction of the consumed energy is spent on data center cooling, which has motivated a large body of work on temperature management in data centers. Interestingly, a key aspect of temperature management has not been well understood: controlling the setpoint temperature at which to run a data center's cooling system. Most data centers set their thermostat based on (conservative) suggestions by manufacturers, as there is limited understanding of how higher temperatures will affect the system. At the same time, studies suggest that increasing the temperature setpoint by just one degree could save 2-5% of the energy consumption.
This paper provides a multi-faceted study of temperature management in data centers. We use a large collection of field data from different production environments to study the impact of temperature on hardware reliability, including the reliability of the storage subsystem, the memory subsystem and server reliability as a whole. We also use an experimental testbed based on a thermal chamber and a large array of benchmarks to study two other potential issues with higher data center temperatures: the effect on server performance and power. Based on our findings, we make recommendations for temperature management in data centers, that create the potential for saving energy, while limiting negative effects on system reliability and performance.

Categories and Subject Descriptors
B.8 [Hardware]: Performance and Reliability—Temperature; C.4 [Computer Systems Organization]: Performance of Systems—Temperature

Keywords
Data Center, Temperature, Reliability, Performance, Energy, LSE, Hard Drive, Memory, DRAM, CPU, Fans

1. INTRODUCTION
Data centers have developed into major energy hogs. The world's data centers are estimated to consume power equivalent to about seventeen 1,000 MW power plants, equaling more than 1% of total world electricity consumption, and to emit as much carbon dioxide as all of Argentina [17]. More than a third, sometimes up to one half of a data center's electricity bill is made up by electricity for cooling [6, 19]. For instance, for a data center consisting of 30,000 square feet and consuming 10MW, the yearly cost of running the cooling infrastructure can reach up to \$4-8 million [23].
Not surprisingly, a large body of research has been devoted to reducing cooling cost. Approaches that have been investigated include, for example, methods to minimize air flow inefficiencies [23, 35], load balancing and the incorporation of temperature awareness into workload placement in data centers [7, 25, 28, 33], and power reduction features in individual servers [14, 15].
Interestingly, one key aspect in the thermal management of a data center is still not very well understood: controlling the setpoint temperature at which to run a data center's cooling system. Data centers typically operate in a temperature range between 20C and 22C, some are as cold as 13C degrees [8, 29]. Due to lack of scientific data, these values are often chosen based on equipment manufacturers' (conservative) suggestions. Some estimate that increasing the setpoint temperature by just one degree can reduce energy consumption by 2 to 5 percent [8, 9]. Microsoft reports that raising the temperature by two to four degrees in one of its Silicon Valley data centers saved \$250,000 in annual energy costs [29]. Google and Facebook have also been considering increasing the temperature in their data centers [29].
While increasing data center temperatures might seem like an easy way to save energy and reduce carbon emissions, it comes with some concerns, the most obvious being its impact on system reliability. Unfortunately, the details of how increased data center temperatures will affect hardware reliability are not well understood and existing evidence is contradicting. A recent study [35] indicated that in order to avoid thermal redlining, a typical server needs to have the air-temperature at its front inlets be in the range of 20C.

Empirical study

- Effects of temperature on reliability are less dramatic than theory suggests
- Temperature variation is more harmful than the temperature level
- Temperature was poorly measured and managed (back then)



DC used to run at 18°C.
Today they often run at 26°C

[Google | DC efficiency]

Most of the energy inefficiency comes from cooling ... so cool less or reuse the heat.



[OCP | Heat reuse]

Operational
Carbon efficiency

=

$$\frac{J \text{ used}}{\text{Task}}$$

×

$$\frac{J \text{ supplied}}{J \text{ used}}$$

×

$$\frac{\text{Carbon}}{J \text{ supplied}}$$

... improves
with better...

