Tomorrow's Internet must sleep more and grow old

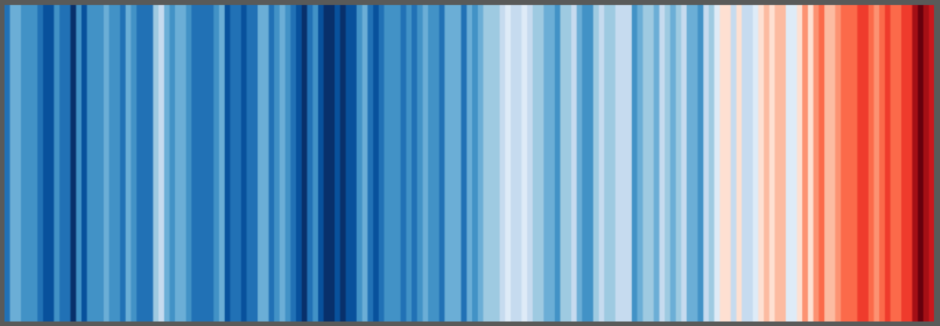

Romain Jacob   Laurent Vanbever
nsg.ee.ethz.ch @ ETH Zürich

RIPE 87   Nov. 27, 2023

What do you think consumes more energy?

Data Centers

Telco Networks

What do you think consumes more energy?

Point this way

Data Centers

or

Telco Networks

Point that way

What do you think consumes more energy?

Data Centers	or	Telco Networks

In 2022	240-340 TWh		260-360 TWh

https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks

What do you think consumes more energy?

Data Centers or **Telco Networks**

	Data Centers		Telco Networks	
In 2022	240-340	TWh	260-360	TWh
In 2015	200	TWh	220	TWh
Change of	+20-70%	in energy	+18-64%	in energy

https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks

What do you think consumes more energy?

Data Centers or **Telco Networks**

	Data Centers		Telco Networks	
In 2022	240-340	TWh	260-360	TWh
In 2015	200	TWh	220	TWh
Change of	+20-70%	in energy	+18-64%	in energy
	+340%	in workload	+600%	in traffic

https://www.iea.org/energy-system/buildings/data-centres-and-data-transmission-networks

Energy efficiency improved a lot

	Data Centers		Telco Networks	
Change in energy	+20-70%	in energy	+18-64%	in energy
is much smaller				
than in work done.	+340%	in workload	+600%	in traffic

Energy efficiency improved a lot but not enough!

Data Centers

Telco Networks

Change in energy is positive!

+20-70% in energy

+18-64% in energy

“With great power comes great responsibility”

- It is easy to keep increasing network capacity
- It is much harder to keep increasing energy efficiency

“With great power comes great responsibility”

- It is easy to keep increasing network capacity
- It is much harder to keep increasing energy efficiency

▶ Total energy usage is likely to keep increasing.

"With great power comes great responsibility" and carbon footprint.

- It is easy to keep increasing network capacity
- It is much harder to keep increasing energy efficiency

▶ Total energy usage is likely to keep increasing.

▶ Producing energy emits carbon.

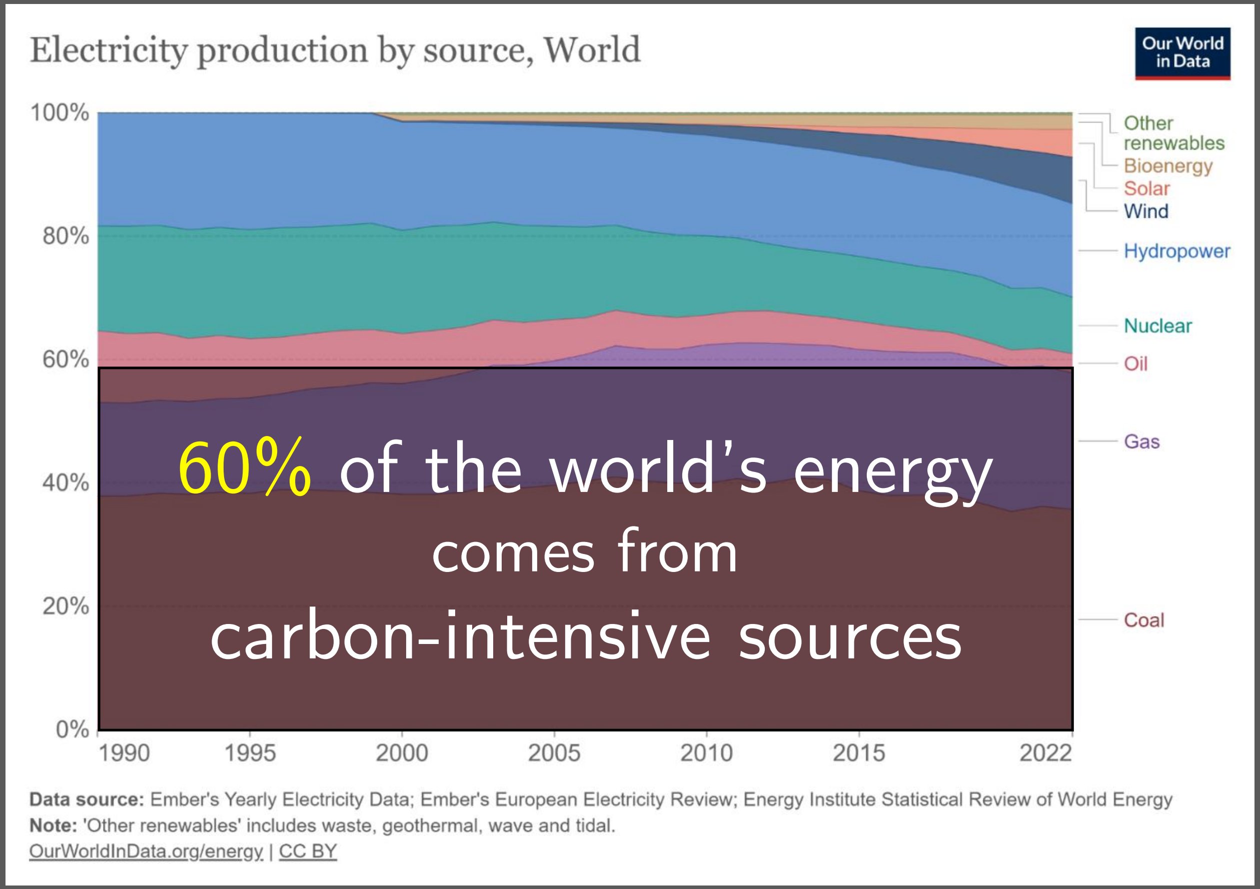

https://ourworldindata.org/grapher/electricity-prod-source-stacked

Tomorrow's Internet must
sleep more and grow old

to reduce its carbon footprint.

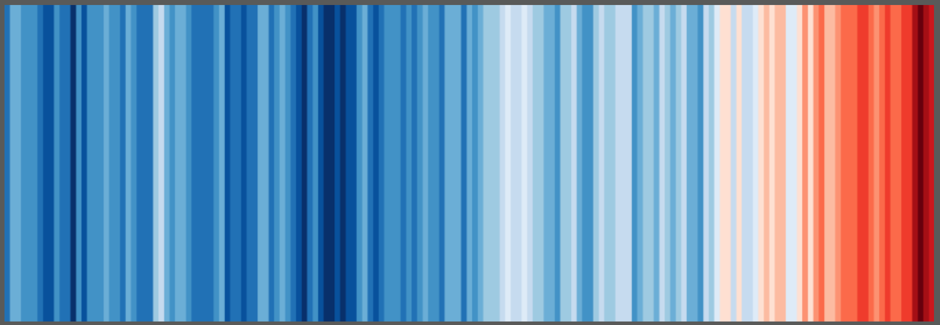

1 Reduce operational footprint
with better proportionality

2 Reduce embodied footprint
with sustainable procurement

Tomorrow's Internet must
sleep more and grow old

to reduce its carbon footprint.

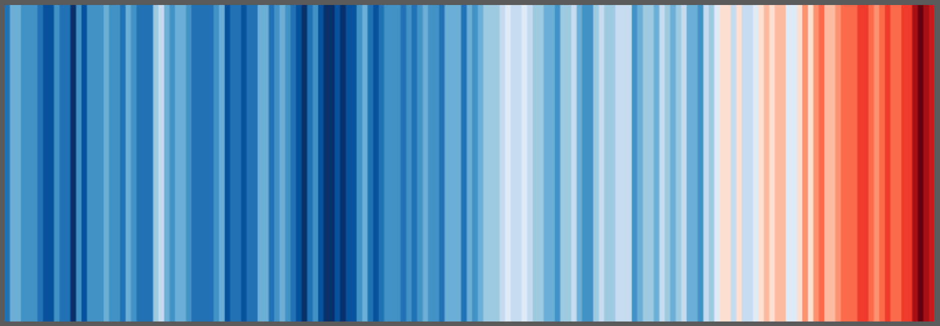

1 Reduce operational footprint
with better proportionality

Reduce embodied footprint
with sustainable procurement

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

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.

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 table[1]. Two energy values missing from the table are the energy cost of *cooling* the equipment and that of UPS (Uninterruptable Power Supplies) equipment[2]. 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 that this is almost negligible in comparison to other energy

[1]Note that the energy draw varies based on load and the values used in this study are based on observed average values.
[2]According to [14], air conditioning in data centers containing routing equipment costs approximately 20 – 60 Watts/ft^2.

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

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

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.

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 table[1]. Two energy values missing from the table are the energy cost of *cooling* the equipment and that of UPS (Uninterruptable Power Supplies) equipment[2]. 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 that this is almost negligible in comparison to other energy

[1] Note that the energy draw varies based on load and the values used in this study are based on observed average values.

[2] According to [14], air conditioning in data centers containing routing equipment costs approximately 20 – 60 Watts/ft^2.

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."

