The Internet is getting emptier; that's a sustainability problem

Romain Jacob
ETH Zürich

GreenIO - Paris
Dec. 4, 2024

What do you think consumes more energy?

Data Centers

Telco Networks

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
≪ work done.	+340%	in workload	+600%	in traffic

Energy efficiency improved a lot but not enough!

	Data Centers	Telco Networks
Change in energy	+20-70% in energy	+18-64% in energy

> 0 !

“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” and carbon footprint.

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

▶ Producing electricity emits carbon.

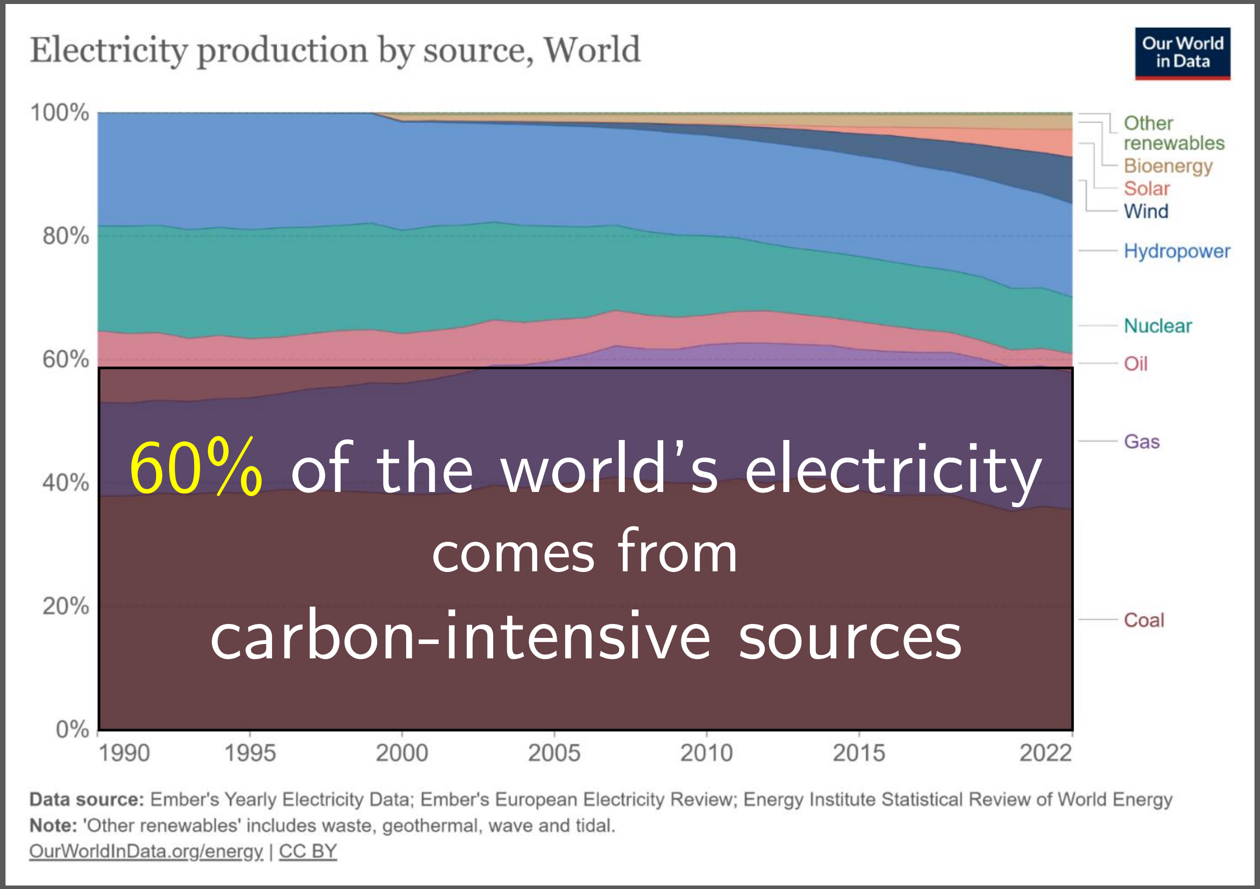

https://ourworldindata.org/grapher/electricity-prod-source-stacked

"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

- Producing electricity emits carbon.
- Total electricity usage is likely to keep increasing.

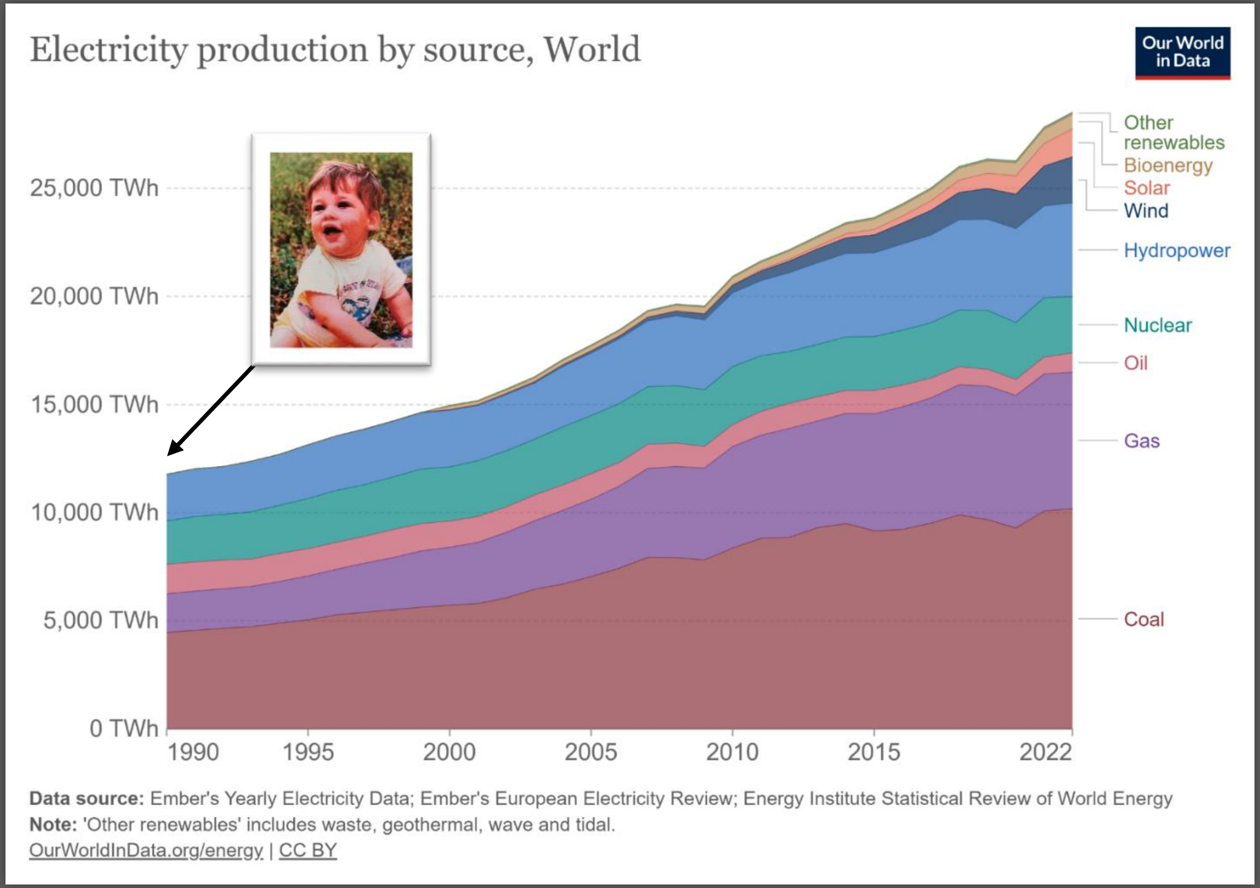

It doubled in my lifetime.

Internet access is
still
far from universal.