Networks
& compute

Infrastructure
& HW design


Application
elasticity



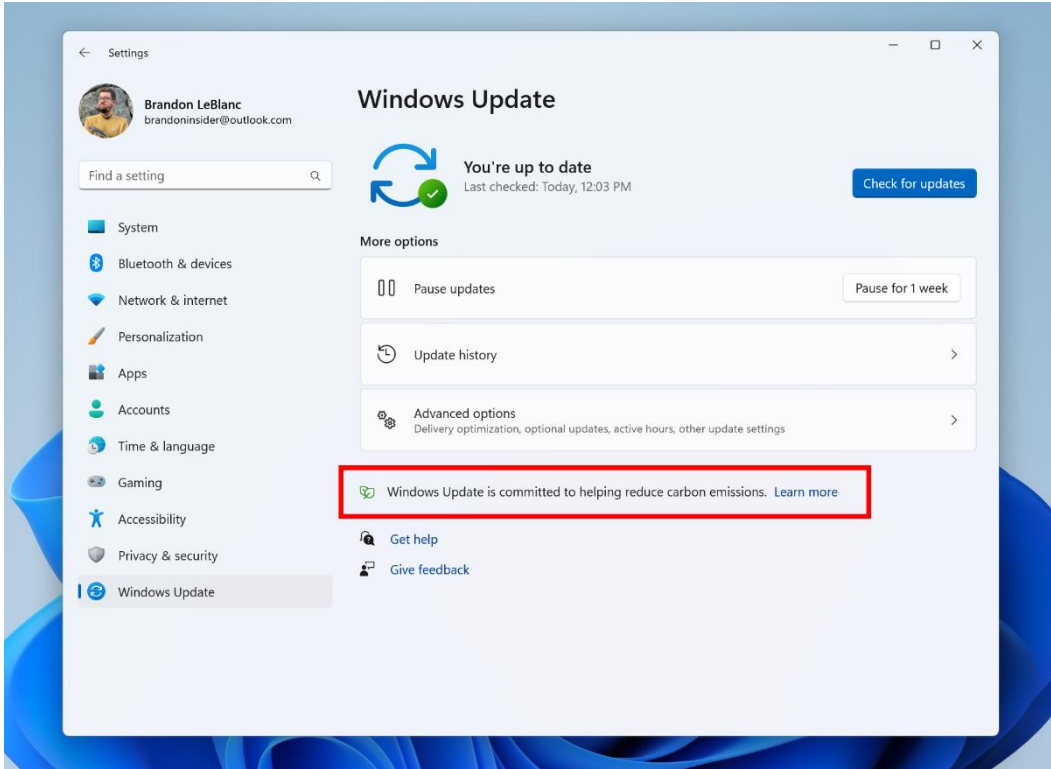
Improves with
better networks

One may reduce the carbon impact by working **when** and **where** the energy is clean.

time-shifting space-shifting



[Google | Time-shifting post]



[Carbon-aware Windows update]

For networks, this idea translates to choosing “greener” paths.

[e-Energy 2023]

[CoNEXT 2023]



Carbon-Aware Global Routing in Path-Aware Networks

Seyedali Tabaeiaghdaei
ETH Zürich

Jonghoon Kwon
ETH Zürich

Simon Scherrer
ETH Zürich

Adrian Perrig
ETH Zürich

ABSTRACT

The growing energy consumption of Information and Communication Technology (ICT) has raised concerns about its environmental impact. However, the carbon footprint of data transmission over the Internet has so far received relatively modest attention. This carbon footprint can be reduced by sending traffic over carbon-efficient inter-domain paths. However, challenges in estimating and disseminating carbon intensity of inter-domain paths have prevented carbon-aware path selection from becoming a reality.

In this paper, we take advantage of path-aware network architectures to overcome these challenges. In particular, we design CIRo, a system for forecasting the carbon intensity of inter-domain paths and disseminating them across the Internet. We implement a proof of concept for CIRo on the codebase of the SCION path-aware Internet architecture and test it on the SCIONLab global research testbed. Further, through large-scale simulations, we demonstrate the potential of CIRo for reducing the carbon footprint of endpoints and end domains: With CIRo, half of domain pairs can reduce the carbon intensity of their inter-domain traffic by at least 47%, and 87% of end domains can reduce their carbon footprint of Internet use by at least 50%.

CCS CONCEPTS

• Applied computing → Forecasting; Multi-criterion optimization and decision-making; • Networks → Network measurement; Network simulations; Network performance modeling; Data path algorithms; Control path algorithms; Topology analysis and generation; • Hardware → Renewable energy.