1 Network devices
are always “on.”

2 Network devices’ energy consumption
is mainly independent of traffic load.

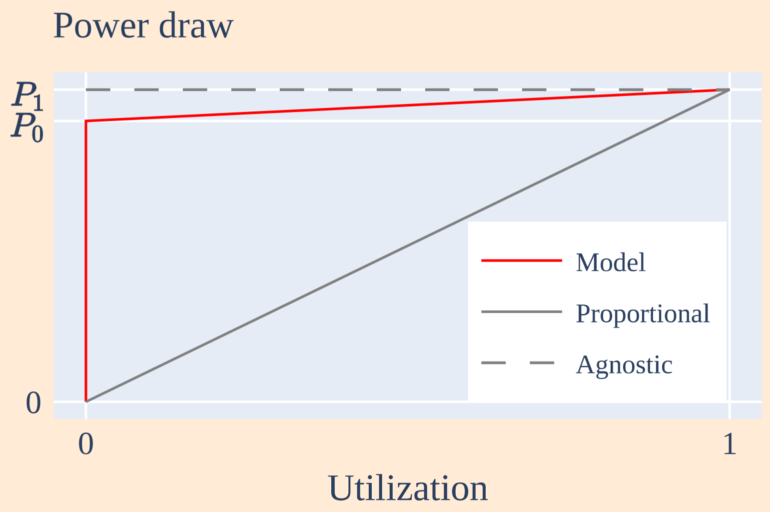

1 Network devices are always "on."

2 Network devices' energy consumption is mainly independent of traffic load.

3 Network devices are under-utilized.

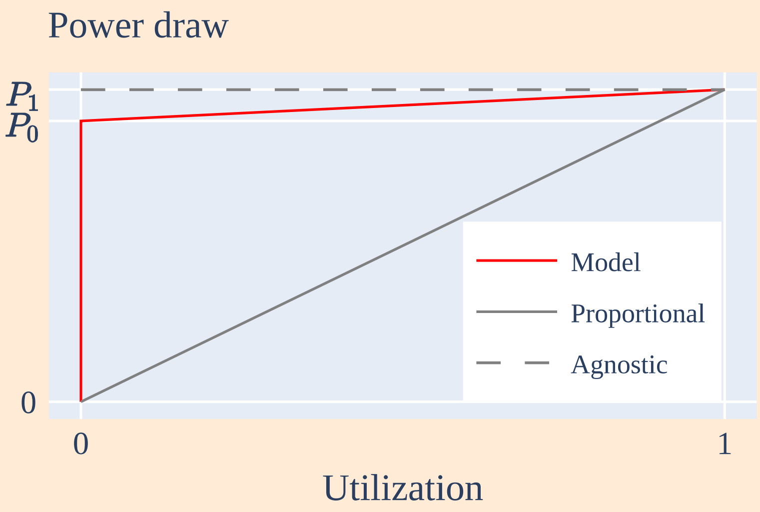

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

1 Network devices are always “on.”

2 Network devices’ energy consumption is mainly independent of traffic load.

3 Network devices are under-utilized.

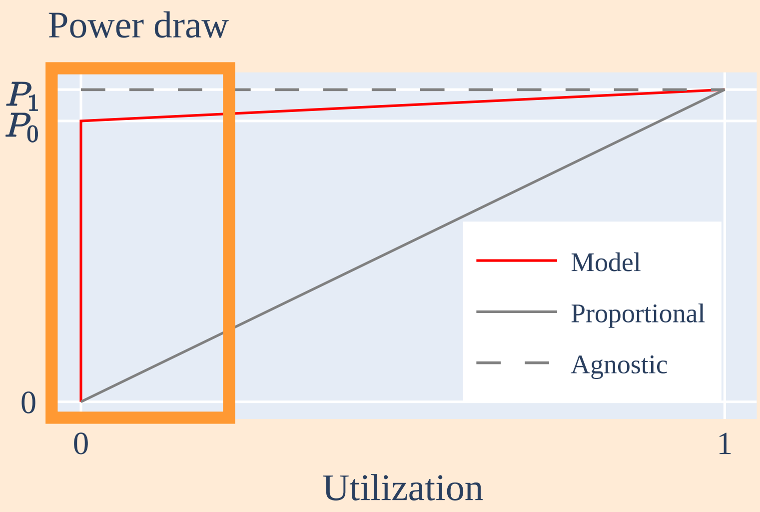

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

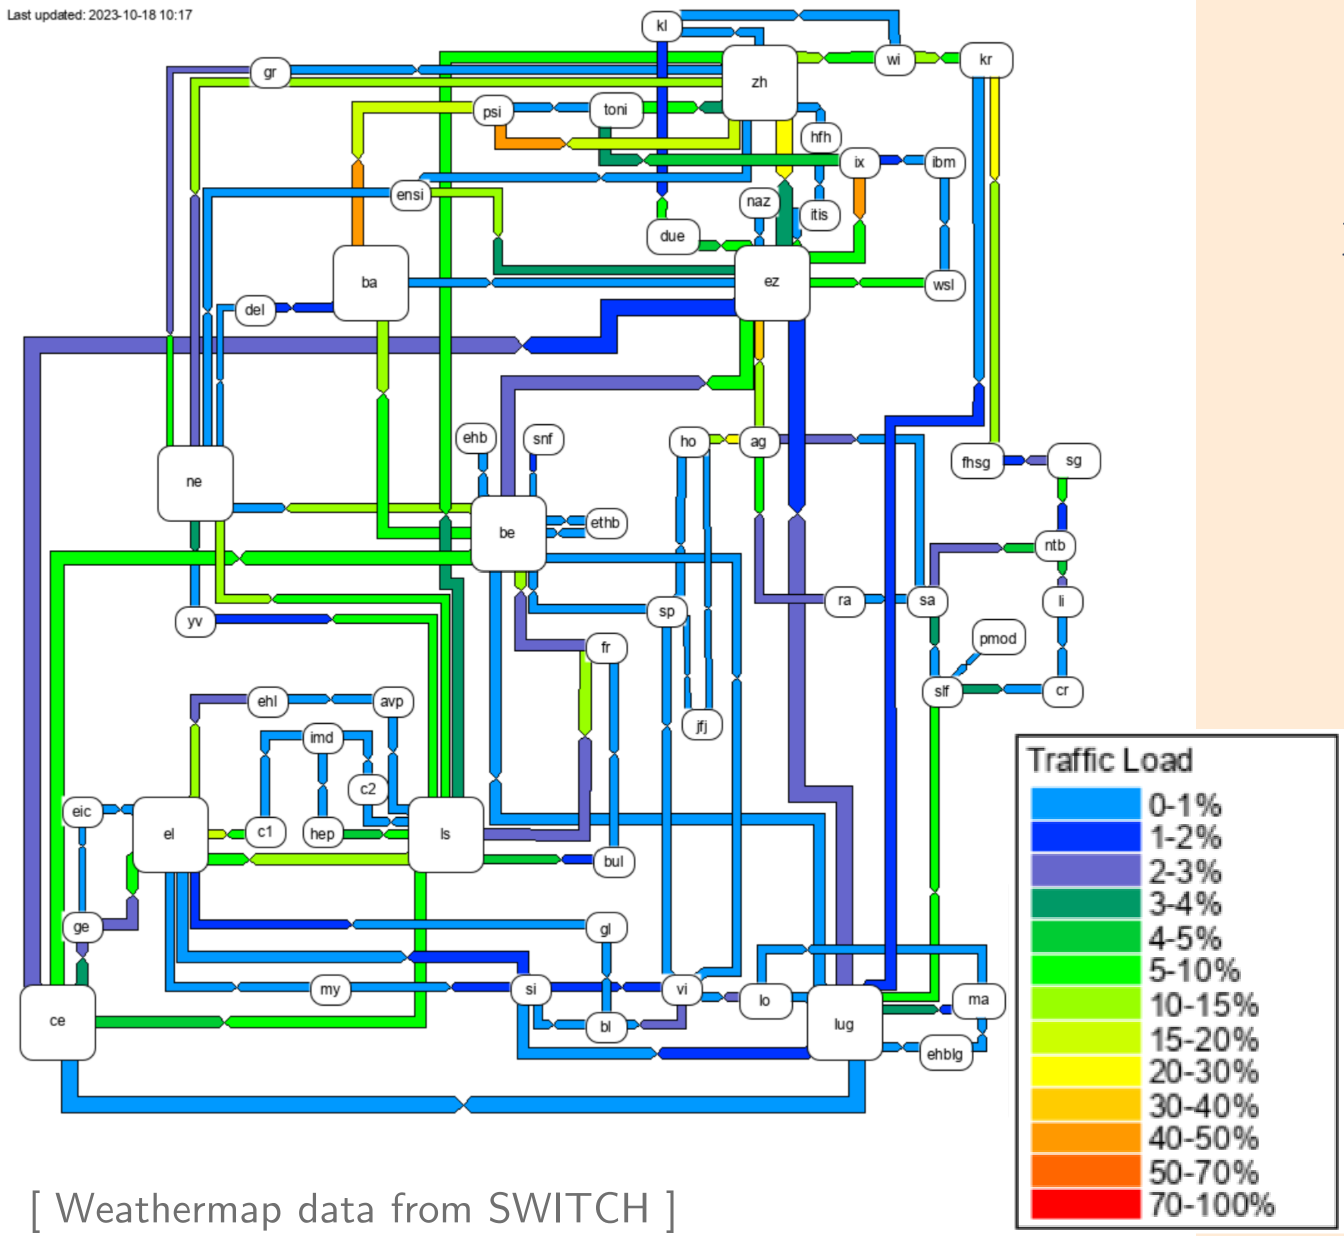

[Weathermap data from SWITCH]

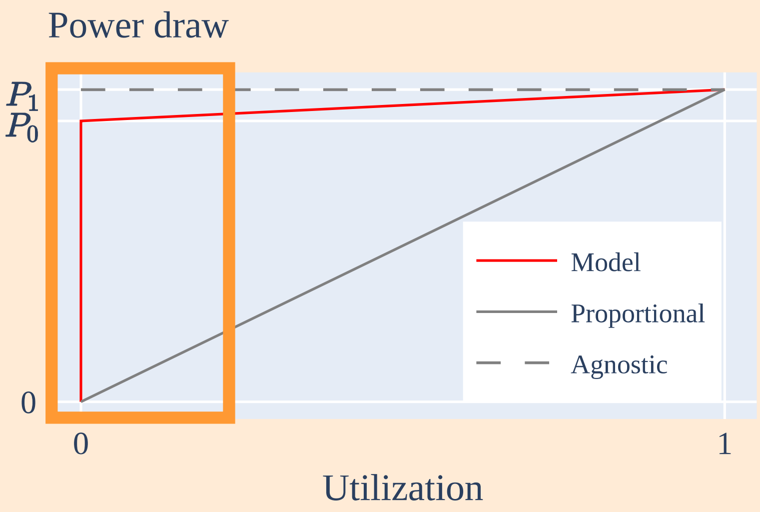

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

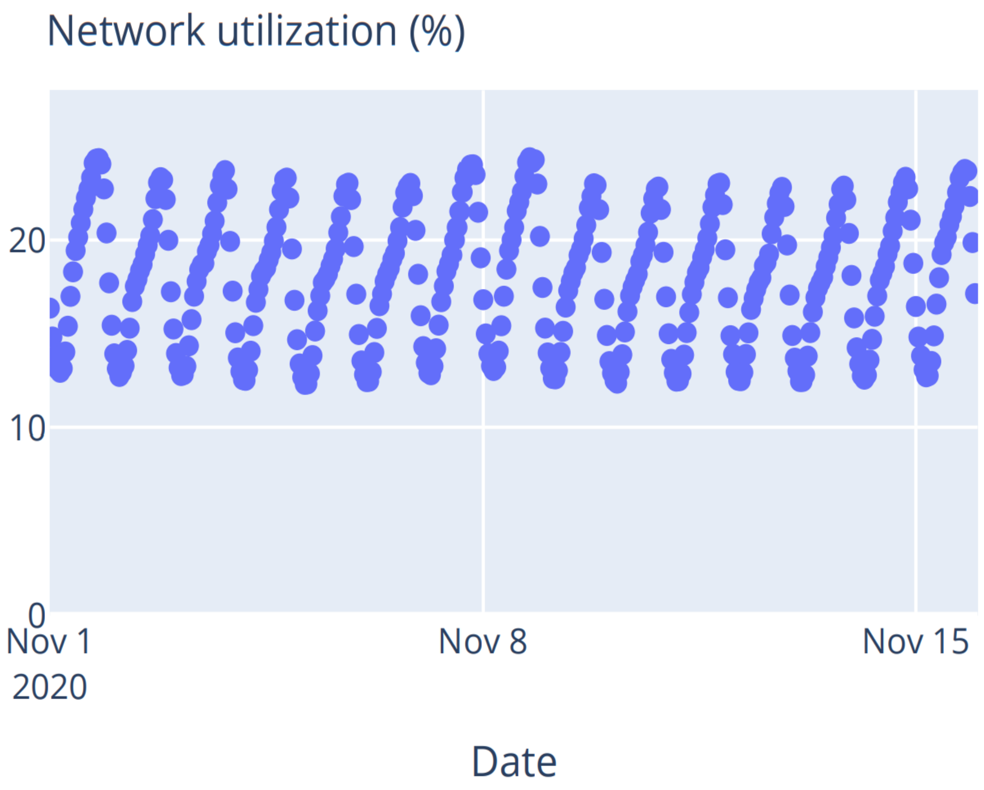

[OVH Weathermap dataset]

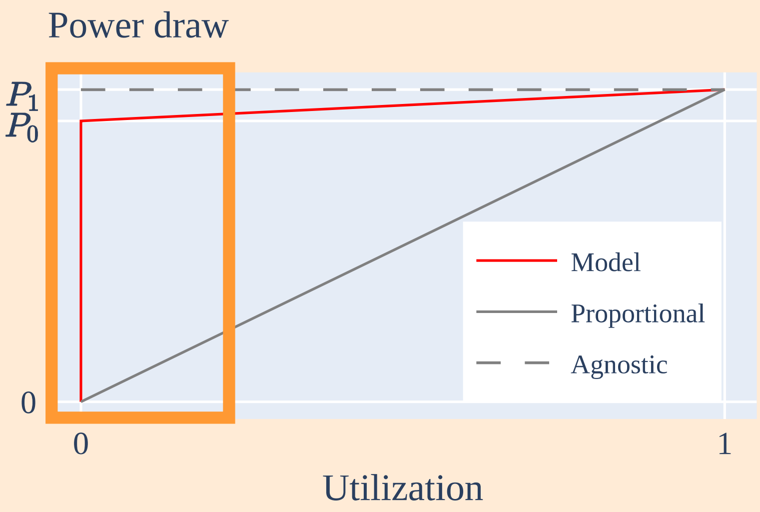

ISP overprovision networks to support

- 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.

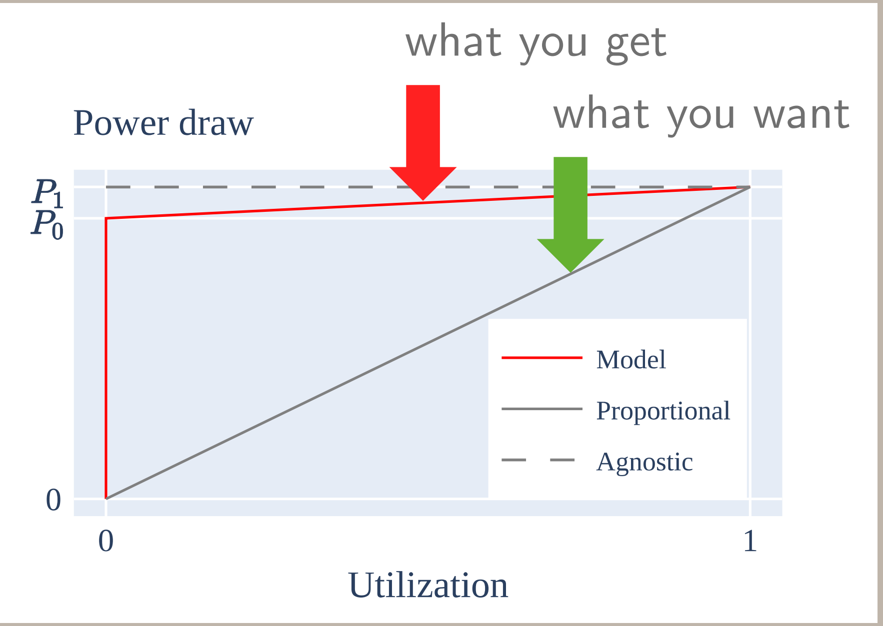

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

There two ways to improve energy efficiency

- Run more often at high utilization

  “Buffer-and-Burst”

  Time-shifting

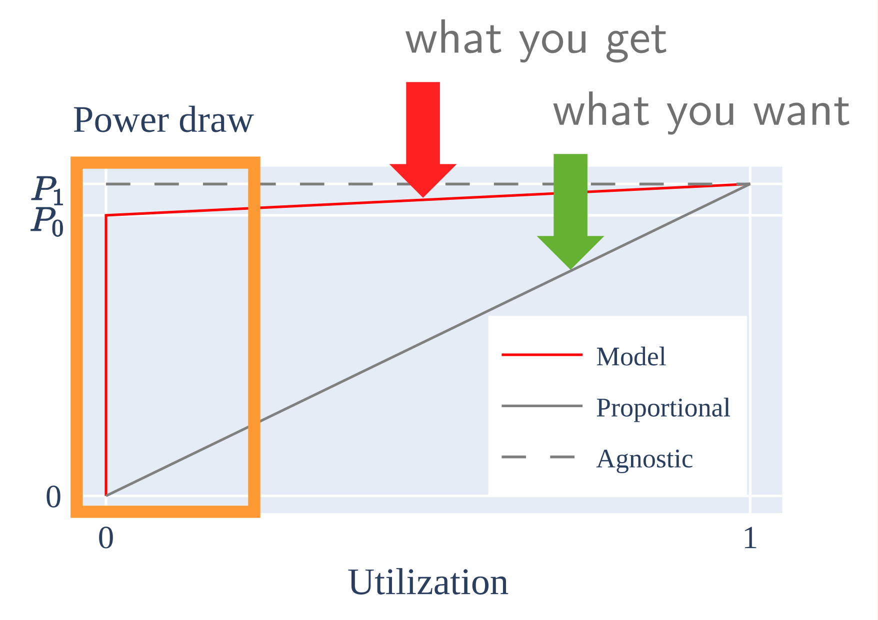

There two ways to improve energy efficiency

- Run more often at high utilization

  "Buffer-and-Burst"

  Time-shifting

- Take low-utilization power down

▶ Ideally, do both.

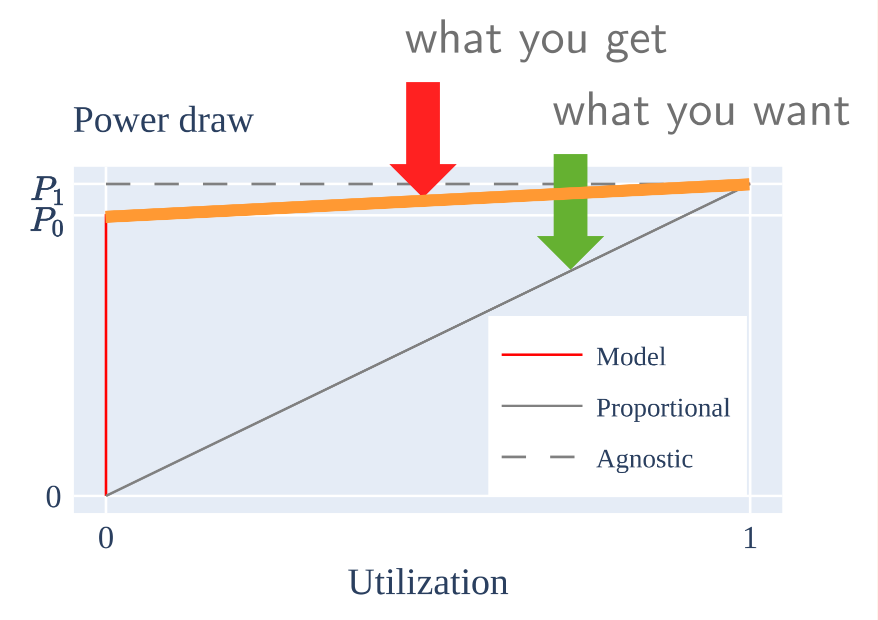

The basic idea is to turn off "stuff" whenever possible.

What can we possibly turn off?

- Ports
- Line cards
- Entire device...

The basic idea is to turn off "stuff" whenever possible.

What can we possibly turn off?

- Ports
- Line cards
- Entire device...

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

The basic idea is to turn off “stuff” whenever possible.

What can we possibly turn off?

- Ports
- Line cards
- Entire device...

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

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

The basic idea is to turn off “stuff” whenever possible. That’s nothing new.

Academia NSDI 2008

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

Sergiu Nedevschi*† 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

RIPE 86

Techniques to reduce network power consumption

Peter Ehiwe, May 2023 @RIPE86

The theory says we can save tens of energy % in ISP networks.