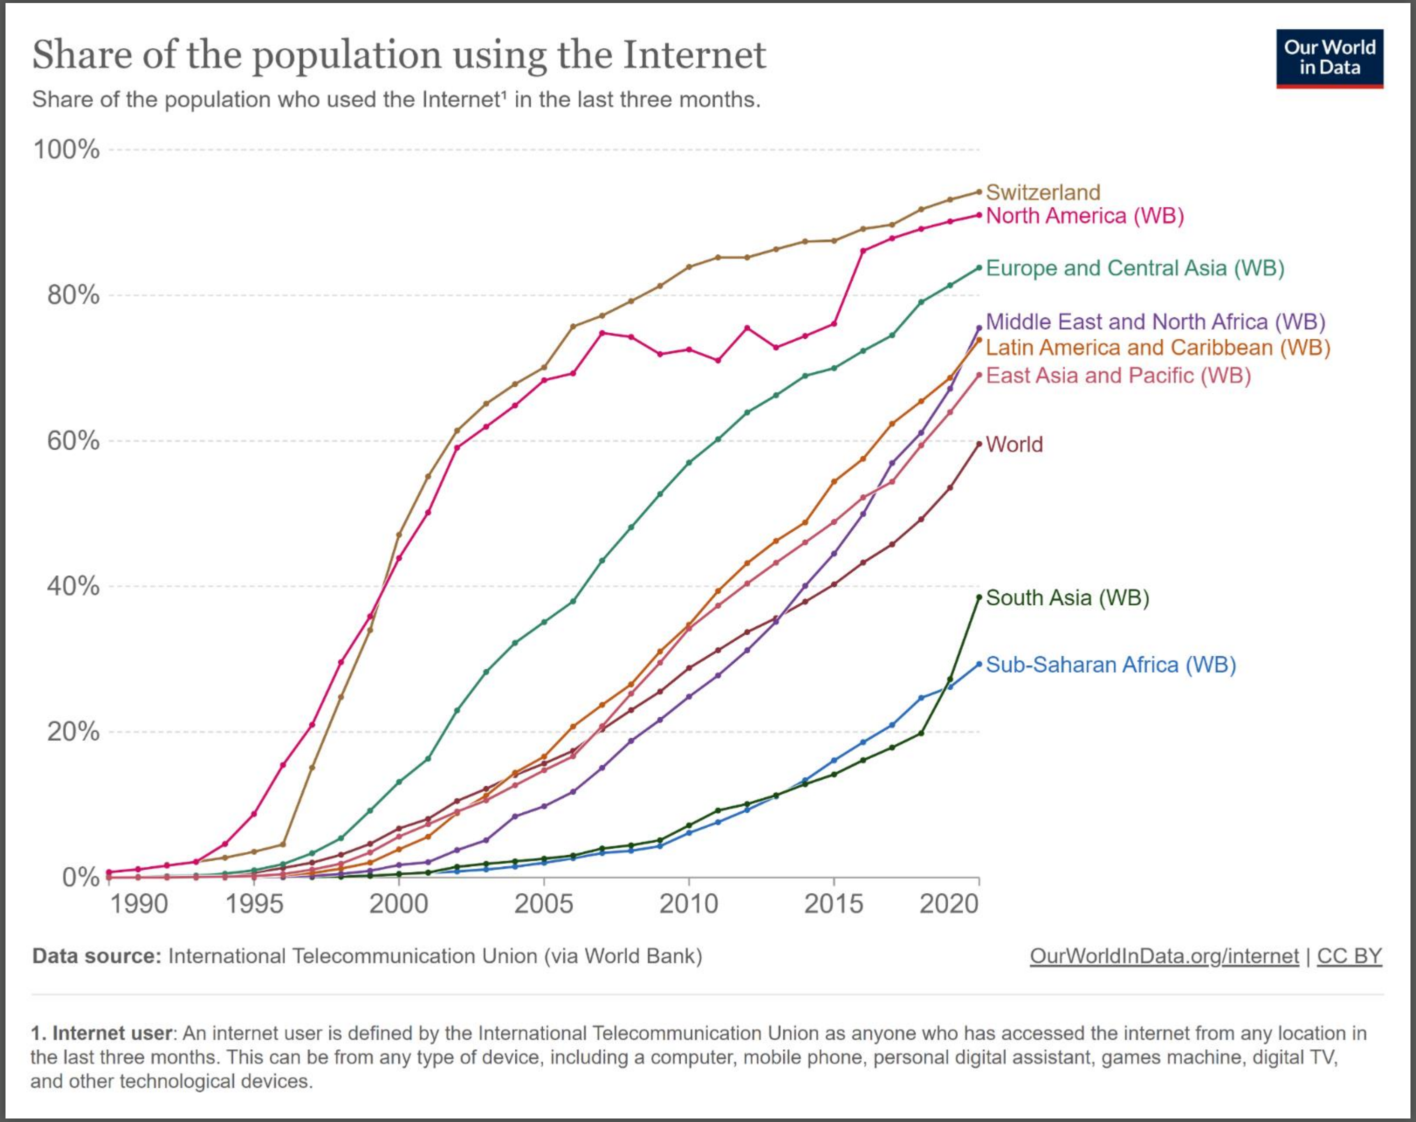

https://ourworldindata.org/grapher/share-of-individuals-using-the-internet

Let’s zoom in.

What’s going on in the Internet?