KEYWORDS

Green Networking, Internet Carbon-Emission Modeling and Measurement, Inter-Domain Routing, Carbon-Aware Routing, SCION

ACM Reference Format:

Seyedali Tabaeiaghdaei, Simon Scherrer, Jonghoon Kwon, and Adrian Perrig. 2023. Carbon-Aware Global Routing in Path-Aware Networks. In *The 14th ACM International Conference on Future Energy Systems (e-Energy '23)*, June 20–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3575813.3595192>

1 INTRODUCTION

Growing concerns regarding climate change encourage companies to measure and reduce their carbon footprint, i.e., the amount of carbon emission that can be attributed to them. This also applies to their use of Information and Communication Technology (ICT), as ICT has a notable contribution of 2.7% to global CO₂ emissions [39], which is expected to grow significantly – approximately four times – until 2030 [3]. Hence, reducing the carbon footprint of ICT use is becoming increasingly relevant for enterprises, manifesting in carbon-neutrality statements of major technology corporations.

While these efforts are laudable and impactful, promising opportunities for further carbon-footprint reduction exist. Indeed, previous research has identified a range of such opportunities. However, most of these proposals apply to local aspects: intra-domain networking (i.e., within a single domain), data-center optimizations, or neighbor-domain cooperation (cf. §8). In contrast, inter-domain networking (i.e., among multiple domains), which accounts for around 13% of total ICT energy consumption, has so far received less attention. An exception is the work by Zilberman et al. [70], who identify carbon-aware networking as a high-potential research area and sketch the concept of carbon-intelligent routing, i.e., to leverage differences in network paths’ carbon intensity (i.e., carbon emission per unit of data transmitted) to reduce the carbon footprint of communications.

Previous research on green inter-domain networking applies carbon efficiency to the optimization metric of the Border Gateway Protocol (BGP) [42]. Unfortunately, this direction faces several challenges. **Inefficient Green Route:** A strict carbon-optimal path can result in a highly inefficient end-to-end path in terms of monetary cost, latency, bandwidth, loss, or jitter (cf. §7.1). Depending on the application requirements, an optimization subject to all these constraints needs to be made, requiring path selection within a fine-grained metric space. **Ossification:** Carbon-optimal paths can thus only be offered as additional options, not as replacements for the conventional BGP route. When using BGP to provide carbon-efficient alternative paths, routers would thus require multiple forwarding tables, and packets would need to indicate the desired optimization criteria. Updating BGP and router hardware represents a challenge – as we have experienced in securing the BGP

Exploring the Benefits of Carbon-Aware Routing

SAWSAN EL-ZAHR, Department of Engineering Science, University of Oxford, UK
PAUL GUNNING, BT Research & Network Strategy, UK
NOA ZILBERMAN, Department of Engineering Science, University of Oxford, UK

Carbon emissions associated with fixed networks can be significant. However, accounting for these emissions is hard, requires changes to deployed equipment, and has contentious benefits. This work sheds light on the benefits of carbon aware networks, by exploring a set of potential carbon-related metrics and their use to define link-cost in carbon-aware link-state routing algorithms. Using realistic network topologies, traffic patterns and grid carbon intensity, we identify useful metrics and limitations to carbon emissions reduction. Consequently, a new heuristic carbon-aware traffic engineering algorithm, CATE, is proposed. CATE takes advantage of carbon intensity and routers’ dynamic power consumption, combined with ports power down, to minimize carbon emissions. Our results show that there is no silver bullet to significant carbon reductions, yet there are promising directions without changes to existing routers’ hardware.

1 INTRODUCTION

The fast development and deployment of the Internet has widely focused on reliability, scalability, speed and security. Starting in 2001, many initiatives tackled the power efficiency of Information and Communications Technology (ICT) for wireless networks [12, 48] and then in 2003 for wired networks [40]. In 2015, the Paris agreement set new sustainability goals of achieving 45% less carbon emissions by 2030 and reaching net zero by 2050 [54]. With this trend, and while ICT carbon footprint contributed to 2% of the overall carbon emissions in 2010 [23], ICT companies try to minimize their carbon emissions. Most works addressed data centers, improving power usage effectiveness (PUE) from 2.0 and above to the order of 1.1 for hyperscale data centers [39], as well as improving across all compute aspects: from CPU design to server, software and data center design.

Compared to data centers, fixed wired networks have seen limited improvement [69]. The improvements in this field are limited by the absence of standard power and carbon accounting metrics [19]. Fast technological advancements affect the contribution of different components within a router to the overall power consumption and thus, power metrics vary substantially with time. On the other hand, carbon metrics require visibility into the energy generation mix of the local power grid [28] which is difficult to integrate into the routing stack of deployed network elements [43].