Academia

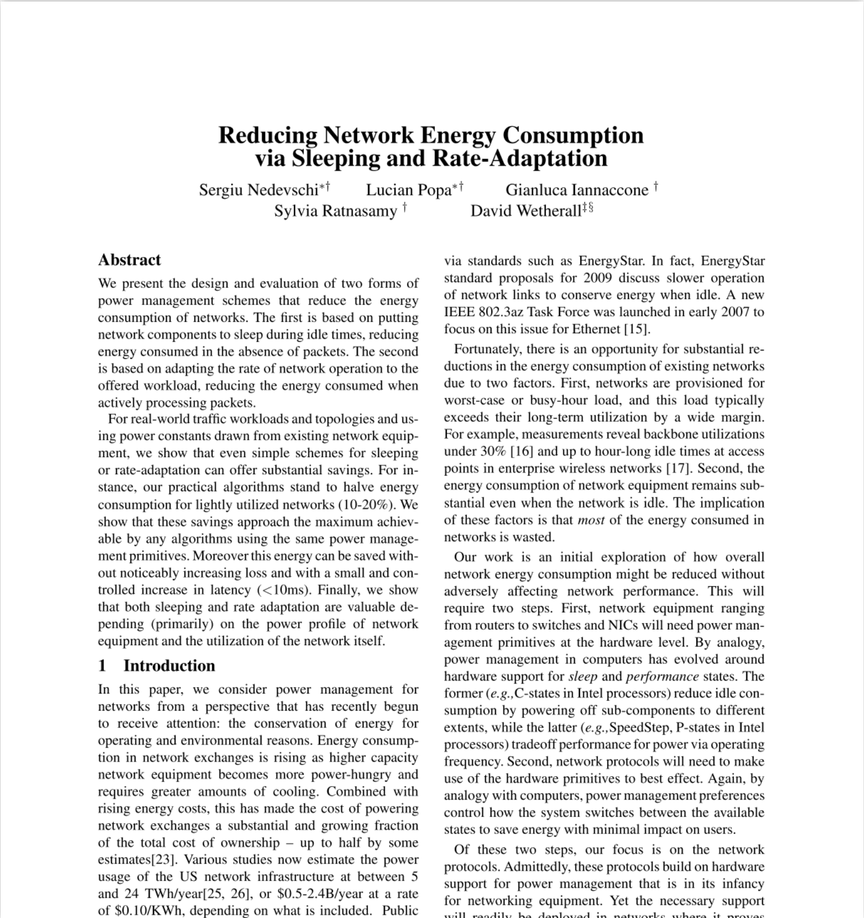

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

Sergiu Nedevschi[*†] 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

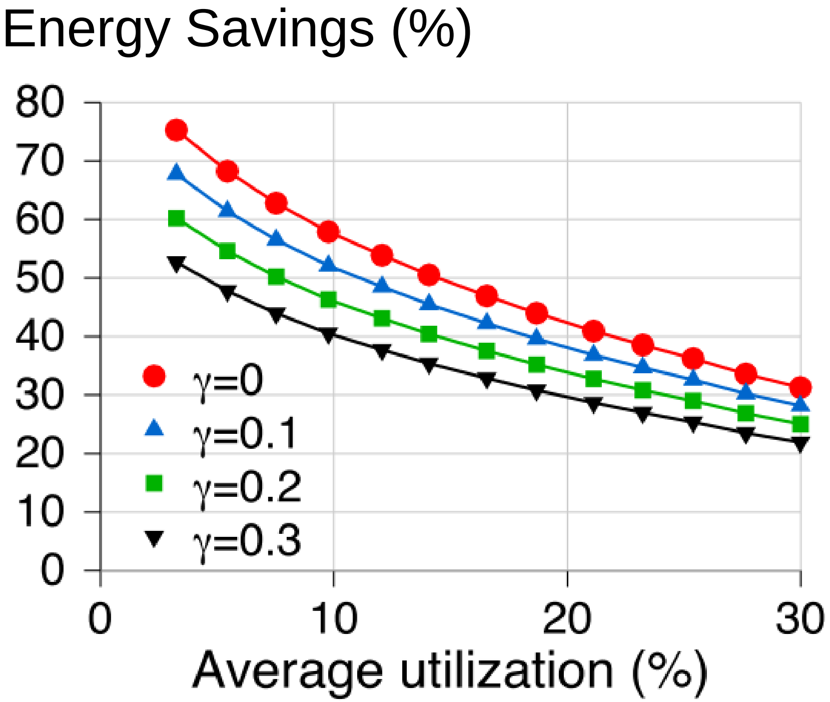

The theory says we can save tens of energy % in ISP networks.

Academia

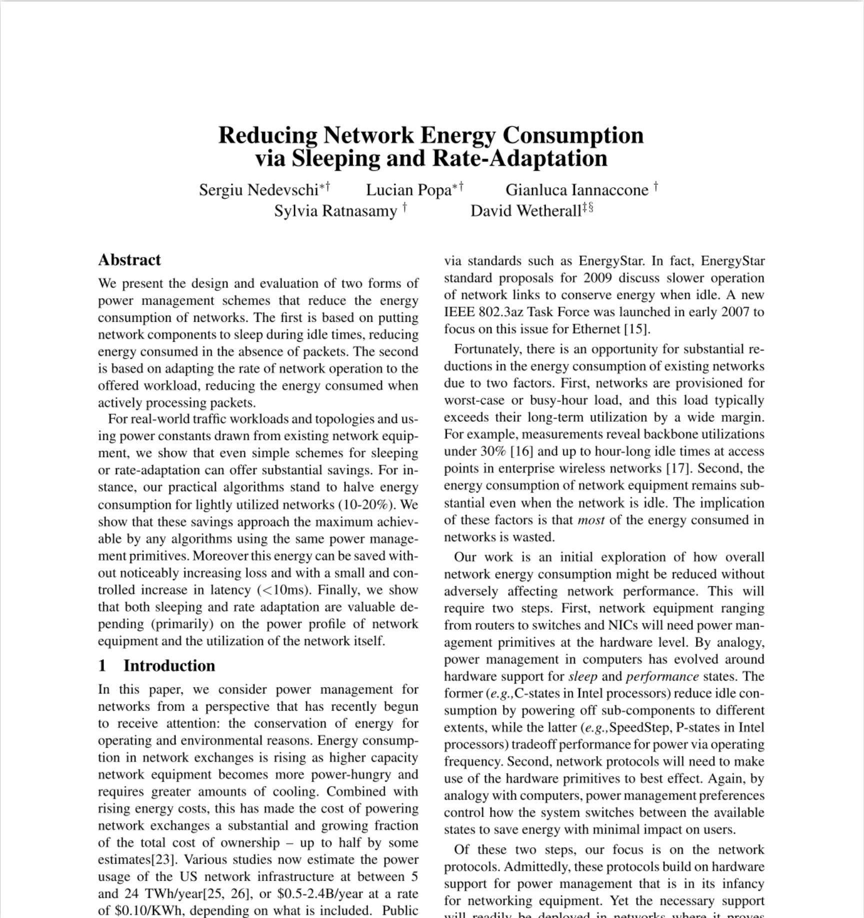

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

Sergiu Nedevschi*† 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

How? Buffer-and-Burst

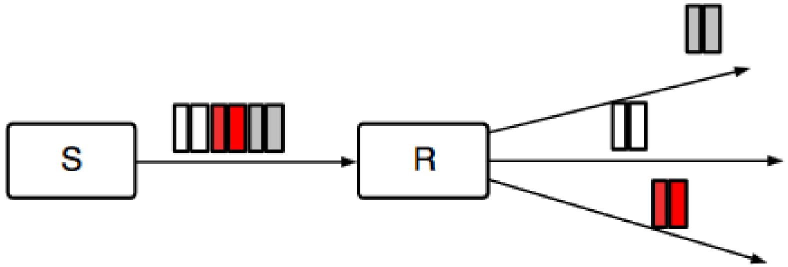

Assuming

- Wake-up delay $1ms$
- Buffering time $10ms$

Practice

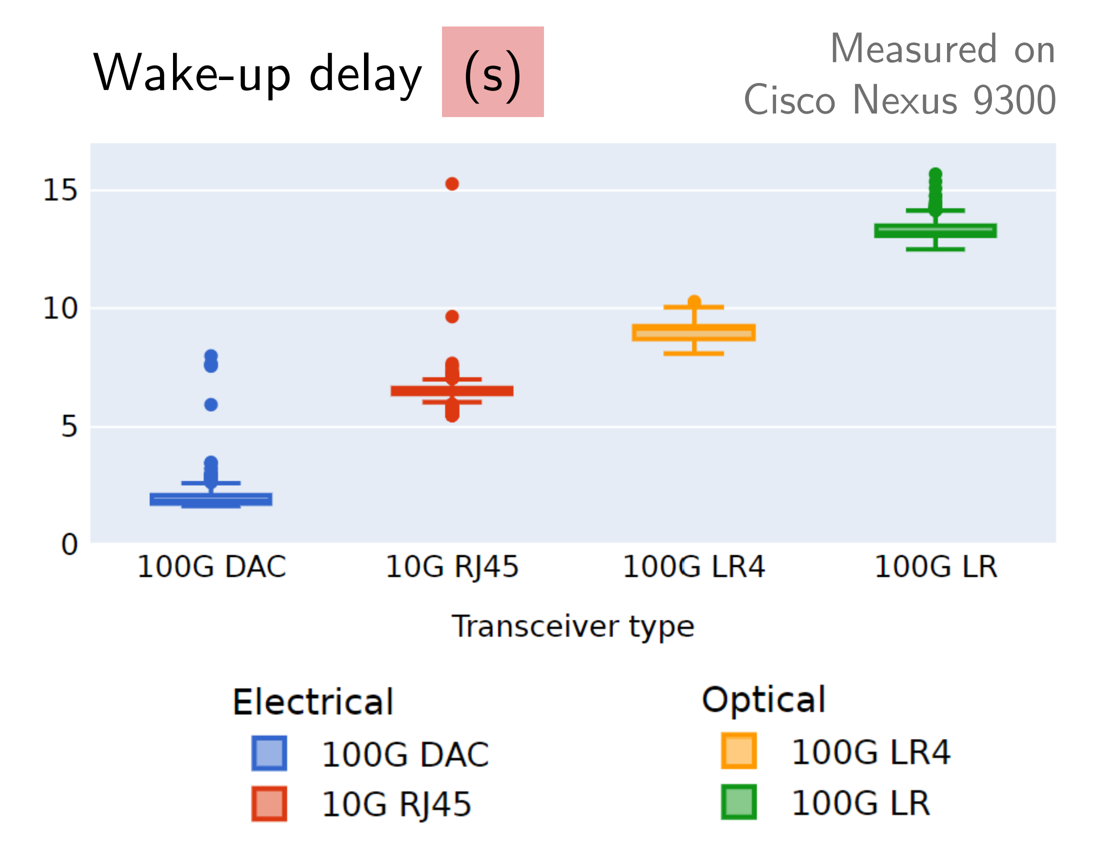

Theory

How? Buffer-and-Burst

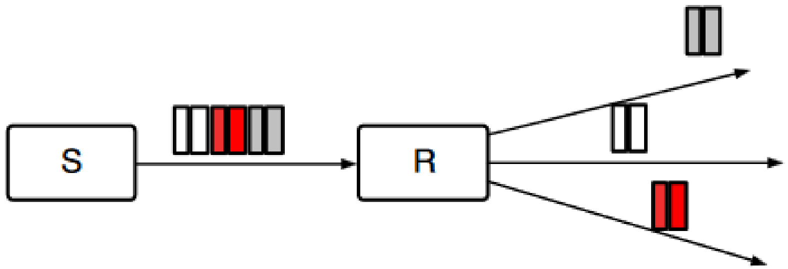

Assuming

- Wake-up delay $1ms$
- Buffering time $10ms$

In practice, transcievers are 1000x slower to start than required for savings via buffering.

Wake-up delay (s)

Measured on Cisco Nexus 9300

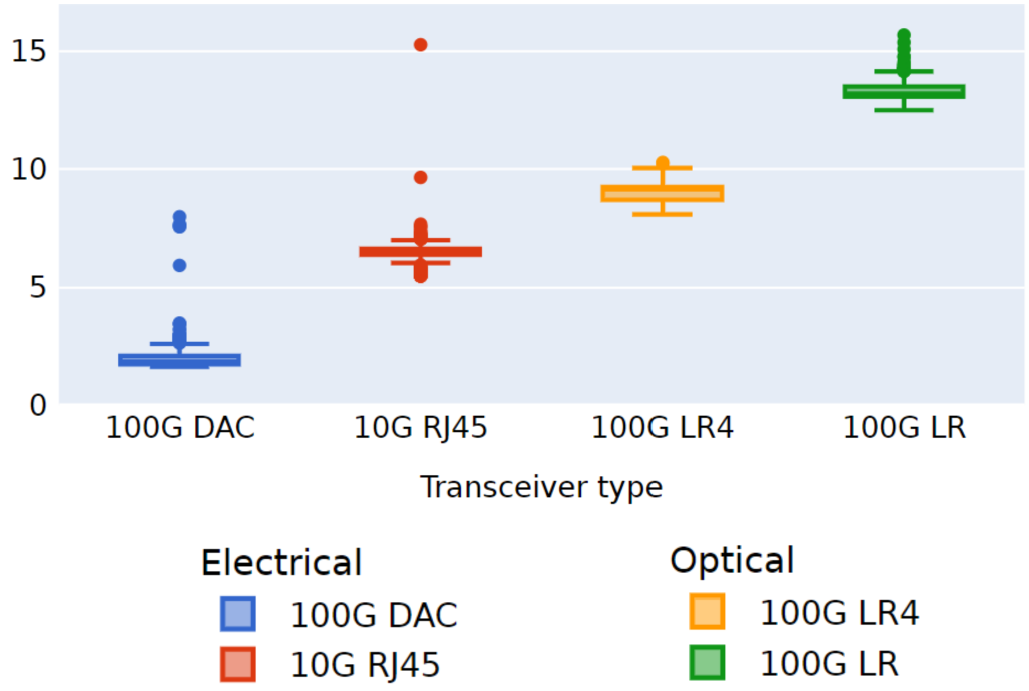

How? Buffer-and-Burst

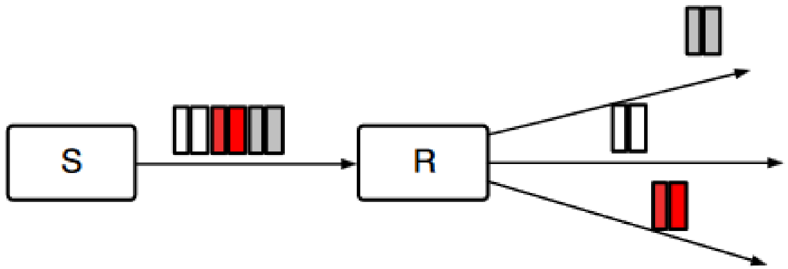

Assuming

- Wake-up delay $1ms$
- Buffering time $10ms$

We can still “sleep” at longer timescales.