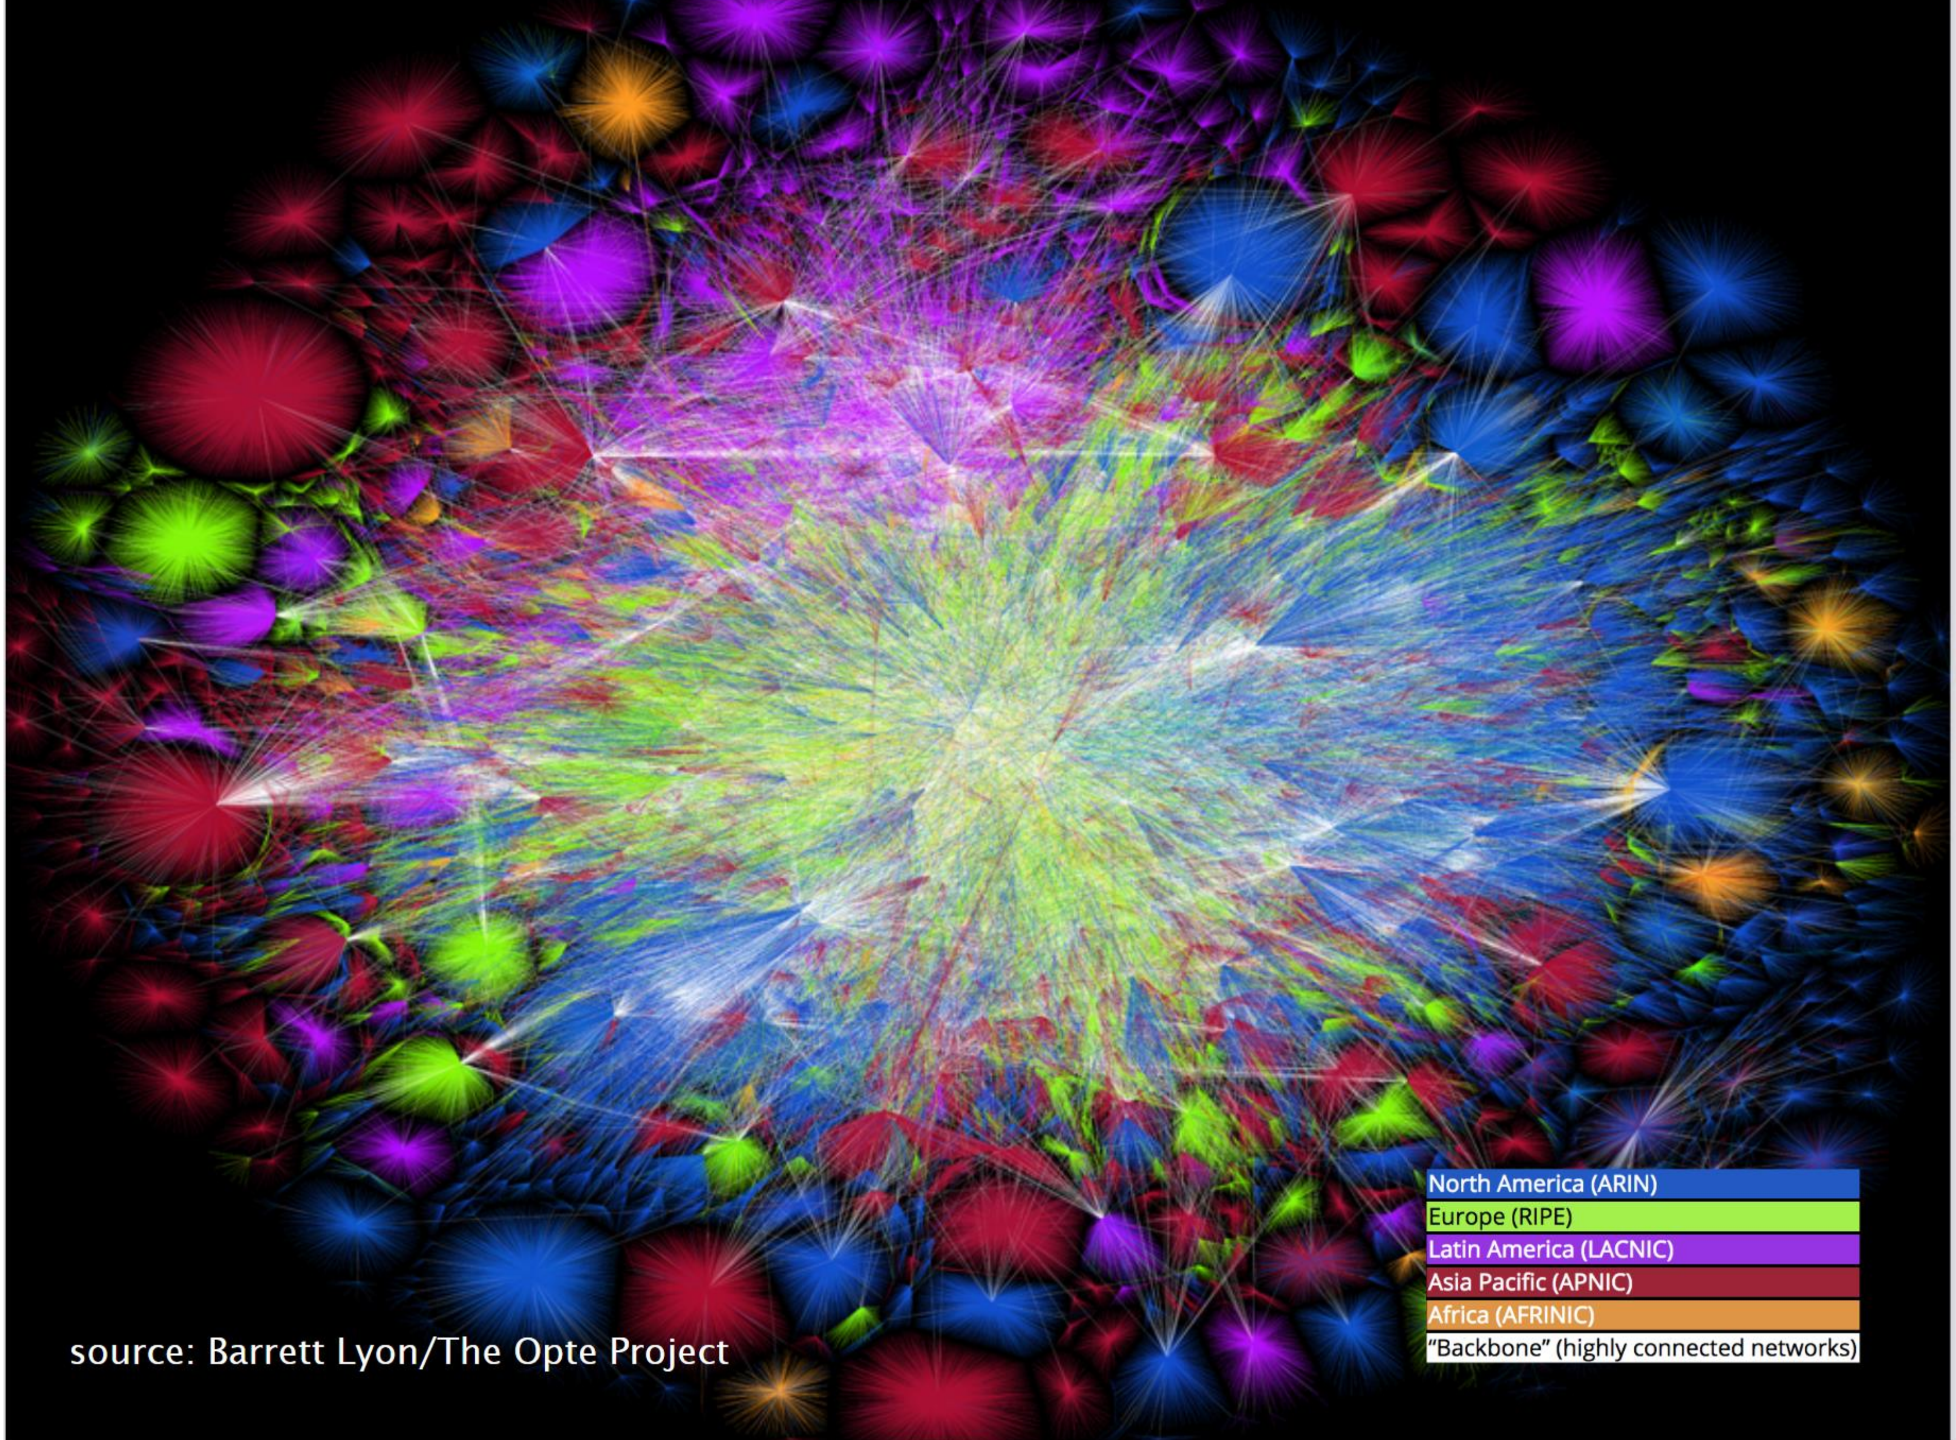
North America (ARIN)
Europe (RIPE)
Latin America (LACNIC)
Asia Pacific (APNIC)
Africa (AFRINIC)
“Backbone” (highly connected networks)
source: Barrett Lyon/The Opte Project

the "map"

Here is the topology of a middle size Swiss network: Switch.