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?



Engraving by [Edward Goodall](#) (1795-1870), original title Manchester, from Kersal Moor after a painting of W. Wyld

◀ Coal-burning factories in 19th-century Manchester, England.

Improved technology allowed coal to fuel the Industrial Revolution.

This greatly increased the consumption of coal.

Improving efficiency of a resource usage may result in increased consumption of that resource.



Engraving by [Edward Goodall](#) (1795-1870), original title Manchester, from Kersal Moor after a painting of W. Wylde

◀ Coal-burning factories in 19th-century Manchester, England.

Improved technology allowed coal to fuel the Industrial Revolution.

This greatly increased the consumption of coal.

▶ Known as the Jevons paradox or rebound effects

The Jevons paradox is observed in the ICT sector.

From 2007
to 2020 Progress in both
hardware and software

From 2015
to 2020 GHG emissions of ICT
increased by 5%

The Jevons paradox is observed in the ICT sector.

From 2007
to 2020

Progress in both
hardware and software

- Energy efficiency increased
 - Energy usage **per subscriber** increased
- ▶ Jevons paradox
on energy

From 2015
to 2020

GHG emissions of ICT
increased by 5%

The Jevons paradox is observed in the ICT sector.

From 2007
to 2020

Progress in both
hardware and software

- Energy efficiency increased
 - Electricity usage **per subscriber** increased
- ▶ Jevons paradox
on electricity

From 2015
to 2020

GHG emissions of ICT
increased by 5%

- More devices are being sold
 - Most consumers power devices
using carbon-intense electricity.
- ▶ Jevons paradox
on carbon

As “we” keep asking for more,
the energy use will keep rising.



Google Play Games Apps

ChatGPT

OpenAI
In-app purchases

4.6★
231K reviews

10M+
Downloads

12
PEGI 12 ⓘ

Install Share

This app is available for all of your devices




But wait, what about networks?



Didn't we say network power was inelastic anyway?

(I'm glad you asked)

The idle power dominates.
I.e., Network power is inelastic.



1. Power increases marginally with utilization.





- ISP overprovision networks to support
- Peak traffic
 - Fault tolerance

2. Average utilization is low in ISP networks.


There two ways to improve energy efficiency.

- Run more often at high utilization
- ▶ Better efficiency
- Increase in total energy...





3. Increasing utilization improves efficiency.

The idle power dominates.
I.e., Network power is inelastic.



1. Power increases marginally with utilization.




ISP overprovision networks to support

- Peak traffic
- Fault tolerance

2. Average utilization is low in ISP networks.


There two ways to improve energy efficiency.

- Run more often at high utilization
- ▶ Better efficiency
- Increase in total energy...