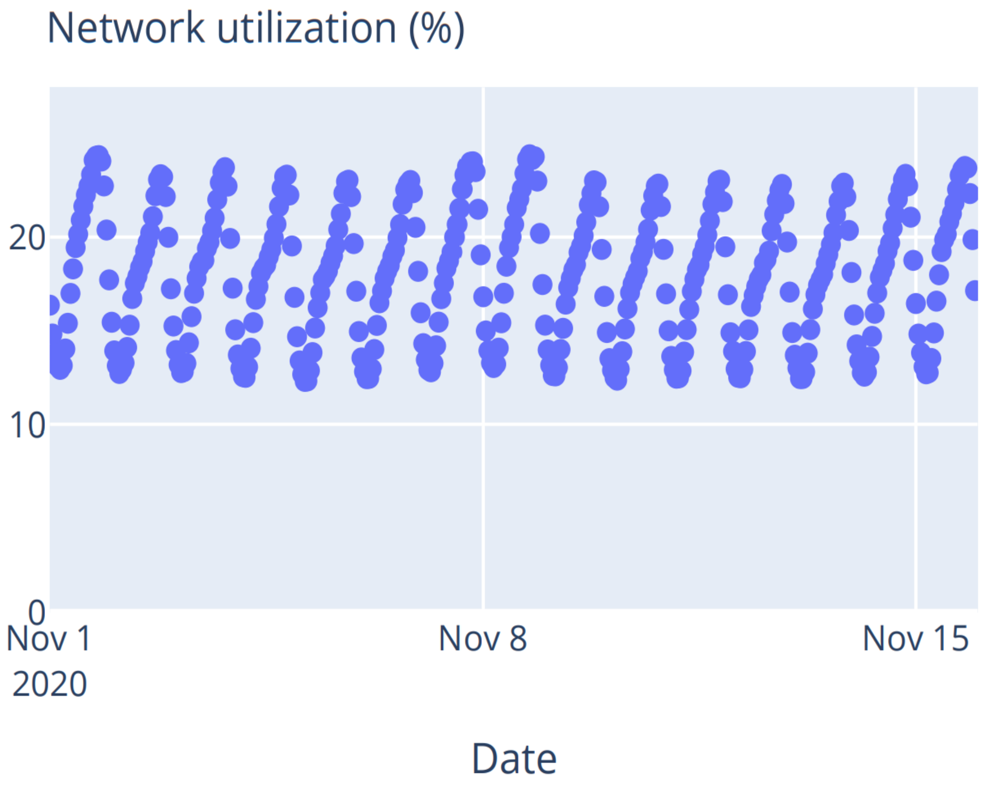

We can still "sleep" at longer timescales.

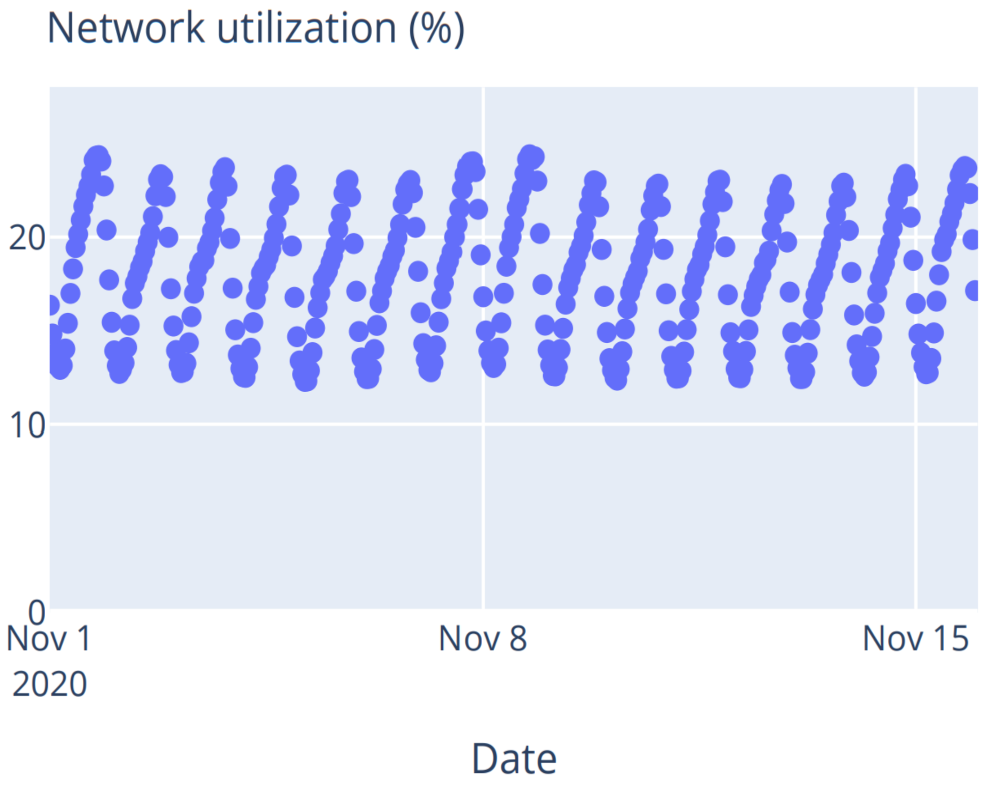

- How to play nice with routing?
- Which signal to use for sleeping and wake-up control?
- How much sleeping affects traffic? And in case of bursts?

We can still “sleep” at longer timescales.

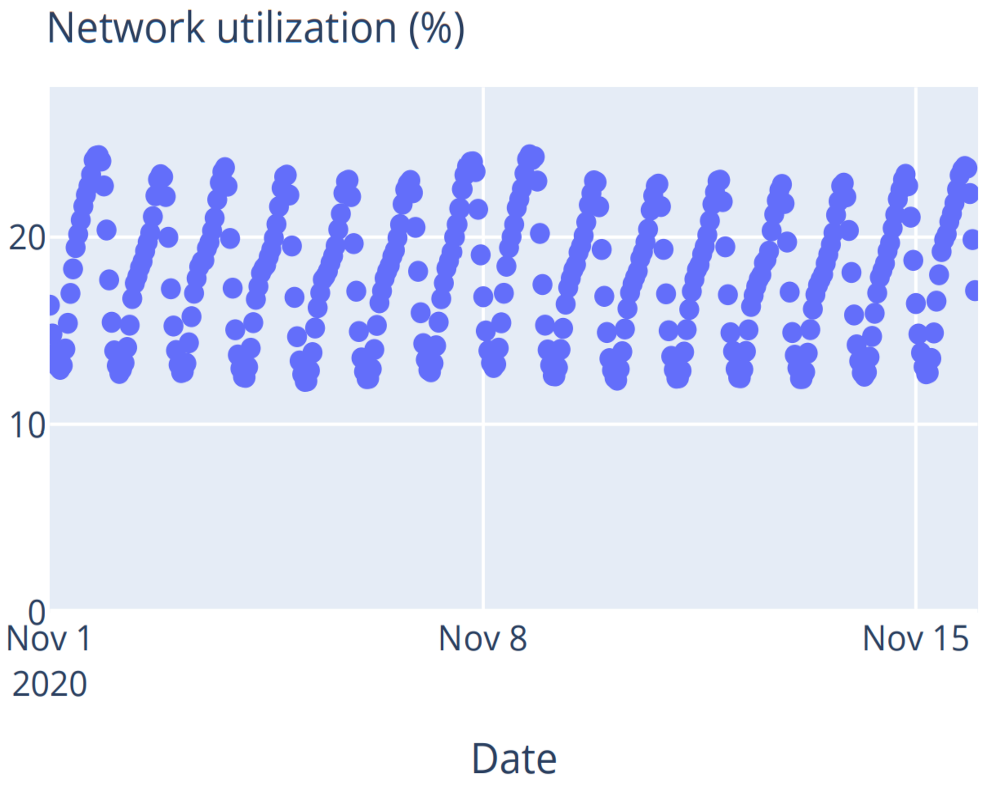

- How to play nice with routing?
  - Easy, because wake-up is slow!
- Which signal to use for sleeping and wake-up control?
- How much sleeping affects traffic? And in case of bursts?

We can still "sleep" at longer timescales.

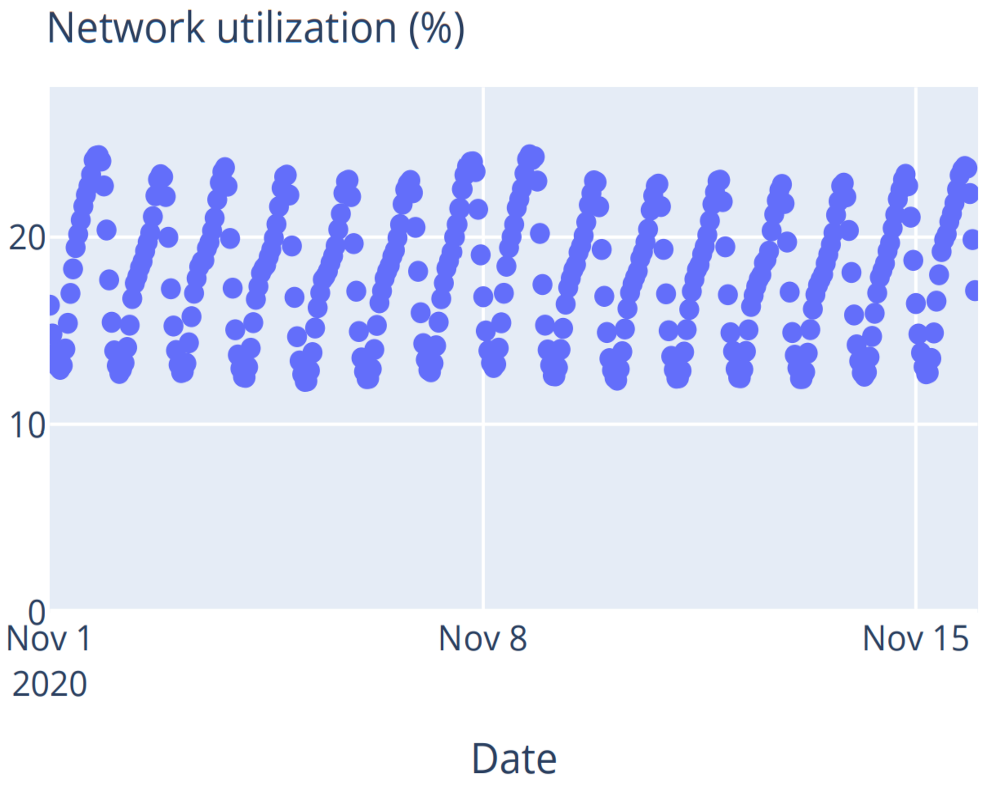

- How to play nice with routing?
  - Easy, because wake-up is slow!
- Which signal to use for sleeping and wake-up control?
  - Link utilization works
    Easy to collect with OSPF-TE
- How much sleeping affects traffic? And in case of bursts?

We can still “sleep” at longer timescales.

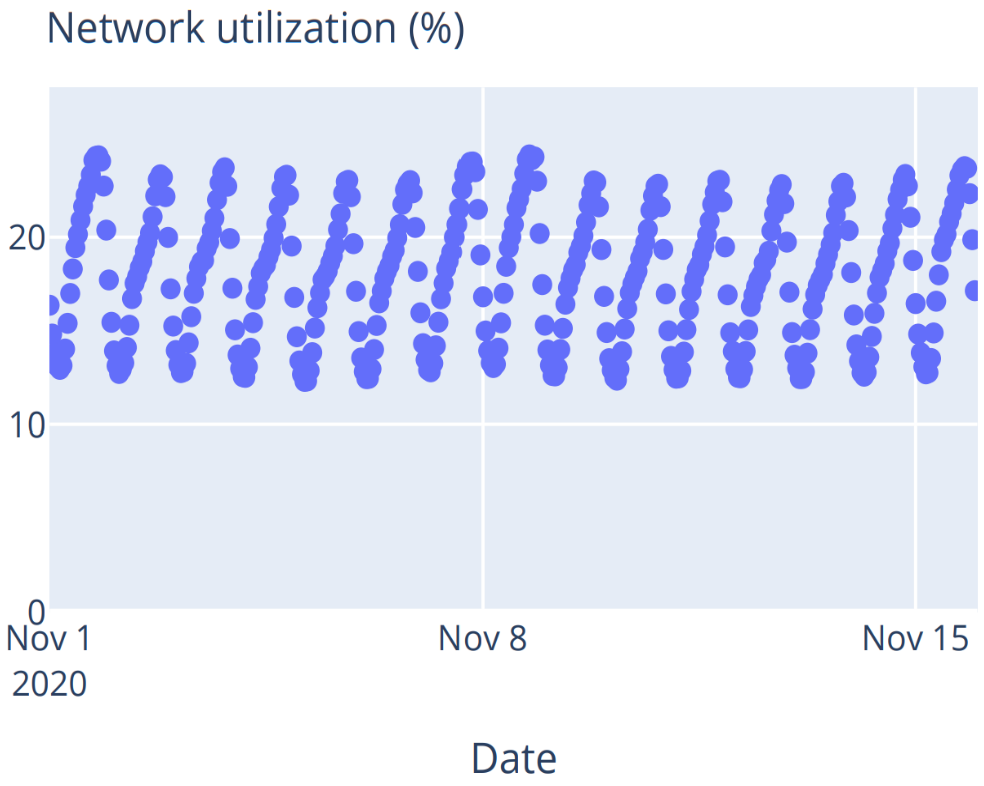

- How to play nice with routing?
  - Easy, because wake-up is slow!
- Which signal to use for sleeping and wake-up control?
  - Link utilization works
    Easy to collect with OSPF-TE
- How much sleeping affects traffic? And in case of bursts?
  - Very little, actually.