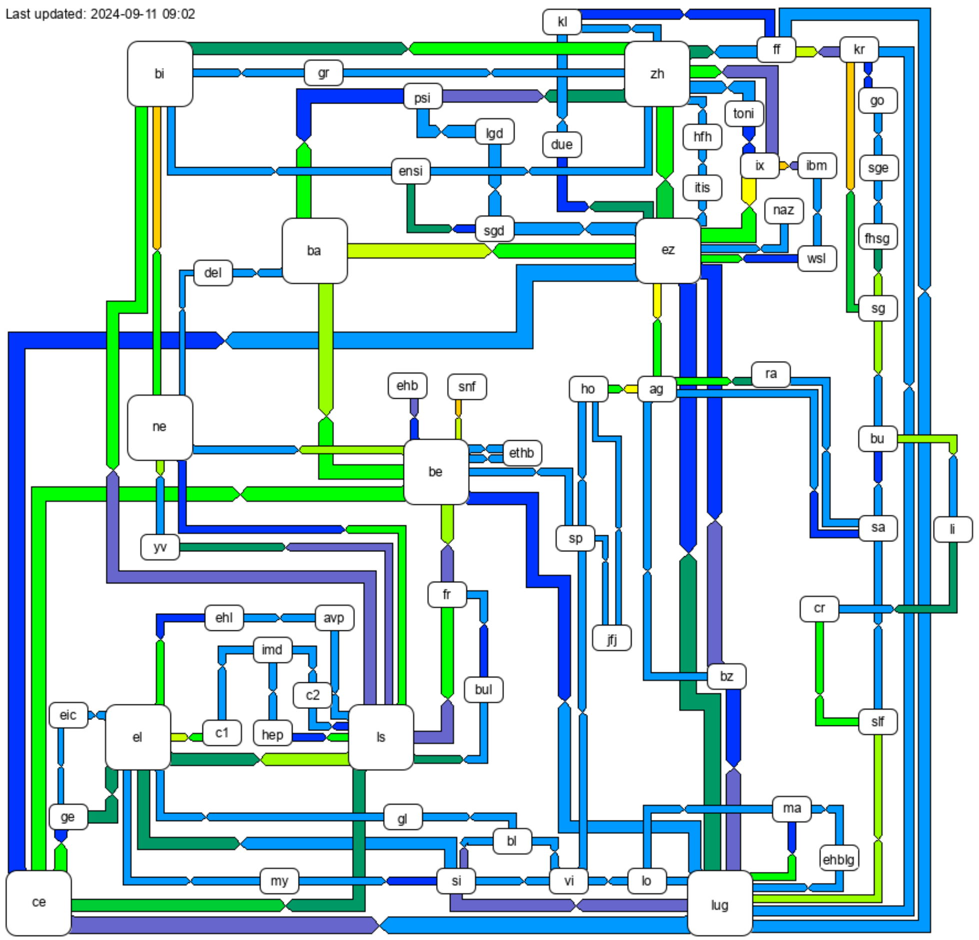

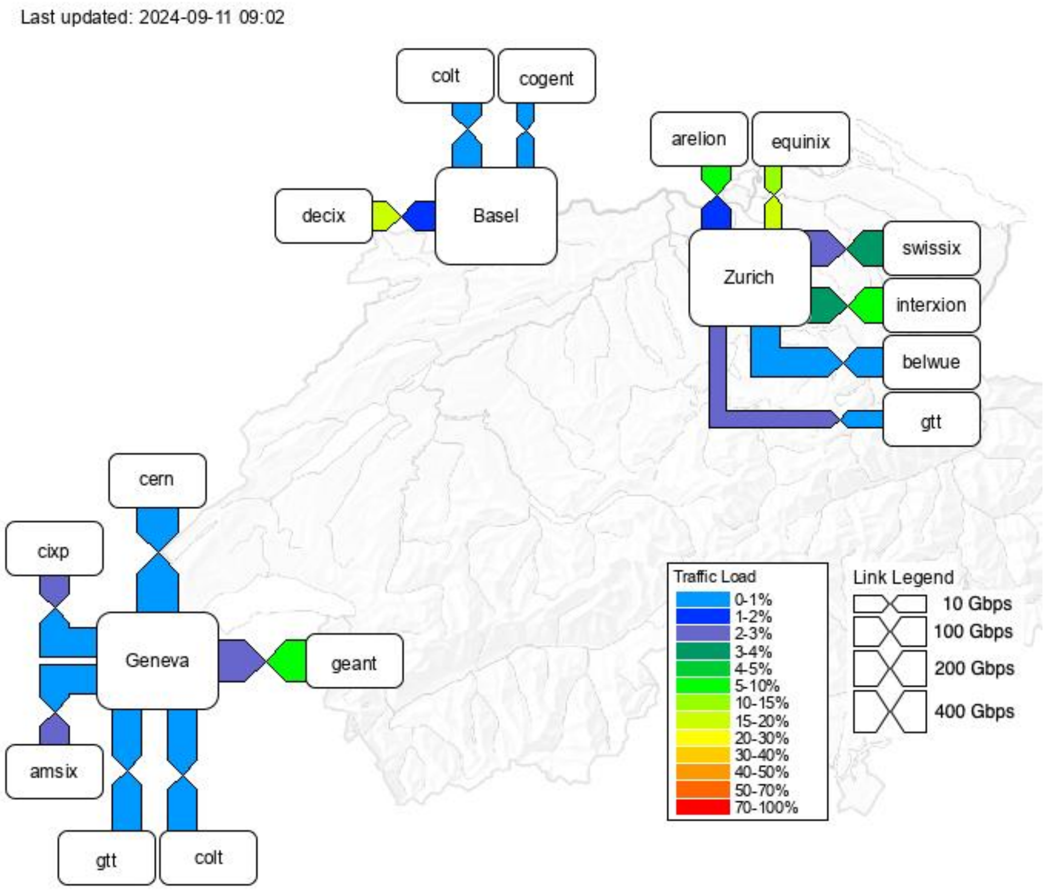

Average link utilization in the Switch network [%]

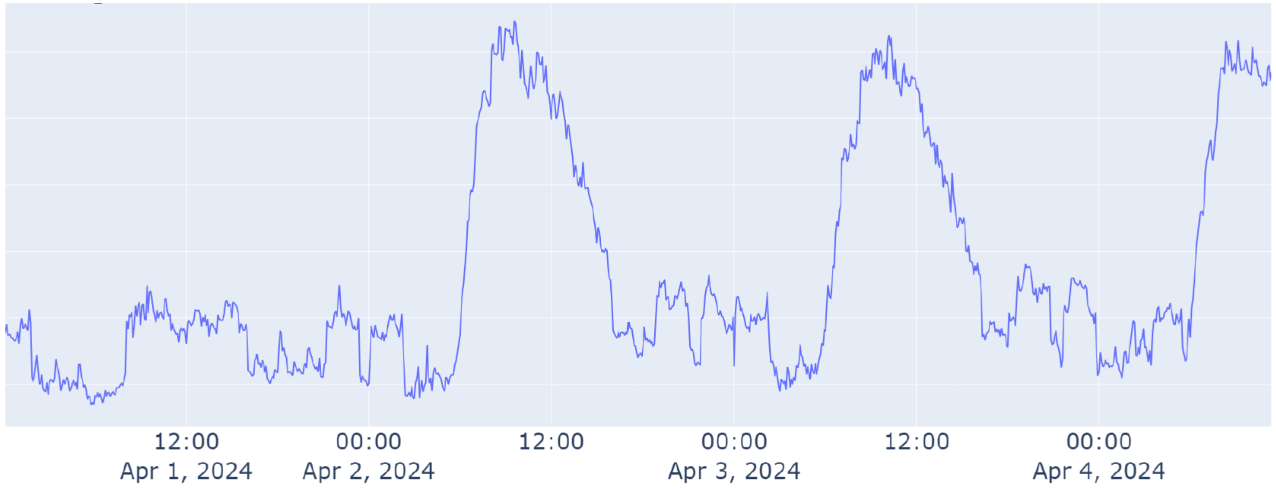

What do you think is the average link load in Switch?

Average link utilization in the Switch network [%]

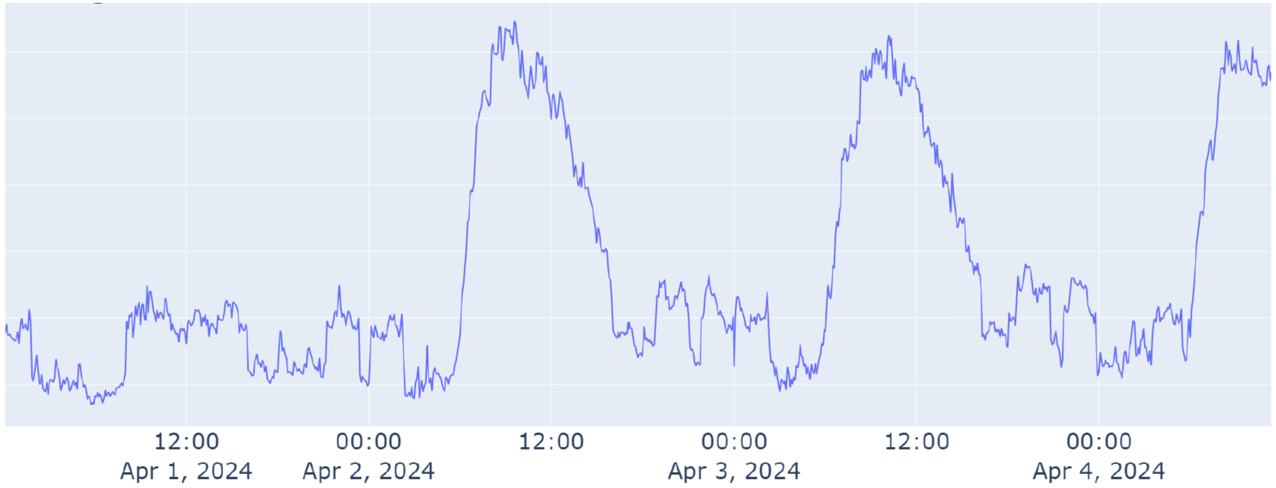

What do you think is the average link load in Switch?

2.1%

2.5 months of data
internal links only

Average link utilization in the Switch network [%]

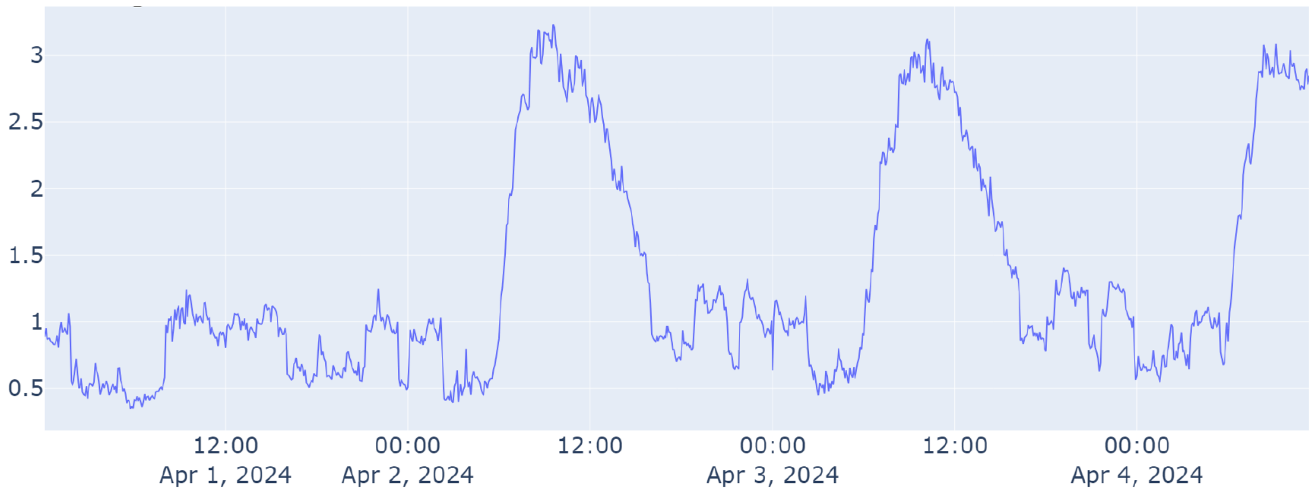

While we cannot generalize, single digit numbers are not rare.