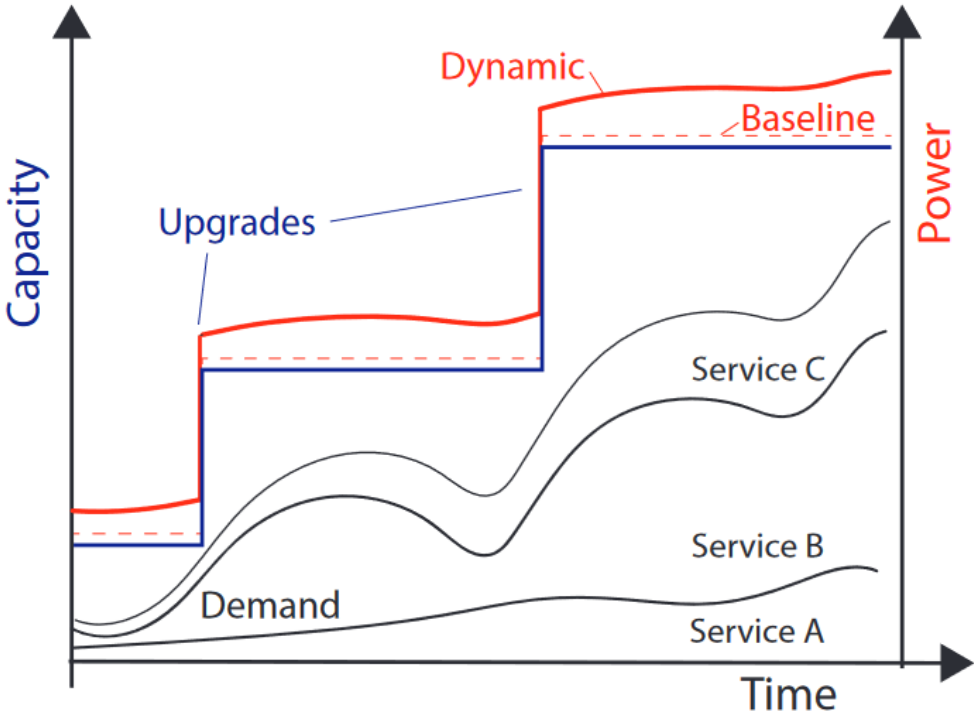



4. Networks are intentionally kept overprovisioned!
3. Increasing utilization improves efficiency.

There is a feedback loop that stimulates network capacity increase



There is a feedback loop that stimulates network capacity increase and energy usage.



We must embrace some **digital sobriety**.

Everything has a **cost**.

- Every picture we upload
- Every app we download
- Every movie we stream
- Every conversation we archive

It is not to say we must not do it
but

we must **be mindful** when doing it
and do it only when actually useful.

We must embrace some **digital sobriety**.

Everything has a cost.

- Every picture we upload
- Every app we download
- Every movie we stream
- Every conversation we archive

It is not to say we must not do it
but

we must **be mindful** when doing it,
and do it only when actually useful.

“We” also includes
the private sector...

1.3 million PB

According to the [World Economic Forum](#), companies generate **1.3 trillion gigabytes** of dark data **every day**. Storing that data for a year using non-renewables generates as much CO2 as three million flights from London to New York.

28×10^9

In 2020, Google said it stored four trillion photos, with **28 billion** new photos and videos uploaded **each week**.

[Dark data]

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

What's the carbon footprint of
one hour streaming Netflix?

55 gCO₂eq.

How do we
measure this?

How can
we improve?

Can technology
save us?

If you are interested in learning more or getting involved, here are some places to start.

Websites

- [ETH NetZero](#)
- [ETH Student Sustainability Committee](#)
- [Greening of Steaming](#)

Podcasts

- [Environement Variables](#)
- [Green I/O](#)
- [My Climate Journey](#)

Tools

- [Ecograder](#)
- [Website Carbon Calculator](#)
- [Electricity Map](#)

Blogs

- [Low-tech magazine](#)
- [Fershad Irani](#)


Research

- [HotCarbon workshop](#)
- [e-Energy conference](#)



I'm interested in your favorites if they are not here already!

If you are interested in learning even more we are offering **projects** on sustainable networking.



Visit [our group website](#) for details.

nsg.ee.ethz.ch

Sustainable Networking

The Internet has grown to become an indispensable part of our lives. It also consumes massive amounts of energy and is accountable for an increasingly large share of the worldwide carbon footprint—each in the order of a couple of percent! The question is not *if* but *how* we can improve Internet sustainability.

“If you are interested in a project or a doctoral position in that area, contact Romain; we often have more project ideas than what is being publicly offered on our website.”

Active Members



Prof. Laurent
Vanbever
Group Leader



Dr. Romain Jacob
Post-doc



You are welcome to come with your own ideas as well 😊

Who cares?


If you are interested, write me 😊



[Who cares?]

Computer Networks

Sustainable Networking



Climate stripes. Ed Hawkins, 2018
portrays the increase of average global temperature

Romain Jacob
jacobr@ethz.ch

I hope you enjoyed the lecture!
Let's improve networks together 😊

References 1

- [UN Sustainability definition] <https://www.un.org/en/academic-impact/sustainability>
- [Internet users] <https://ourworldindata.org/grapher/share-of-individuals-using-the-internet>
- [World energy production] <https://ourworldindata.org/grapher/electricity-prod-source-stacked>
- [IEA | DC and Telco] <https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks>
- [Ireland DC consumption] <https://www.cso.ie/en/releasesandpublications/ep/p-dcmec/datacentresmeteredelectricityconsumption2022/>
- [AWS buys diesel gen] <https://www.datacenterdynamics.com/en/news/aws-buys-105-back-up-diesel-generators-for-new-data-center-in-dublin/>
- [Planetary boundaries] <https://www.stockholmresilience.org/research/planetary-boundaries.html>
- [Does not compute] <https://doi.org/10.1016/j.joule.2021.05.007>
- [Netflix 2022 ESG report] https://downloads.ctfassets.net/4cd45et68cgf/7rnC6zK537cM8zAGrXA90E/3c654a2d0023a4dac26a20b2fff39855/Netflix_2022-ESG-Report-FINAL.pdf
- [GHG equivalencies calculator] <https://www.epa.gov/energy/greenhouse-gas-equivalencies-calculator>
- [Methane emissions] https://energy.ec.europa.eu/topics/oil-gas-and-coal/methane-emissions_en