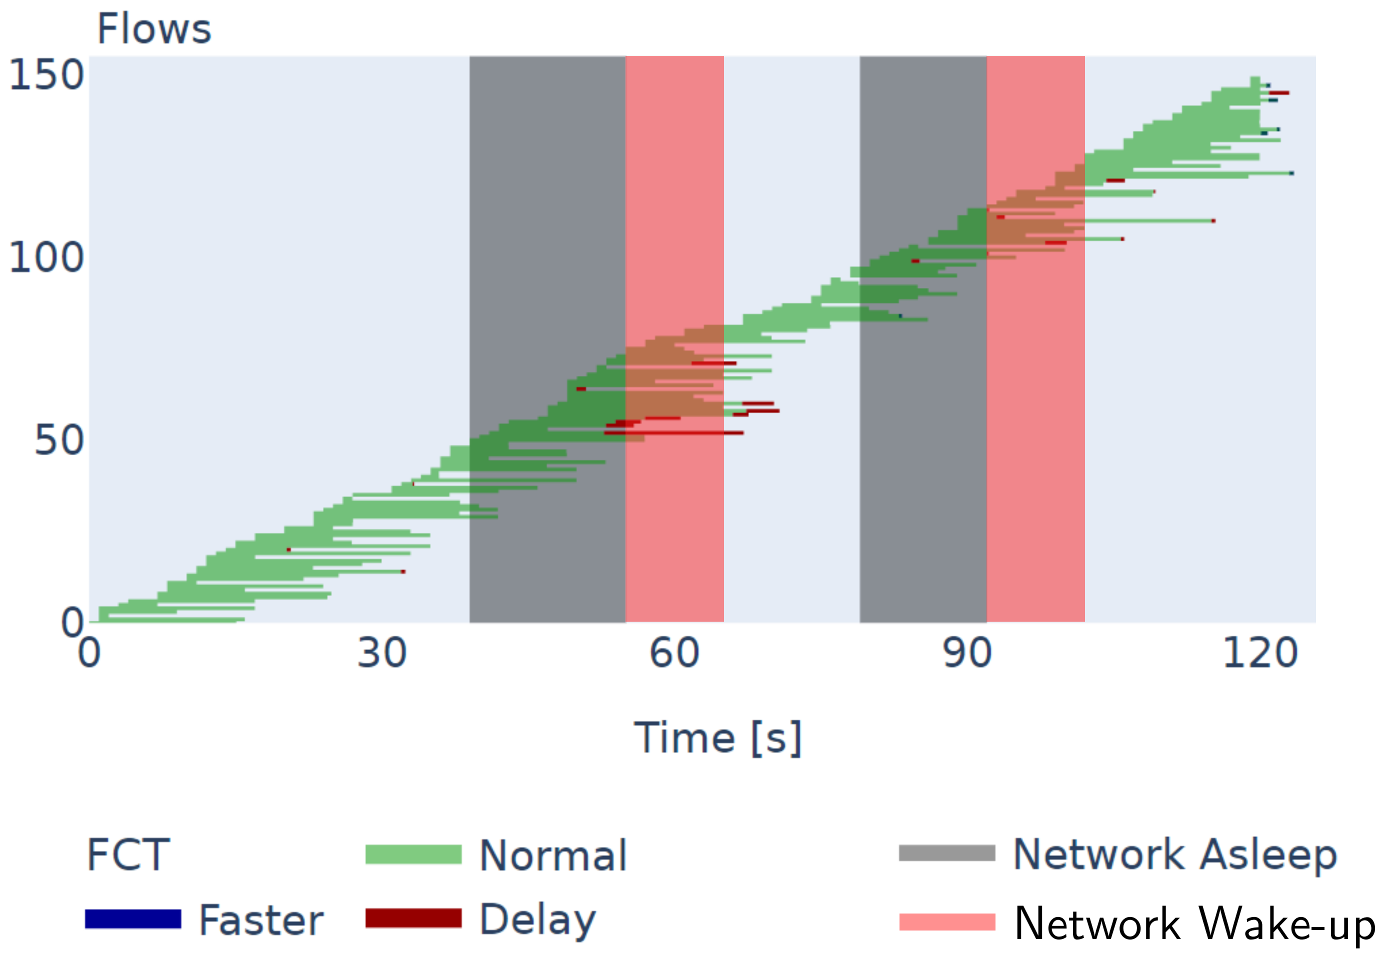

Flows
150
100
50
0
0
30
60
90
120
Time [s]
FCT
Faster
Normal
Delay
Network Asleep
Network Wake-up

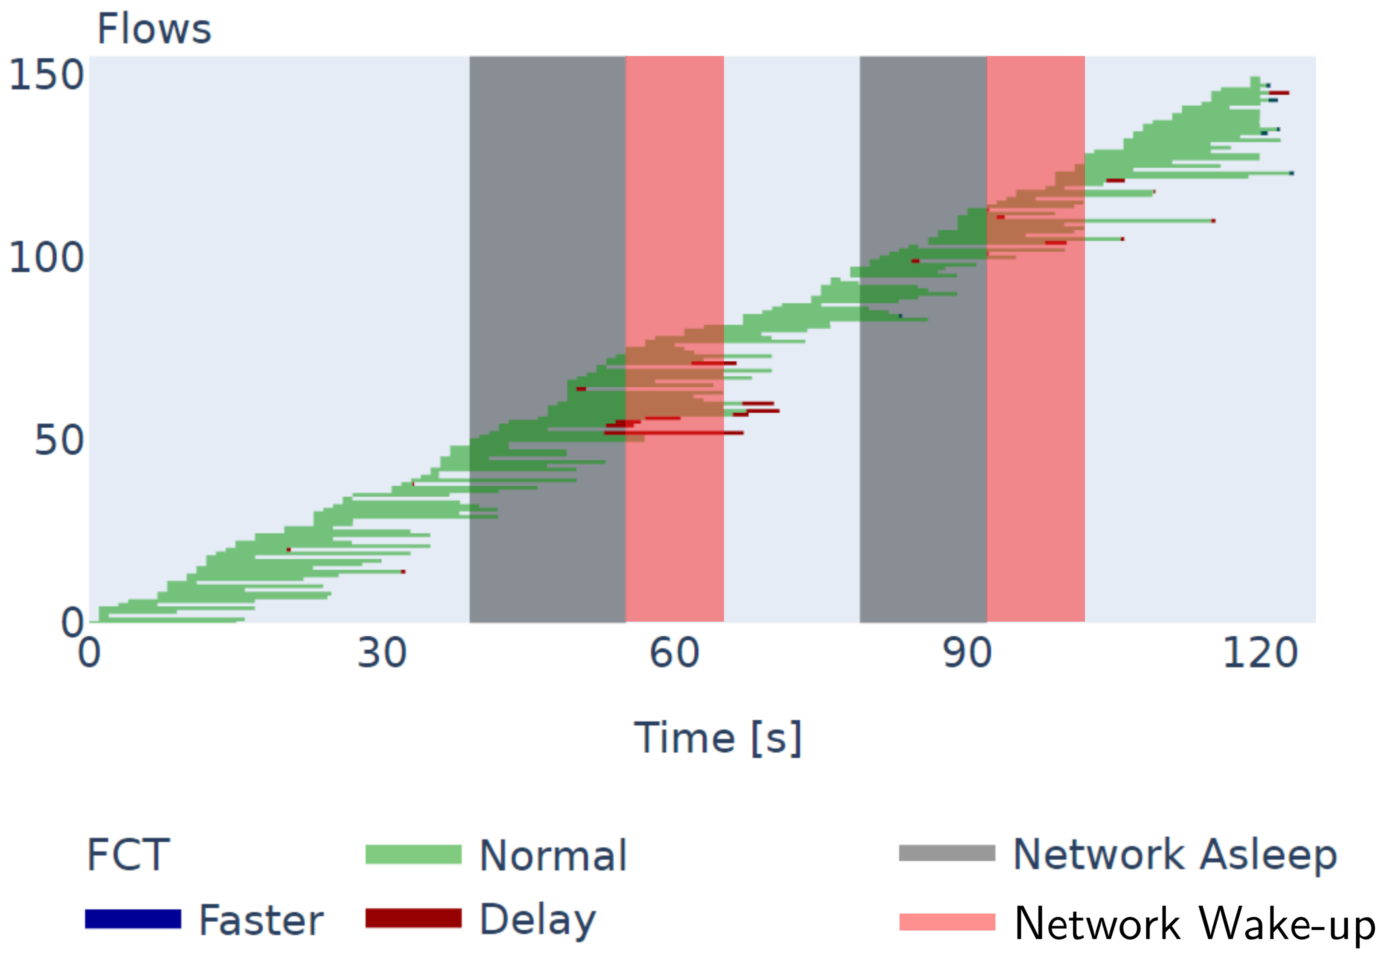

Only the flows that would finish when the network wakes up may suffer some FCT increase

We see no extra loss and little FCT increase
because TCP is doing its job decently well.

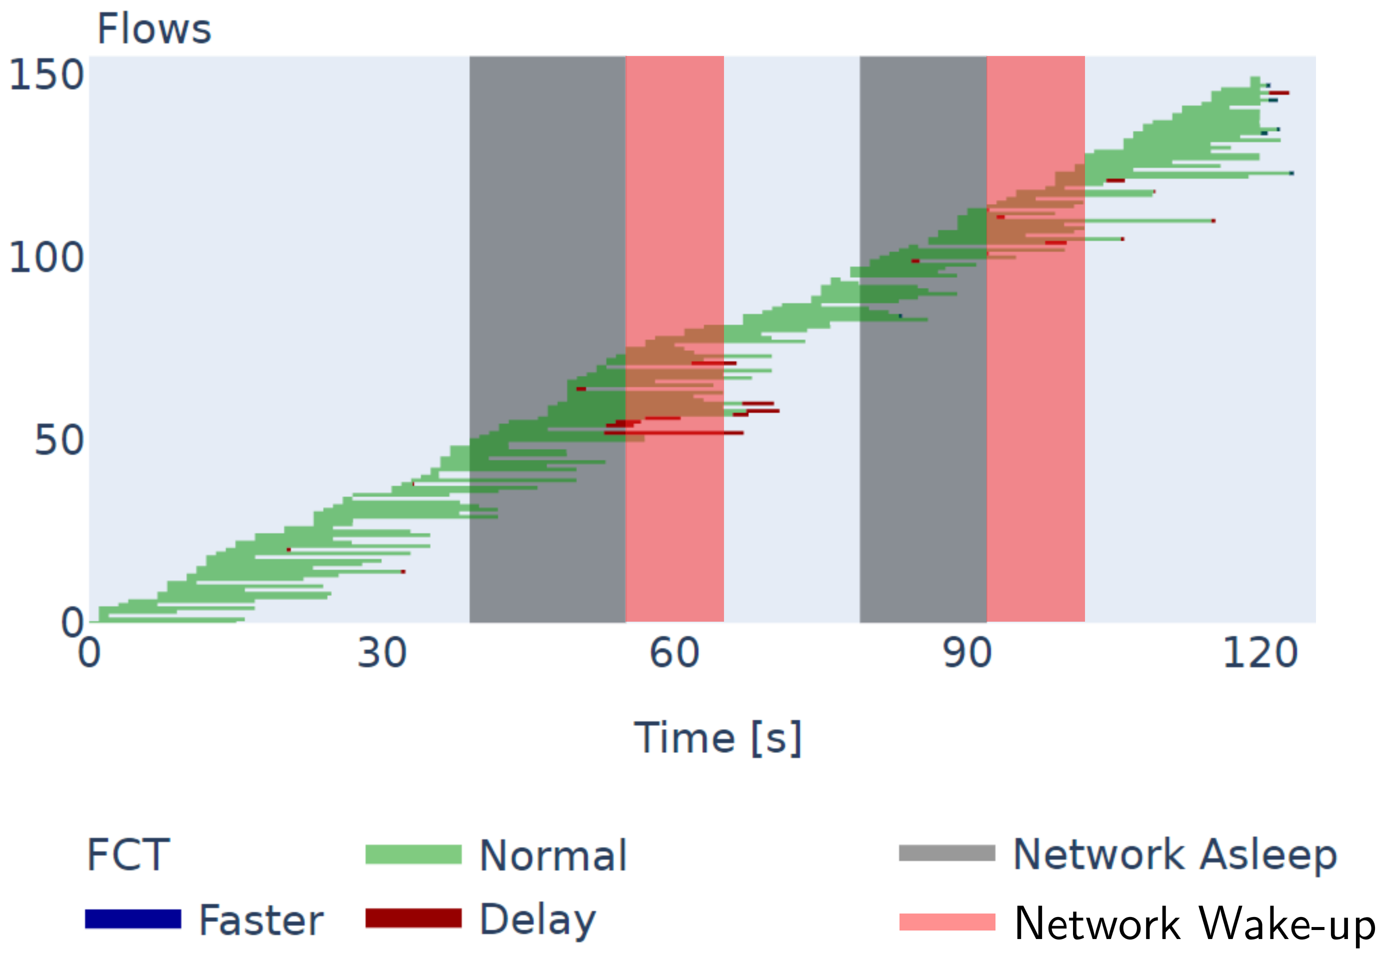

Only the flows that would finish when the network wakes up may suffer some FCT increase

We had to work “hard” to even see an effect...

We can still “sleep” at longer timescales

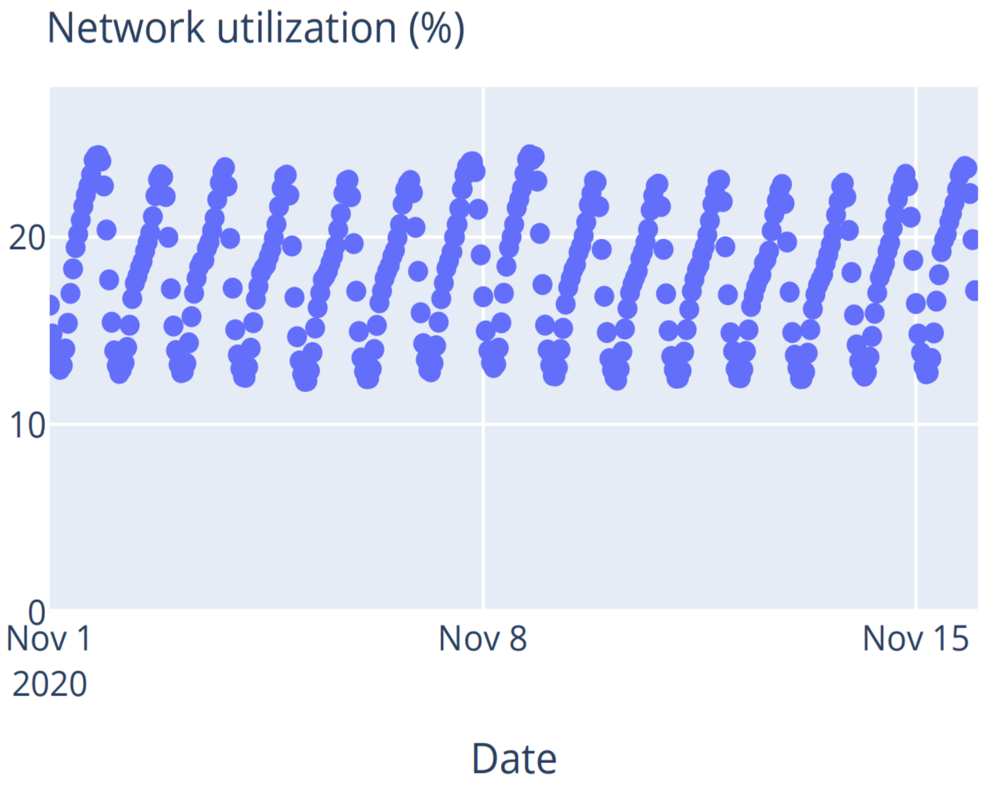

Ultimately, it is very similar to a traditional TE problem.

How much energy can we really save?

The theory says we can save **tens of energy %** in ISP networks.

Academia


Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

Sergiu Nedevschi Lucian Popa Gianluca Iannaccone
Sylvia Ratnasamy David Wetherall

Abstract

1 Introduction

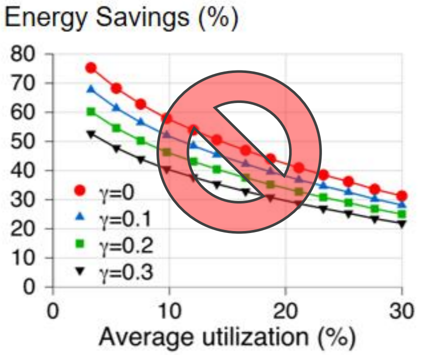

How much energy can we really save?