2.1%

2.5 months of data
internal links only

Average link utilization in the Switch network [%]

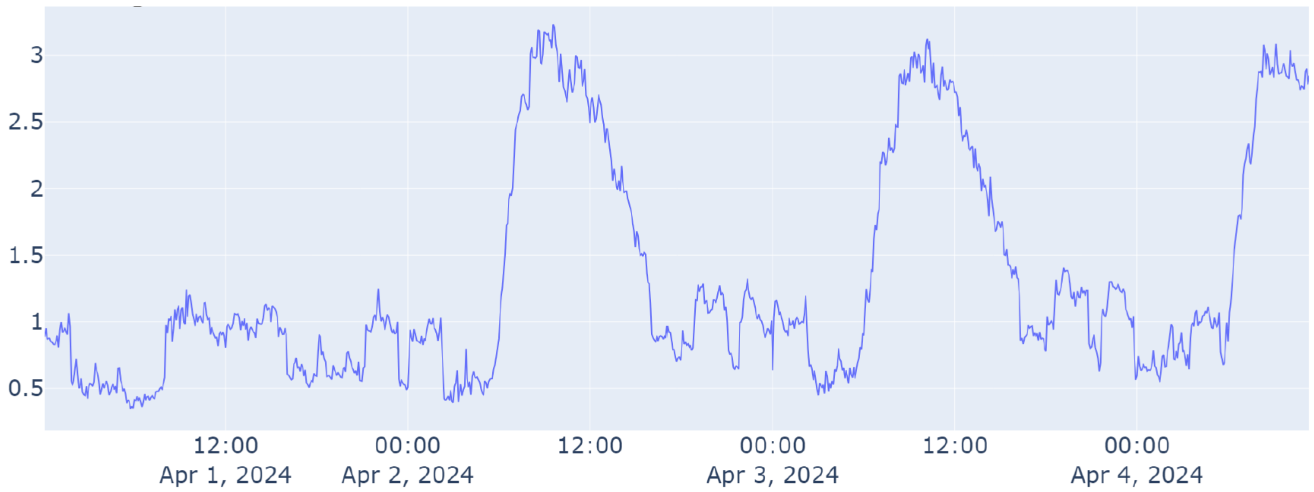

The total volume of Internet traffic increases over time.

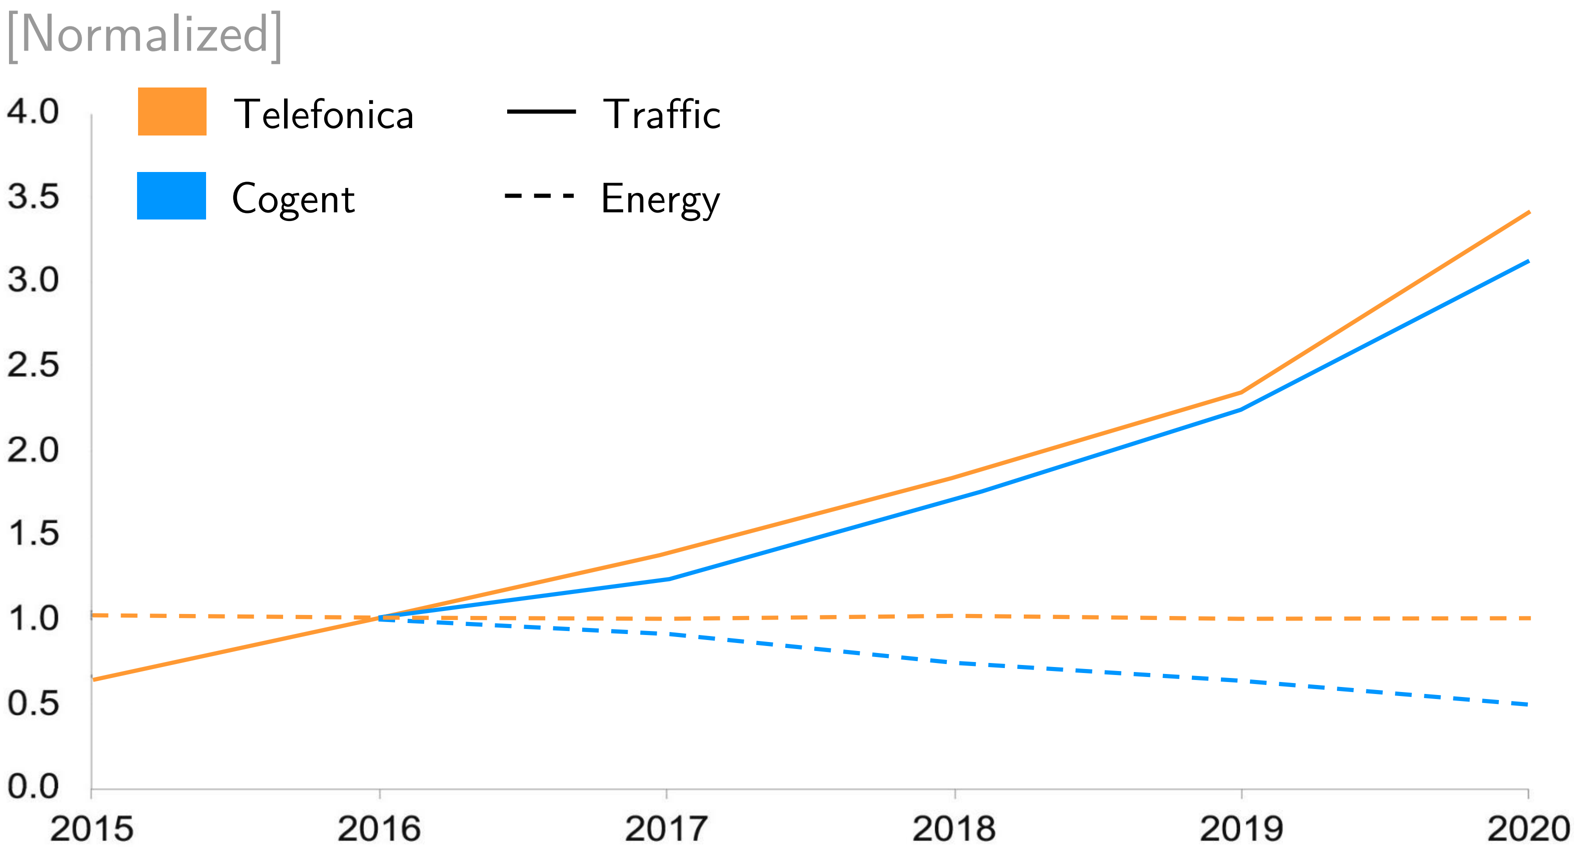

https://www.sciencedirect.com/science/article/pii/S2542435121002117

But that increase happens only at some times of the day.