References 2

- [GSF practioner course] <https://learn.greensoftware.foundation/>
- [Rel. energy prod. | World] <https://ourworldindata.org/grapher/electricity-prod-source-stacked?stackMode=relative>
- [Rel. energy prod. | CH] <https://ourworldindata.org/grapher/electricity-prod-source-stacked?stackMode=relative&country=~CHE>
- [LCA] <https://www.oneclicklca.com/life-cycle-assessment-explained/>
- [Ericsson ICT report] <https://www.ericsson.com/en/reports-and-papers/industrylab/reports/a-quick-guide-to-your-digital-carbon-footprint>
- [Consequential LCA] <https://consequential-lca.org/>
- [GHG protocol] <https://ghgprotocol.org/>
- [GSF procurement post] <https://www.thegreenwebfoundation.org/news/place-policy-procurement-more-levers-for-a-lower-carbon-internet/>
- [HotCarbon 2023, A] <https://dl.acm.org/doi/10.1145/3604930.3605710>
- [DIMPACT] <https://dimpact.org/>
- [Uptime | PUE goes up] <https://journal.uptimeinstitute.com/is-pue-actually-going-up/>
- [Data Center Energy, 2016] <https://doi.org/10.1109/COMST.2015.2481183>

References 3

- [Google | DC efficiency] <https://www.google.com/about/datacenters/efficiency/>
- [Electricity Maps] <https://app.electricitymaps.com/zone/CH>
- [Carbon Trust report] <https://ctprodstorageaccountp.blob.core.windows.net/prod-drupal-files/documents/resource/public/Carbon-impact-of-video-streaming.pdf>
- [SIGCOMM 2003] <https://doi.org/10.1145/863955.863959>
- [SWITCH weathermap] <https://network.switch.ch/pub/swiss-map/>
- [HotCarbon 2023, B] <https://doi.org/10.1145/3604930.3605713>
- [NSDI 2008] <https://www.usenix.org/conference/nsdi-08/reducing-network-energy-consumption-sleeping-and-rate-adaptation>
- [Cisco 8800 power post] <https://xrdocs.io/8000/blogs/cisco-8800-power-provisioning/>
- [Lukas | Master thesis] <https://doi.org/10.3929/ethz-b-000637984>
- [Lukas | Poster] <https://dl.acm.org/doi/10.1145/3624354.3630092>
- [SIGMETRICS 2012] <https://doi.org/10.1145/2254756.2254778>
- [OCP | Heat reuse] <https://www.opencompute.org/projects/heat-reuse>

References 4

- [Google | Time-shifting post] <https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/>
- [Carbon-aware Windows update] <https://arstechnica.com/gadgets/2022/03/microsoft-is-trying-to-lower-carbon-emissions-via-windows-update-of-all-things/>
- [e-Energy 2023] <https://dl.acm.org/doi/10.1145/3575813.3595192>
- [CoNEXT 2023] https://ora.ox.ac.uk/objects/uuid:504d0a87-6742-414f-b2c2-be74f2e5b579/download_file?file_format=application%2Fpdf&safe_filename=El_Zahr_et_al_2023_exploring_the_benefits.pdf&type_of_work=Journal+article
- [Wikipedia | Jevons paradox] https://en.wikipedia.org/wiki/Jevons_paradox
- [SSRN 2023] <https://ssrn.com/abstract=4424264>
- [CHI 2016] <https://doi.org/10.1145/2858036.2858378>
- [IAB 2022] <https://research-information.bris.ac.uk/en/publications/rethinking-allocation-in-high-baseload-systems-a-demand-proportio>
- [Dark data] <https://www.datacenterdynamics.com/en/opinions/the-elephant-in-the-data-center-shedding-light-on-dark-data/>
- [Who cares?] <https://whocares.ethz.ch/>
- [Climate stipes] <https://www.reading.ac.uk/planet/climate-resources/climate-stripes>