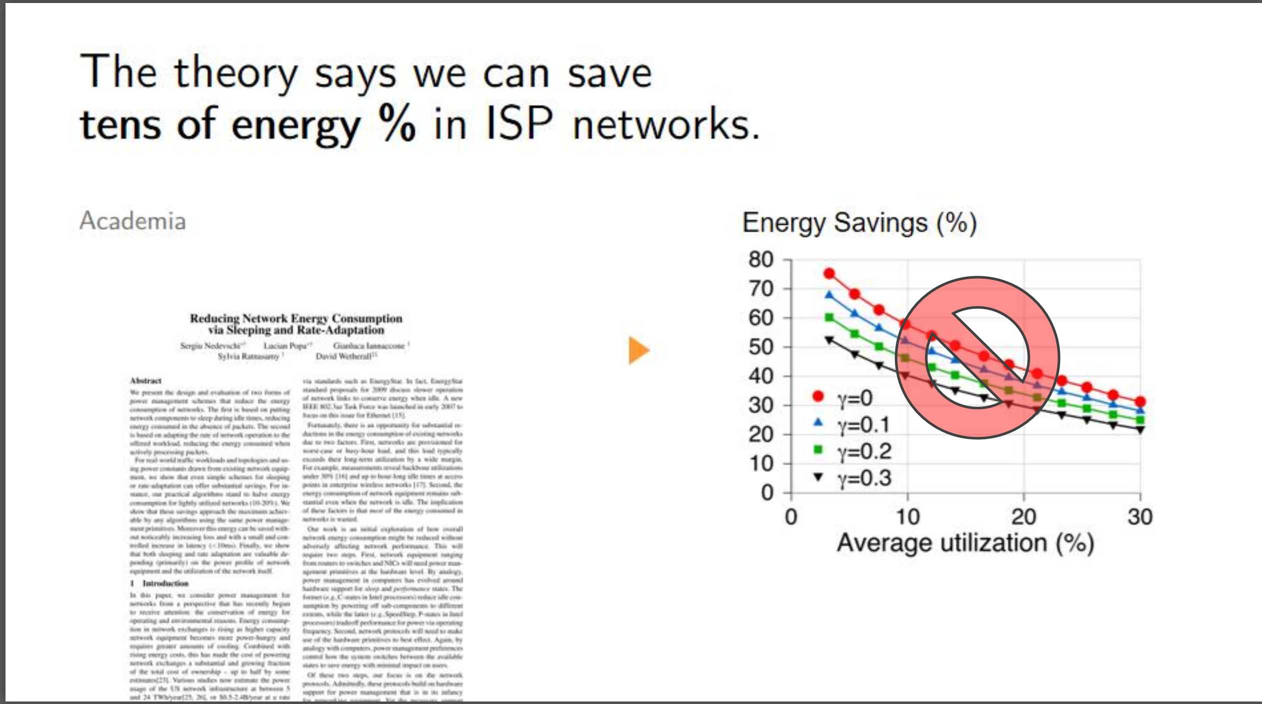

Hard to say because we lack

1 Measurements

2 Test cases

Energy savings are hard to estimate because we lack good power models.

- Datasheets only talk about the max power
- Devices are never under full load

Energy savings are hard to estimate because we lack good power models.

- Datasheets only talk about the max power
- Devices are never under full load

? How much power is drawn under “typical” load?

Energy savings are hard to estimate because we lack good power models.

... so we are building our own ...

Profiling a Tofino switch

WEDGE 100BF-32X

Wedge switch

Power meter

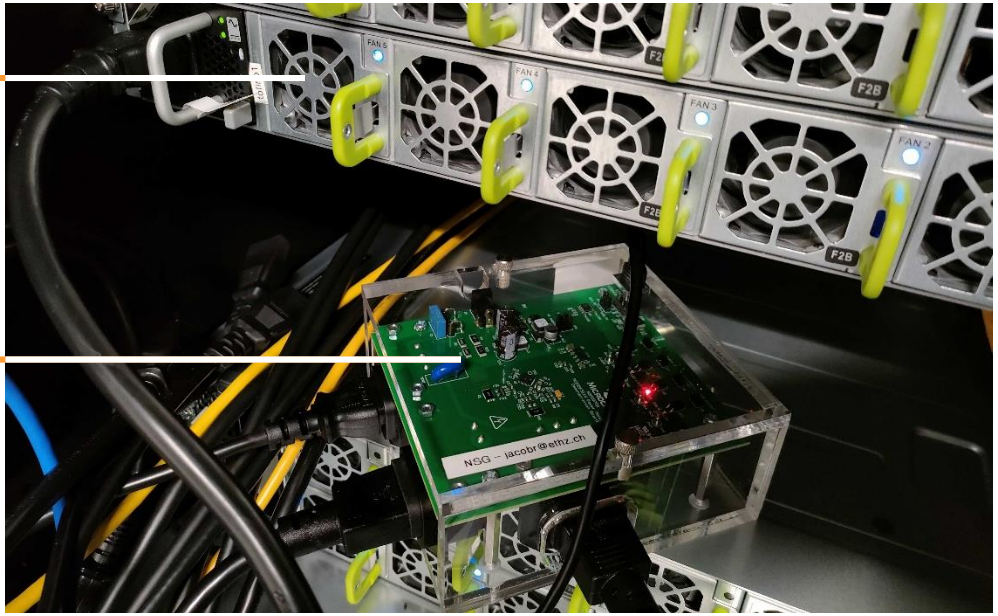

Energy savings are hard to estimate because we lack good power models.

… so we are building our own …

Device power	$=$	Static power	f(device config)
	$+$	Energy per bit * bit rate	
	$+$	Energy per packet * packet rate	
	$+$	Fan power	$\sim$ f(temperature)
	$+$	Power conversion losses	f(power demand)

We discuss with the IETF to establish a benchmark for instantiating such models.

```
Benchmarking Methodology Working Group                         V. Manral
Internet-Draft                                                 P. Sharma
Intended status: Informational                               S. Banerjee
Expires: September 13, 2013                                           HP
                                                                 Y. Ping
                                                                     H3C
                                                          March 12, 2013


            Benchmarking Power usage of networking devices
                   draft-manral-bmwg-power-usage-04

Abstract

   With the rapid growth of networks around the globe there is an ever
   increasing need to improve the energy efficiency of network devices.
   Operators are begining to seek more information of power consumption
   in the network, have no standard mechanism to measure, report and
   compare power usage of different networking equipment under different
   network configuration and conditions.

   This document provides suggestions for measuring power usage of live
   networks under different traffic loads and various switch router
   configuration settings.  It provides a benchmarking suite which can
```

Reviving this expired draft

with inputs from

- Carlos Pignataro
- <add-your-name-here>

We have a modelling approach.
You have devices that need modeling.

Academics have limited access
to devices used in the field.

> ? Can we measure yours?

We have a modelling approach.
You have devices that need modeling.

Academics have limited access to devices used in the field.

? Can we measure yours?

- We sent you hardware
- You plug it in
- Everyone gets data!

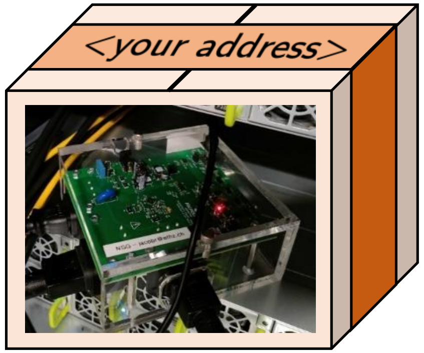

Academics have limited access to devices used in the field.

? Can we measure yours?

- We sent you hardware
- You plug it in
- Everyone gets data!

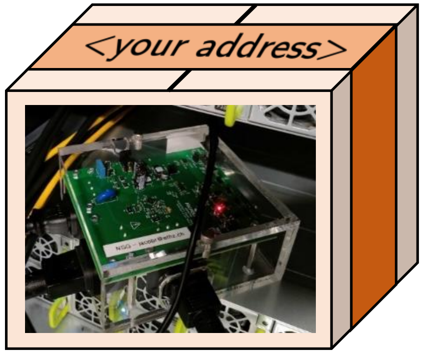

Vision

RIPE Atlas for Power Data

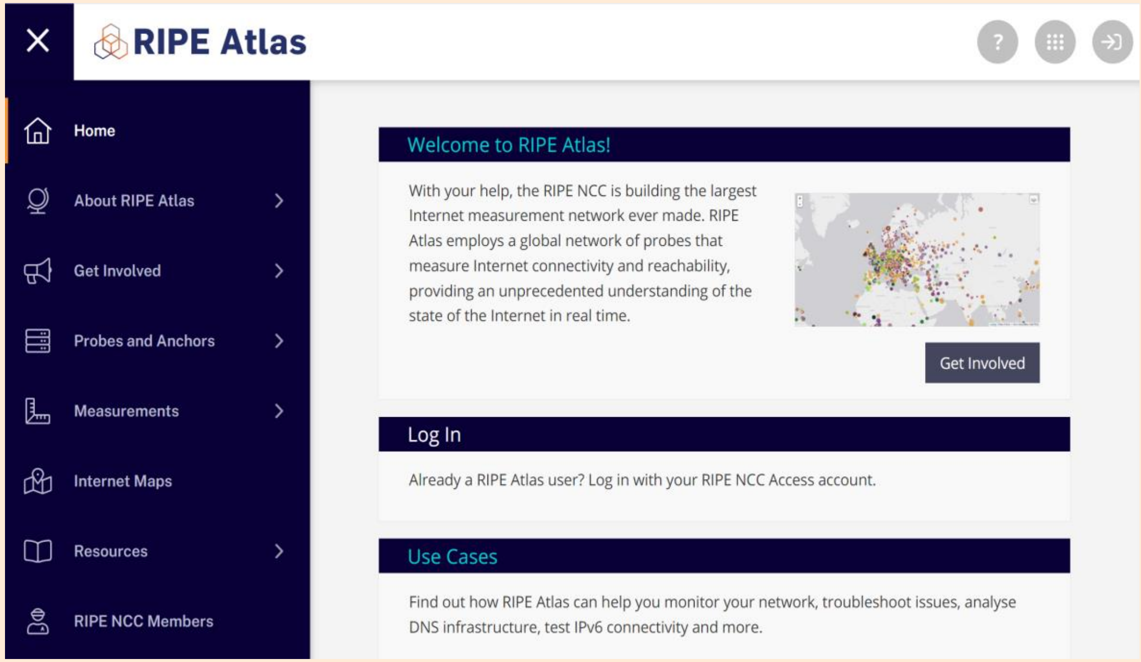

How much energy can we really save?

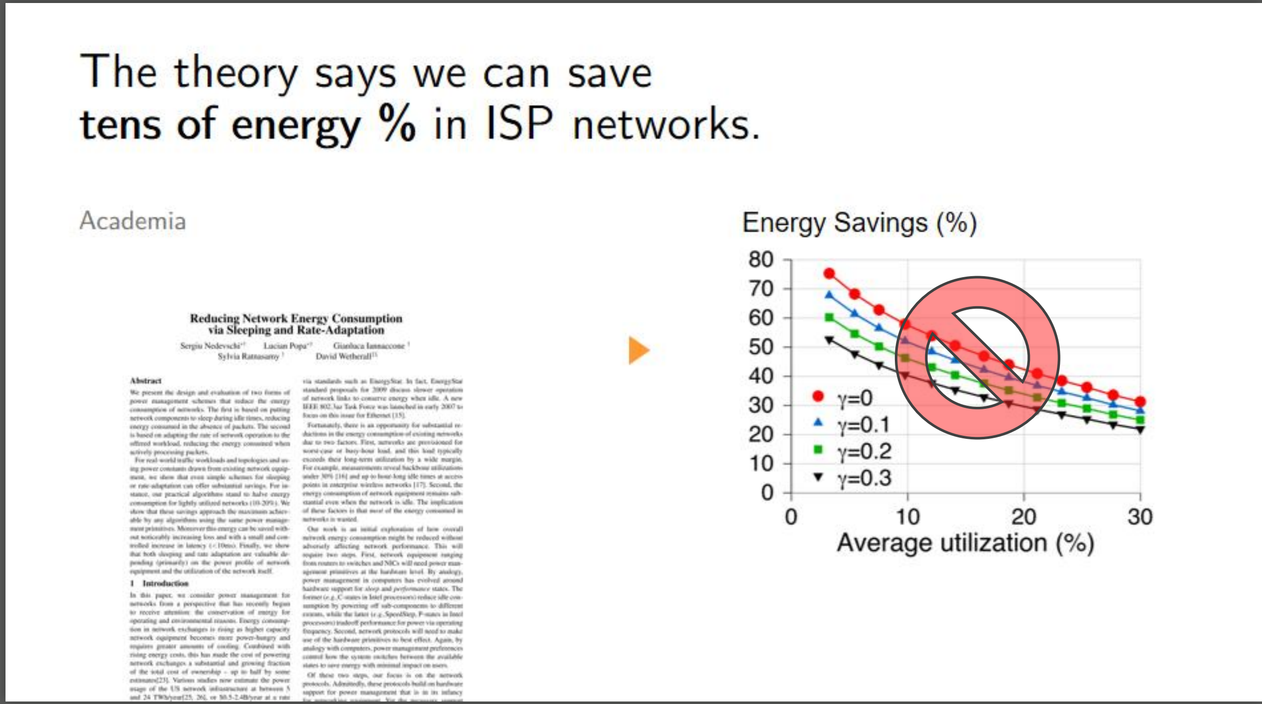

Hard to say because we lack

1 Measurements

2 Test cases

Energy savings are hard to estimate because they depend on the network.

We see no extra loss and little FCT increase
because TCP is doing its job decently well.

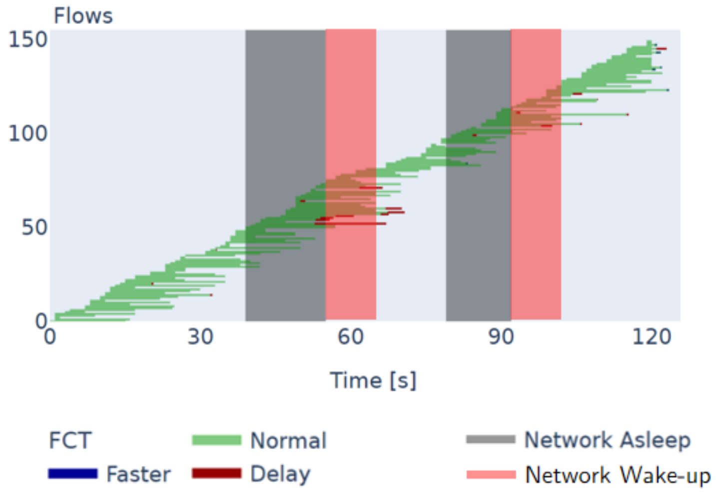

Only the flows that would finish when the network wakes up may suffer some FCT increase

We had to work "hard" to even see an effect...

- Anything can happen in simulation.
- We need real traffic dynamics to accuratly assess the impact of sleeping.

? Can we get yours?

Tomorrow's Internet must sleep more and grow old

to reduce its carbon footprint.

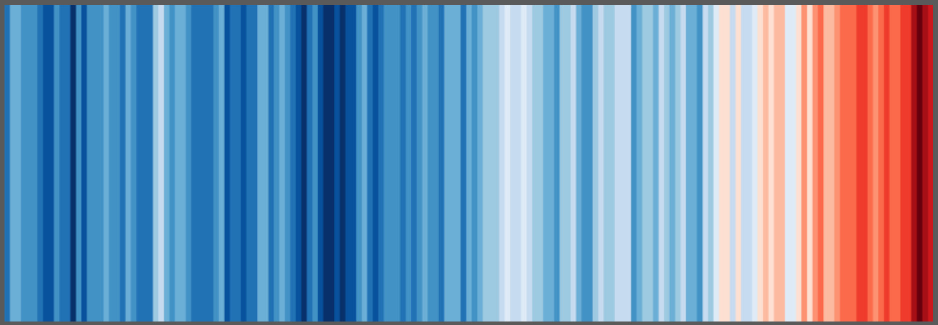

1 Reduce operational footprint
with better proportionality

- We can "sleep" at daily timescales
  one in many ideas for better proportionality
- We need your help
  to know if it is worth it

Tomorrow's Internet must
sleep more and grow old

to reduce its carbon footprint.