[Normalized]

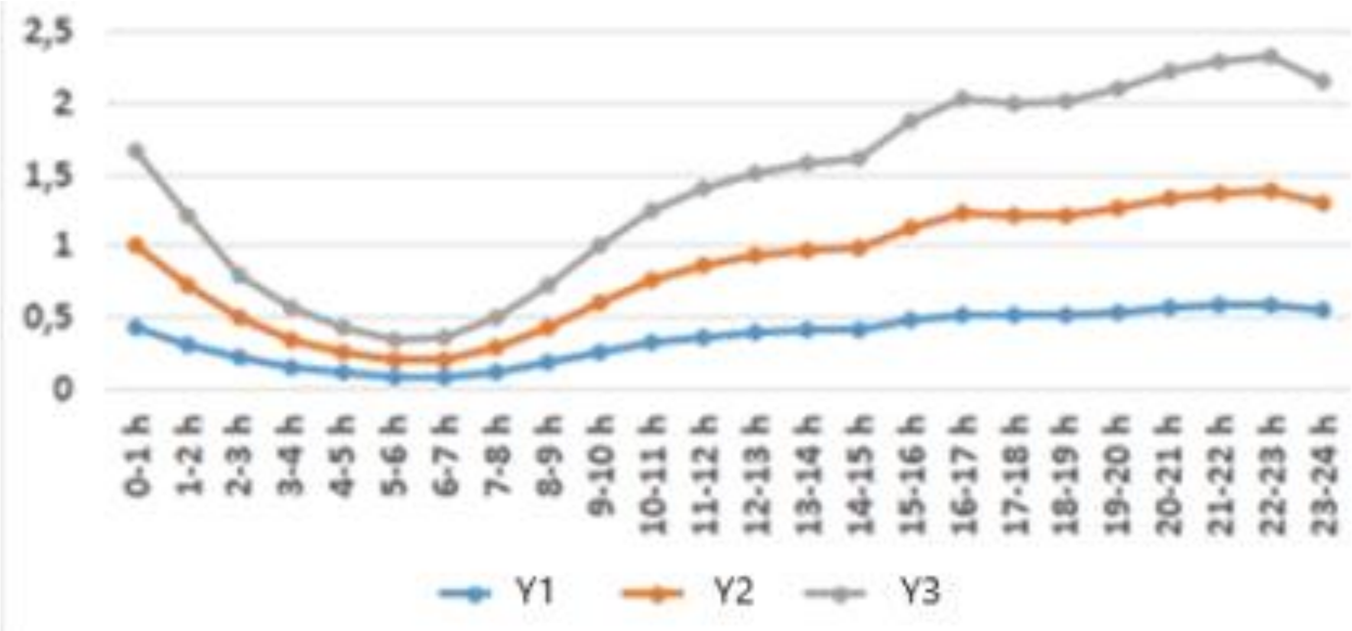

Telefonica

https://eng.ox.ac.uk/media/qhwpxjua/david_sustainability_on_carrier_transport_networks.pdf

Today

Empty

Full

As the Internet gets bigger, the difference between “empty” and “full” increases.

Today

Tomorrow

Empty

Full

Networks are getting emptier, on average.

Why does it matter?

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 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 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.

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).

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.

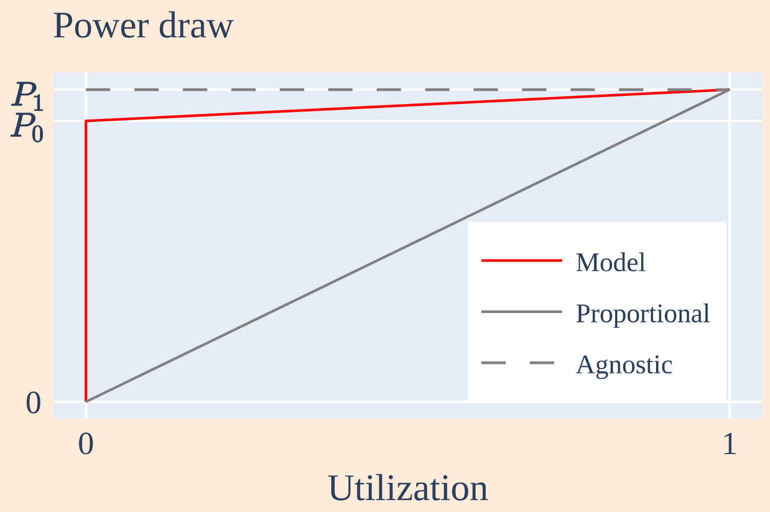

The Internet is not really a network of roads

Roads do not consume energy
when they are unused.

The Internet is not really a network of roads; more like Hyperloop tubes.

Hyperloop tubes require loads of power to create vacuum, whether used or not.

Roads do not consume energy when they are unused.

1 Network devices are always "on."

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

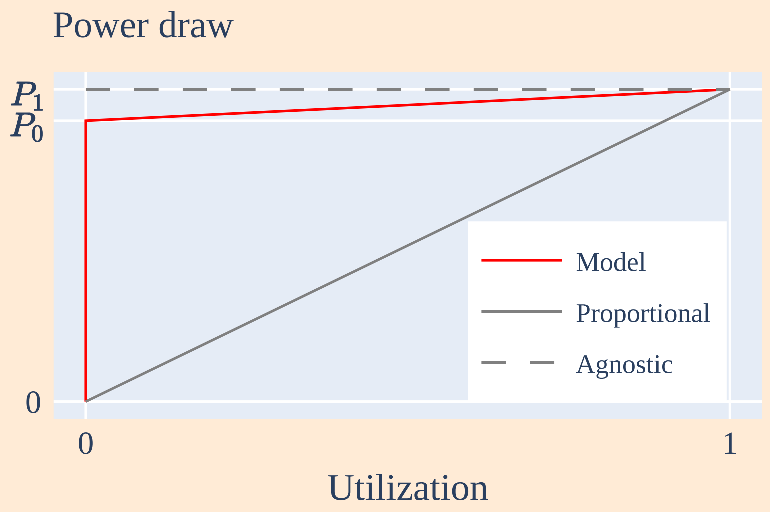

1 Network devices are always "on."

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

3 Network devices are under-utilized.

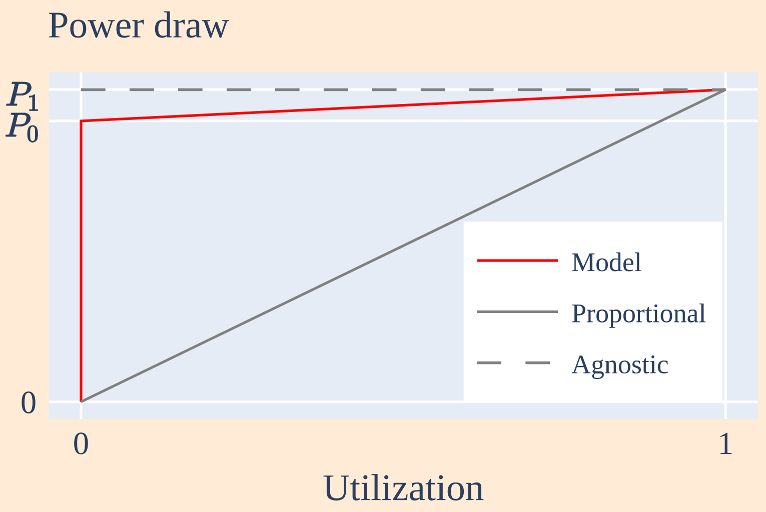

ISP overprovision networks to support

- 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.

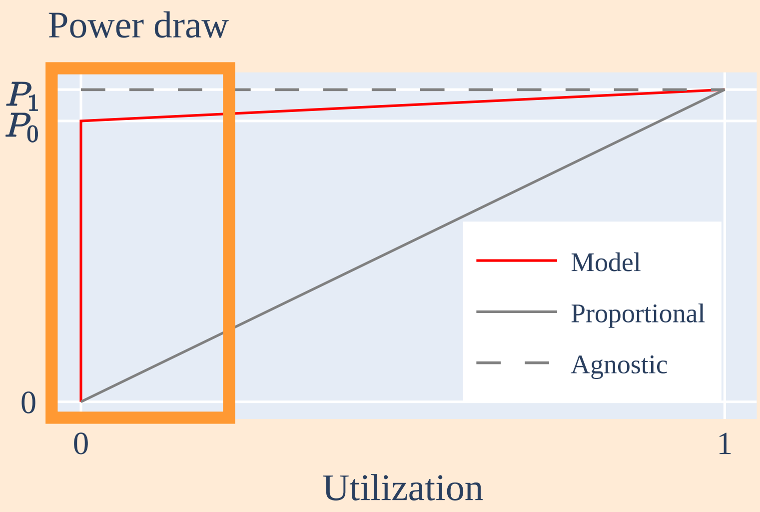

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

1 Network devices are always “on.”

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

3 Network devices are under-utilized.

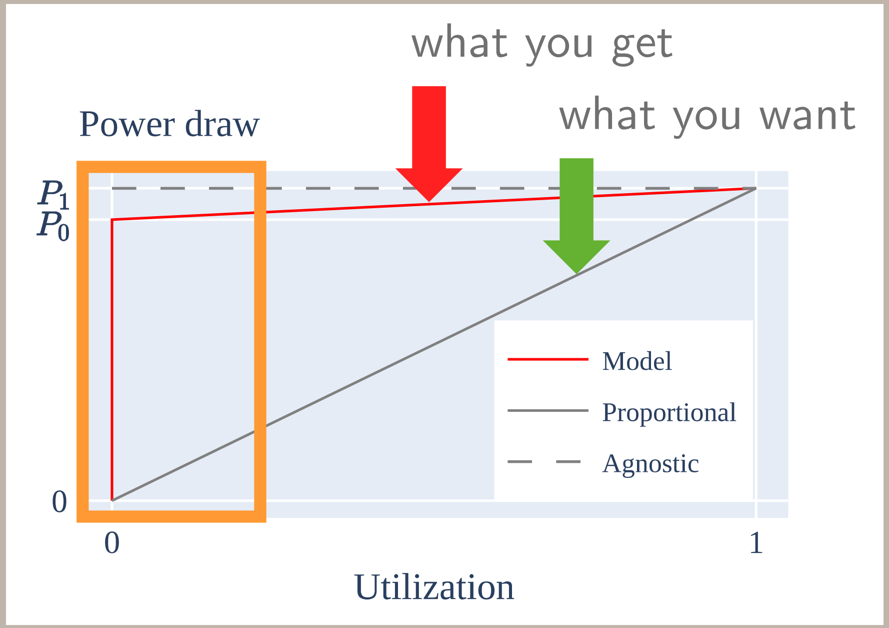

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

There two ways to improve energy efficiency

- Run more often at high utilization

  "Buffer-and-Burst"

  Time-shifting

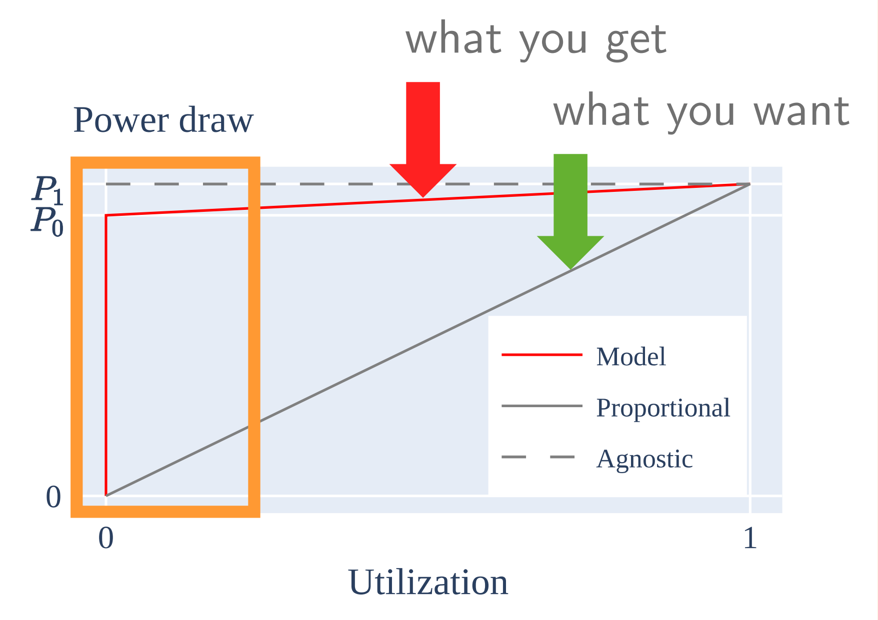

There two ways to improve energy efficiency

- Run more often at high utilization

  "Buffer-and-Burst"

  Time-shifting

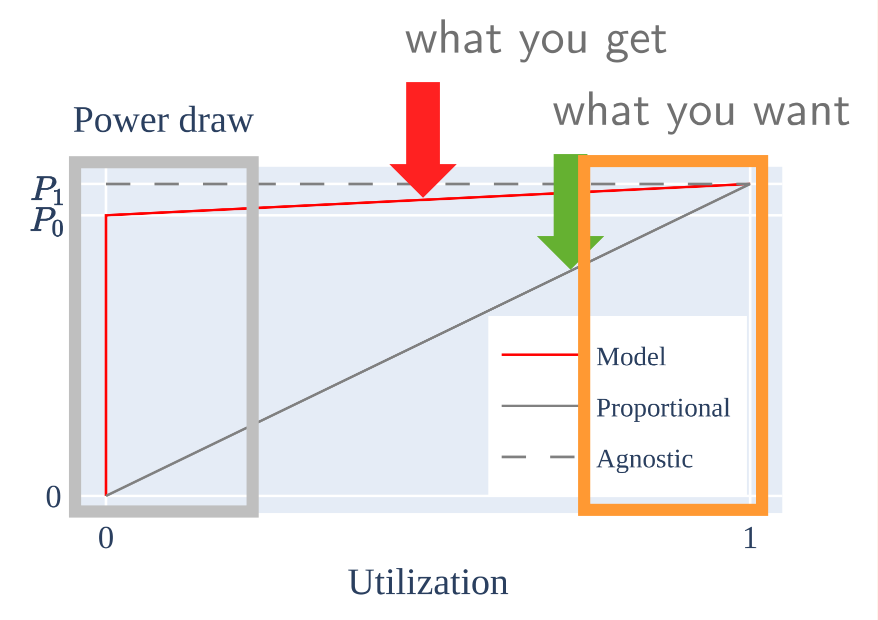

There two ways to improve energy efficiency

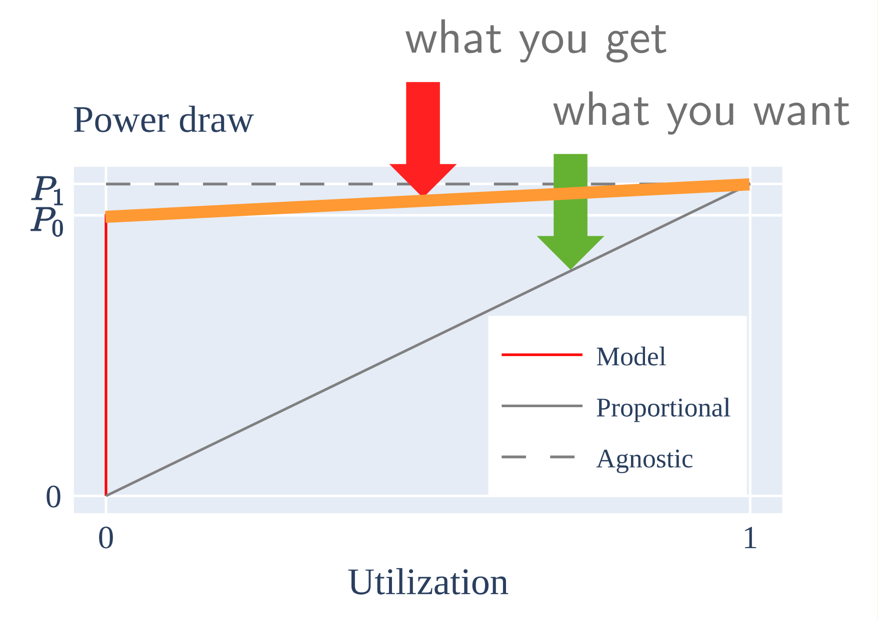

- Run more often at high utilization

  "Buffer-and-Burst"

  Time-shifting

- Take low-utilization power down

There two ways to improve energy efficiency

- Run more often at high utilization

  “Buffer-and-Burst”

  Time-shifting

- Take low-utilization power down

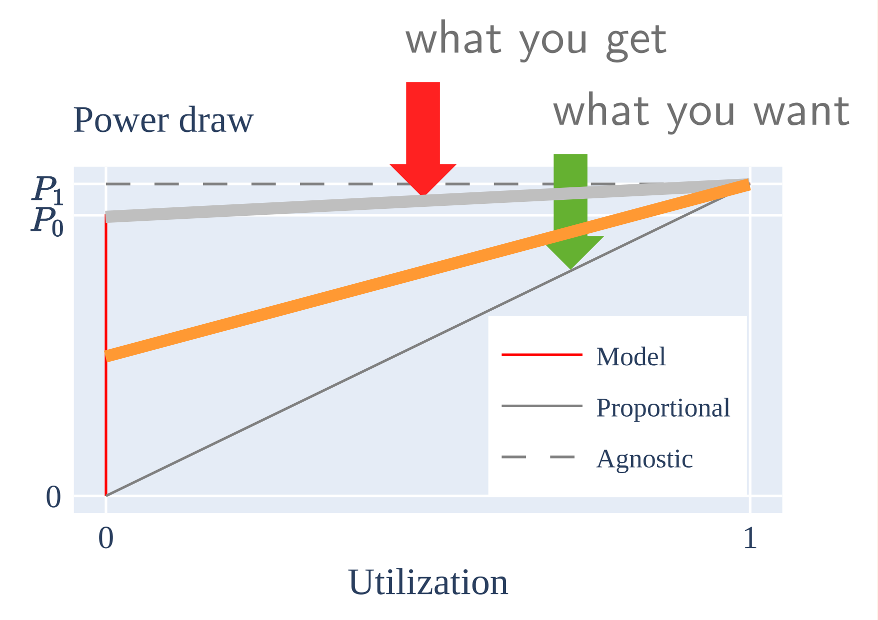

The basic idea is to turn off "stuff" whenever possible.