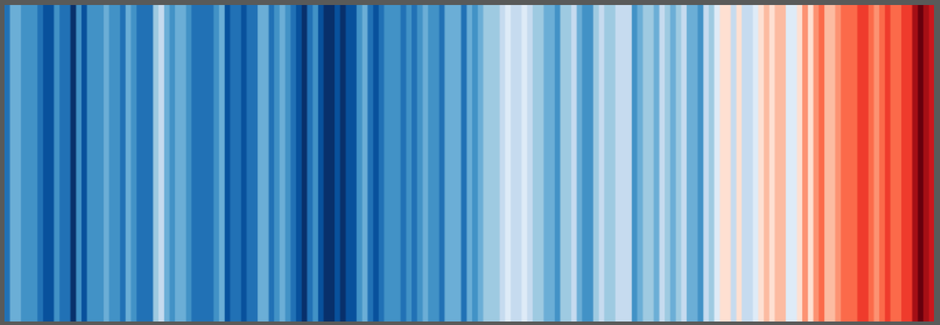

Reduce operational footprint
with better proportionality

2 Reduce embodied footprint
with sustainable procurement

Embodied carbon refers to the footprint of producing and recycling a product.

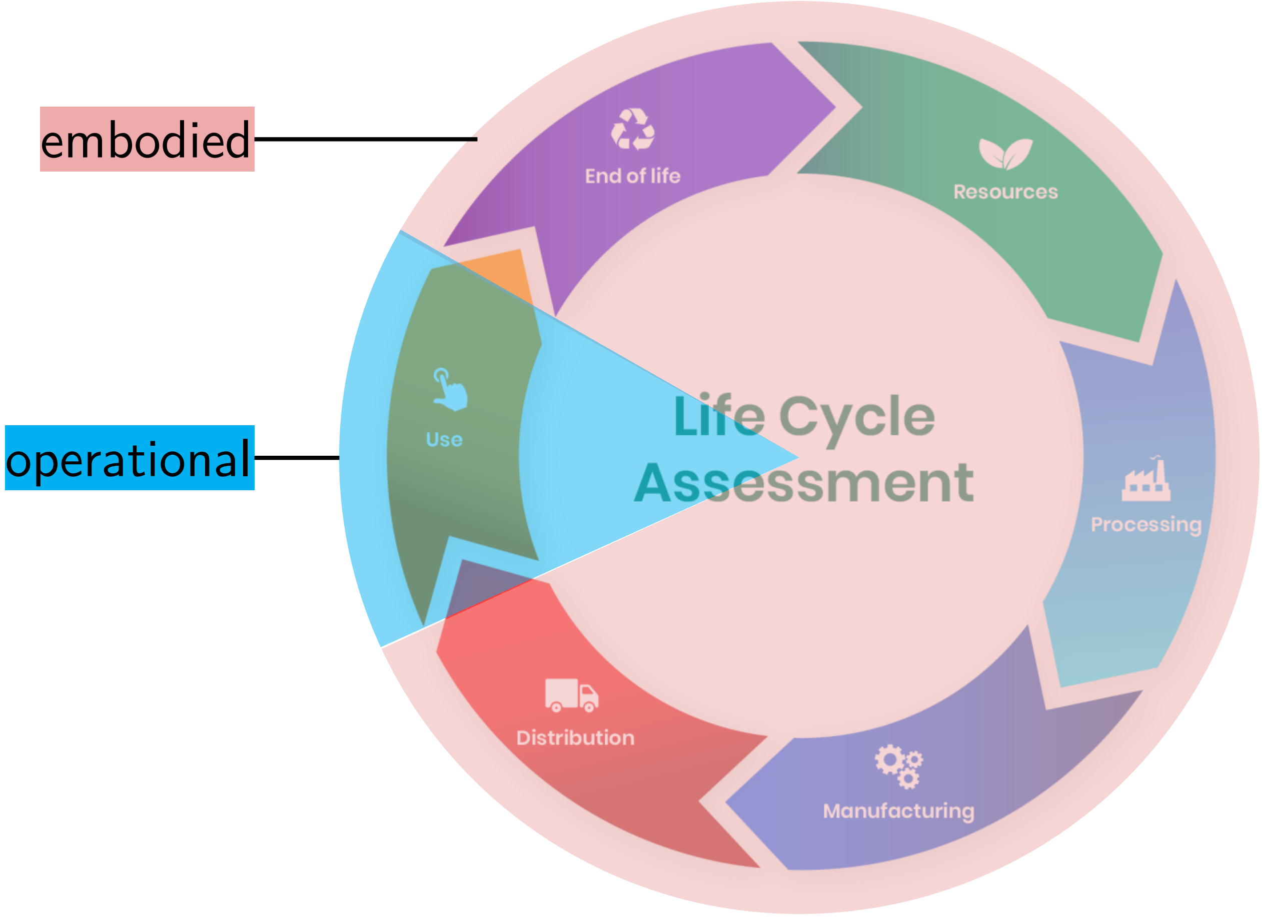

https://www.oneclicklca.com/life-cycle-assessment-explained/

For consumer devices, the embodied footprint dominates.

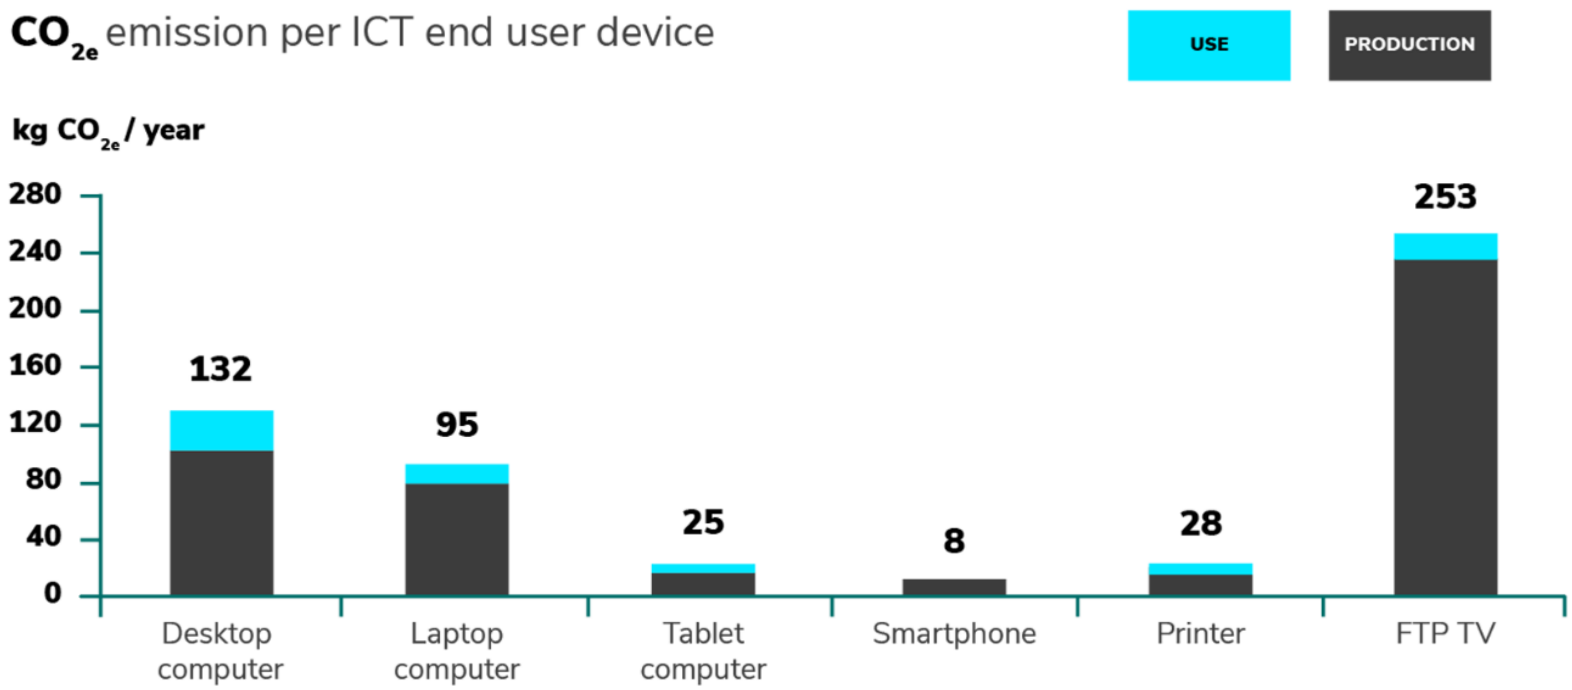

Green Software Foundation
greensoftware.org

https://learn.greensoftware.foundation/

For networked devices, it tends to be the opposite

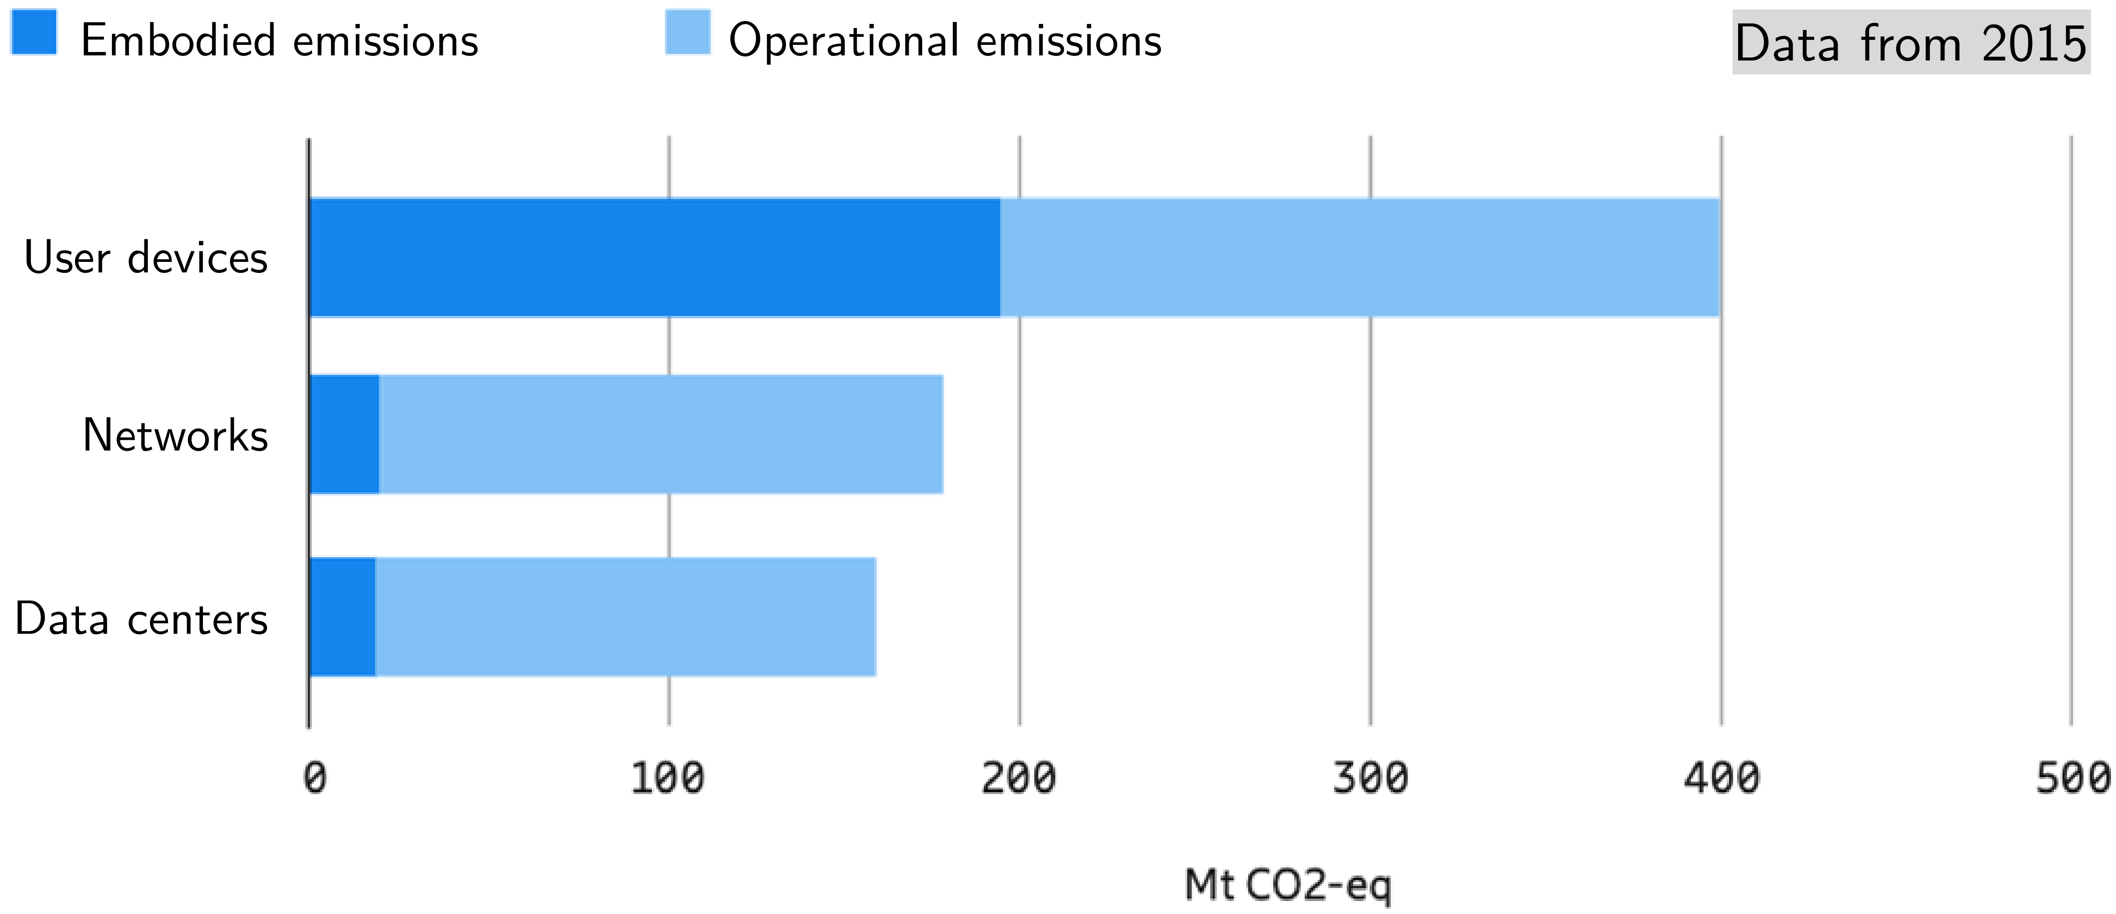

https://www.ericsson.com/en/reports-and-papers/industrylab/reports/a-quick-guide-to-your-digital-carbon-footprint

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

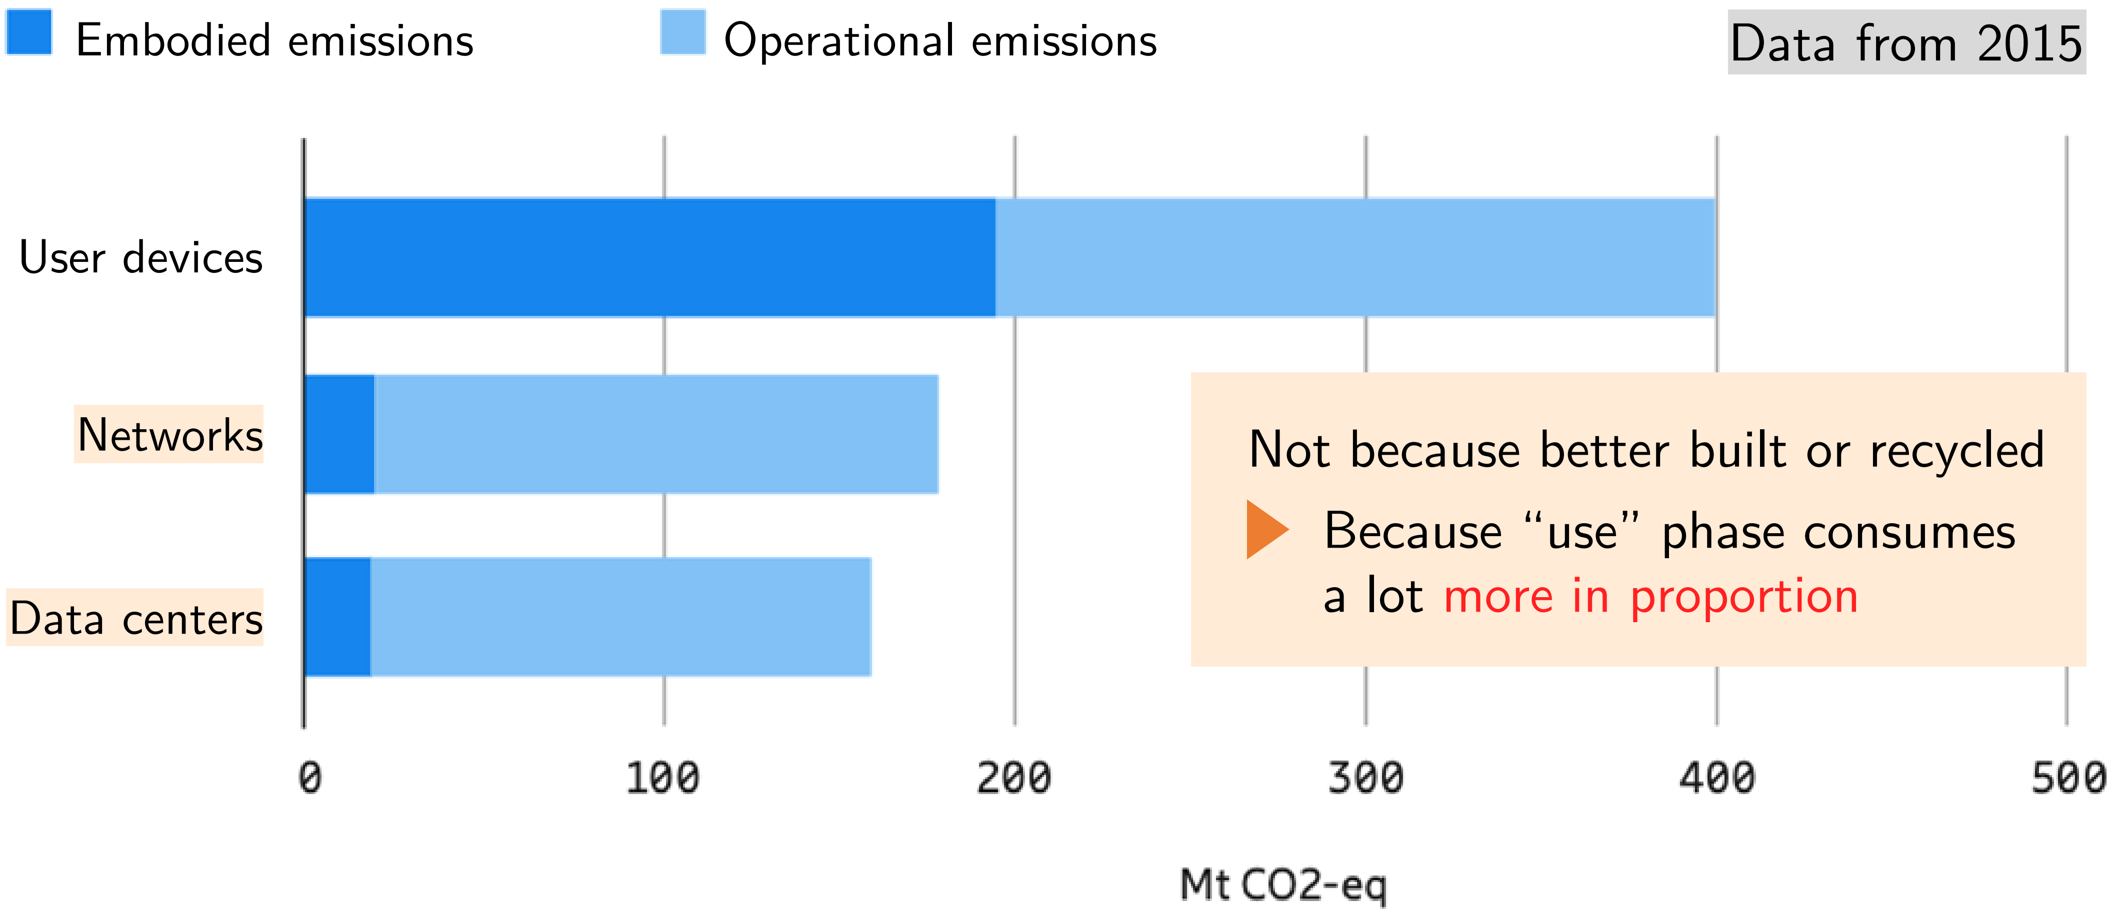

Not because better built or recycled

- Because "use" phase consumes a lot more in proportion

https://www.ericsson.com/en/reports-and-papers/industrylab/reports/a-quick-guide-to-your-digital-carbon-footprint

Reducing the embodied footprint is simple:
Use hardware longer.

Today

Refresh rates are
around 3-5 years only.

▶ Easy to extend

Useful Life of IT Network Equipment: Assets & Perspective
icorps Technologies, 02/2015, Online.

Reducing the embodied footprint is simple:
Use hardware longer.

~~Today~~ Recently	Refresh rates ~~are~~ were around 3-5 years only.

- Easy to extend

Useful Life of IT Network Equipment: Assets & Perspective
icorps Technologies, 02/2015, Online.

~~Today~~ Recently Refresh rates ~~are~~ were around 3-5 years only.

- Easy to extend

Okay, but

> Wouldn't this make networks
>
> - Less reliable
> - Less secure
> - Harder to manage
>
> ?

~~Today~~ Recently Refresh rates ~~are~~ were around 3-5 years only.

- Easy to extend

Okay, but

> Wouldn't this make networks
> - Less reliable
> - Less secure
> - Harder to manage
>
> ?

Not necessarily.

“Older” networks are not necessarily less reliable.

> The vast majority of network hardware failures take place within the first 30 days of installing brand new, out-of-the-box network hardware.

CXTEC

Surprising truth about network hardware failures.
CXTEC, 03/2022, Online.

“Older” networks are not necessarily less reliable.

> The vast majority of network hardware failures take place within the first 30 days of installing brand new, out-of-the-box network hardware.
>
> CXTEC

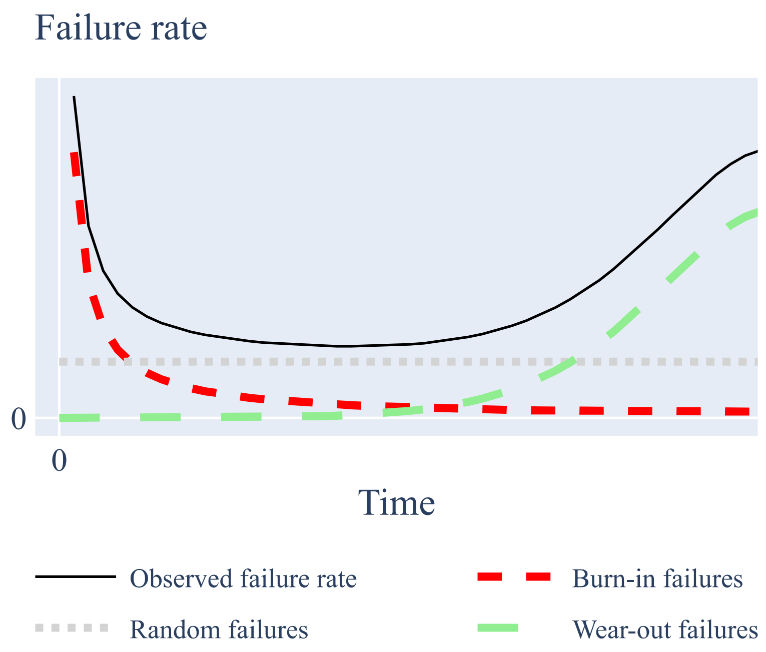

Surprising truth about network hardware failures. CXTEC, 03/2022, Online.

Manufactured products typically fail following a “bathtub” profile.

Devices that never failed in 3 years are unlikely to fail anytime soon after.

Two hints in that direction

- Main vendors usually provide 5 year support after end-of-sale.

Devices that never failed in 3 years are unlikely to fail anytime soon after.

Two hints in that direction

- Main vendors usually provide 5 year support after end-of-sale.
- Specilized companies even provide unlimited warranties for refurbished network hardware.

We must understand better the aging of networking devices.

- What are the practical consequences of operating older devices?
- When do aging effects appear?

We must understand better the aging of networking devices.

- What are the practical consequences of operating older devices?
- When do aging effects appear?

When does it really make sense to renew networking hardware?

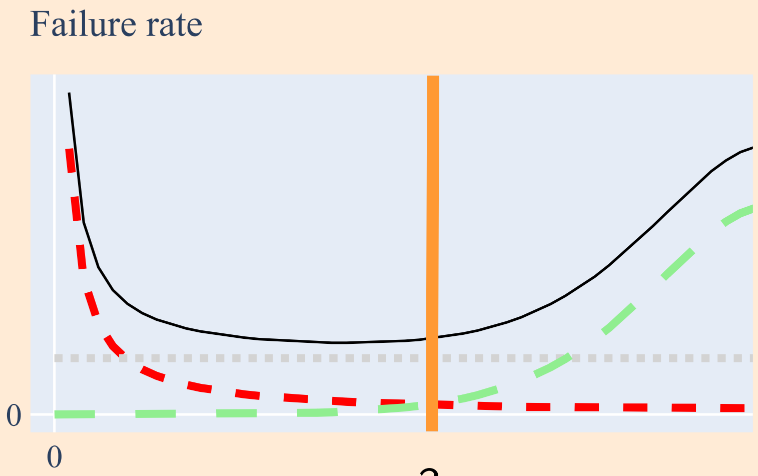

3 years
5
15
more ?

To answer that question, we need data.

When do you renew your hardware?

Why do you renew?

What failures do you see in practice?

When and where do they occur?

When does it really make sense to renew networking hardware?

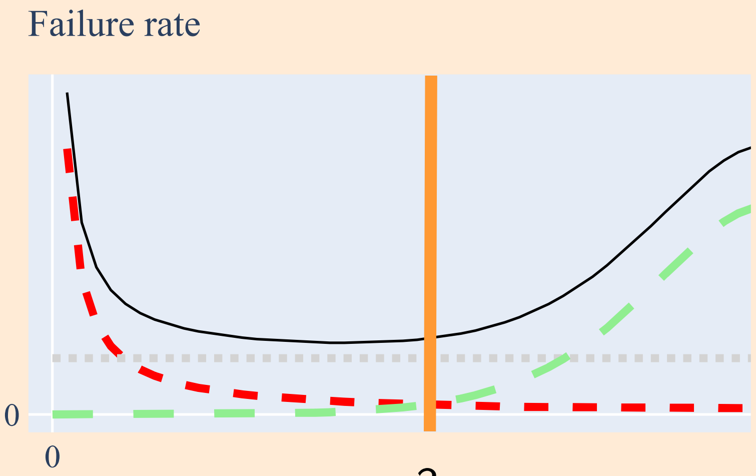

3 years
5
15
more ?

Okay, but

> Wouldn't this make networks
>
> - Less reliable
> - Less secure
> - Harder to manage

Let's talk offline.

Not necessarily.

Tomorrow's Internet must

sleep more and grow old

to reduce its carbon footprint.

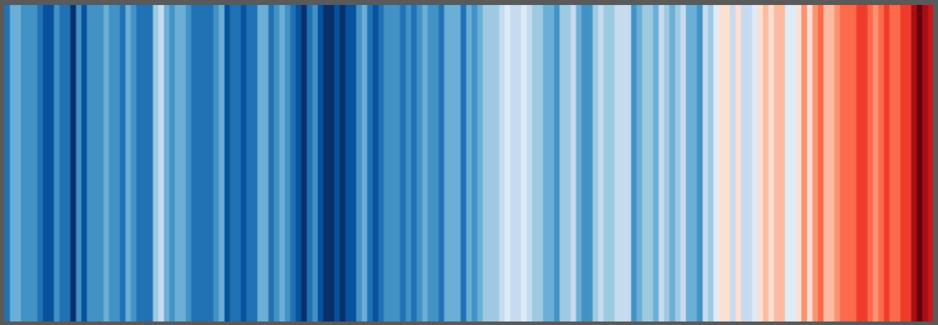

2 Reduce embodied footprint
with sustainable procurement

- You should renew when really needed
  which saves both carbon and money
- We can help you assess when that is

We need your help
to help your network.

We need your help to help your network.

We need data.

Testing	Measurement	Procurement
Dev. data	Ops. data	CapEx data

We need your help
to help your network.

We need data.

- Academics have ideas
  sometimes even good ones!
- Operators have power

We need your help to help your network.

We need data.

- Academics have ideas
  sometimes even good ones!
- Operators have power to pay for every month.

We need your help to help your network.

We need data.

- Academics have ideas
  sometimes even good ones!
- Operators have power to pay for every month.
  to change things in their network.

Let's work together

Yes, we know what NDAs are.

See you @ the BoF
Today, 18h, Side room

Tomorrow's Internet must sleep more and grow old to reduce its carbon footprint.

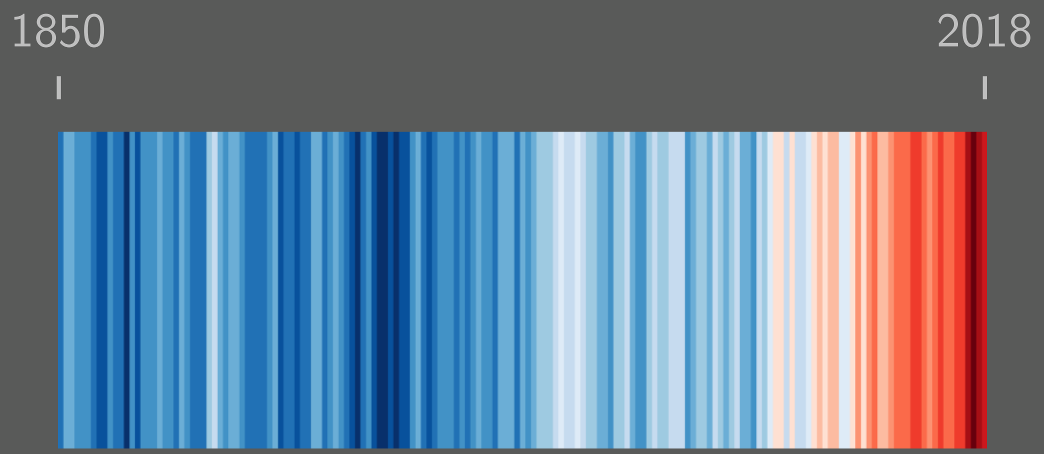

Climate stripes. Ed Hawkins, 2018
portrays the increase of average global temperature

Romain Jacob
jacobr@ethz.ch

Laurent Vanbever
lvanbever@ethz.ch

RIPE 87

Nov. 27, 2023