Router

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

Router

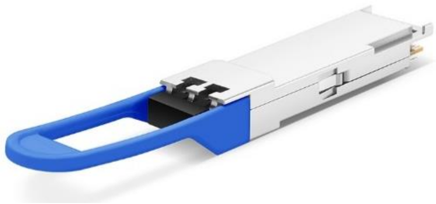

Transceiver

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

What can we possibly turn off?

- Ports
- Line cards
- Entire device...

Router

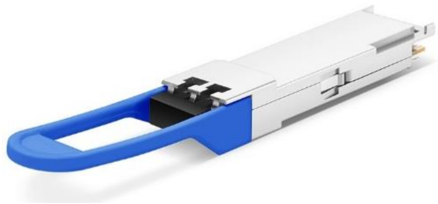

Transceiver

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.

Router

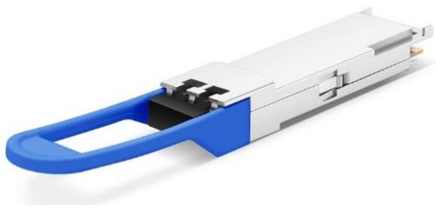

Transceiver

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

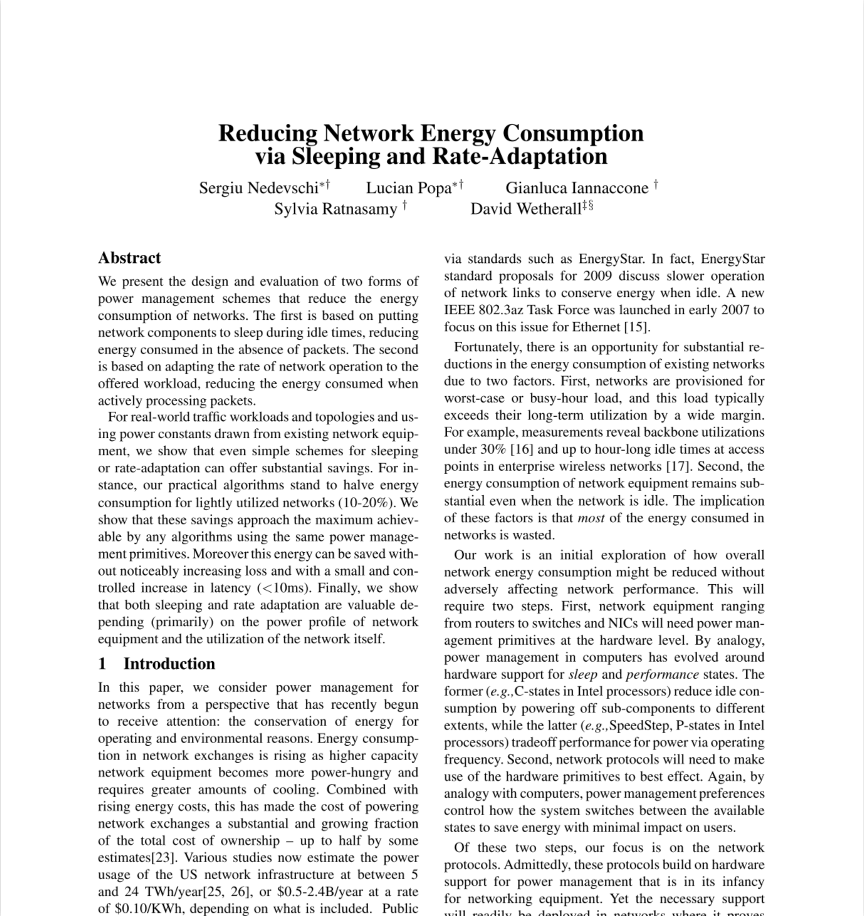

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

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

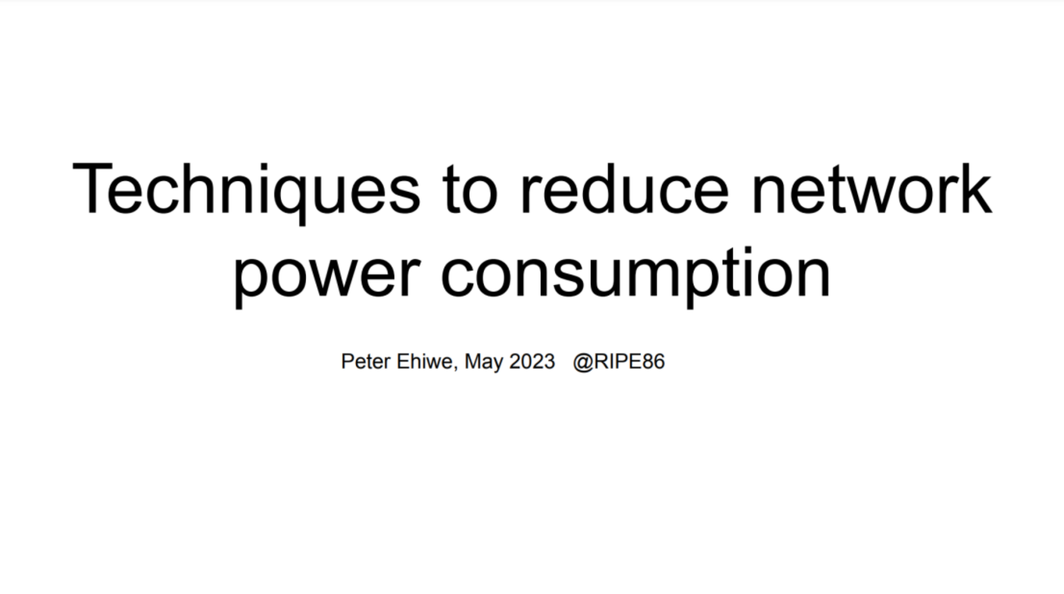

The theory says we can save tens of energy % in ISP networks.

Academia NSDI 2008

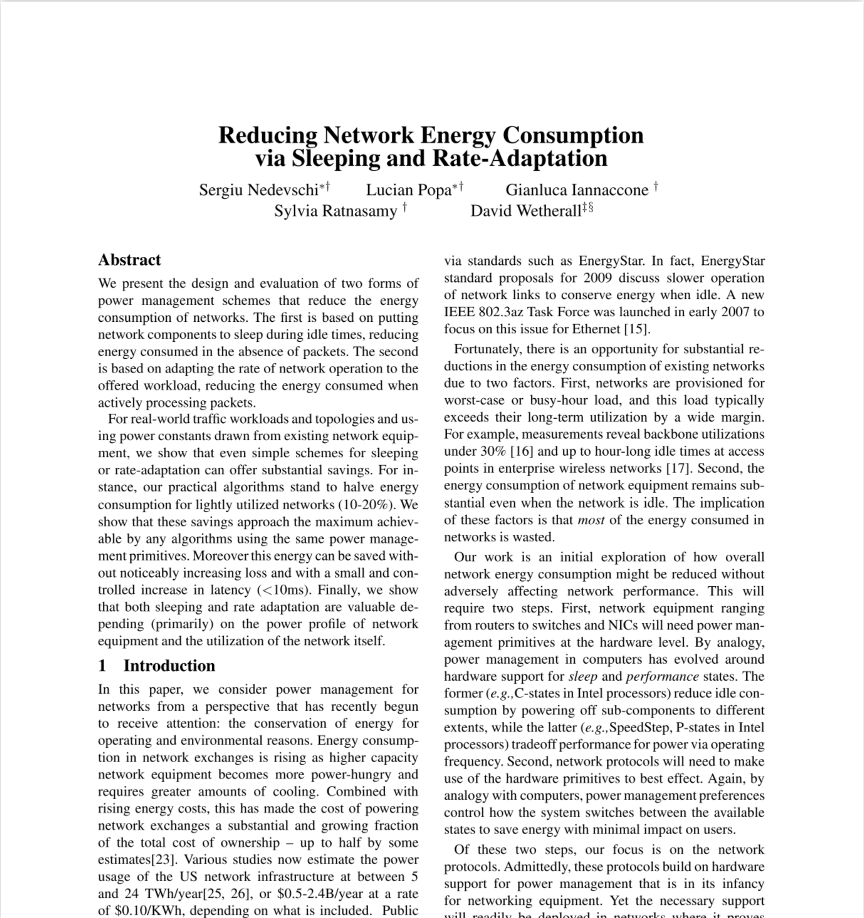

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

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

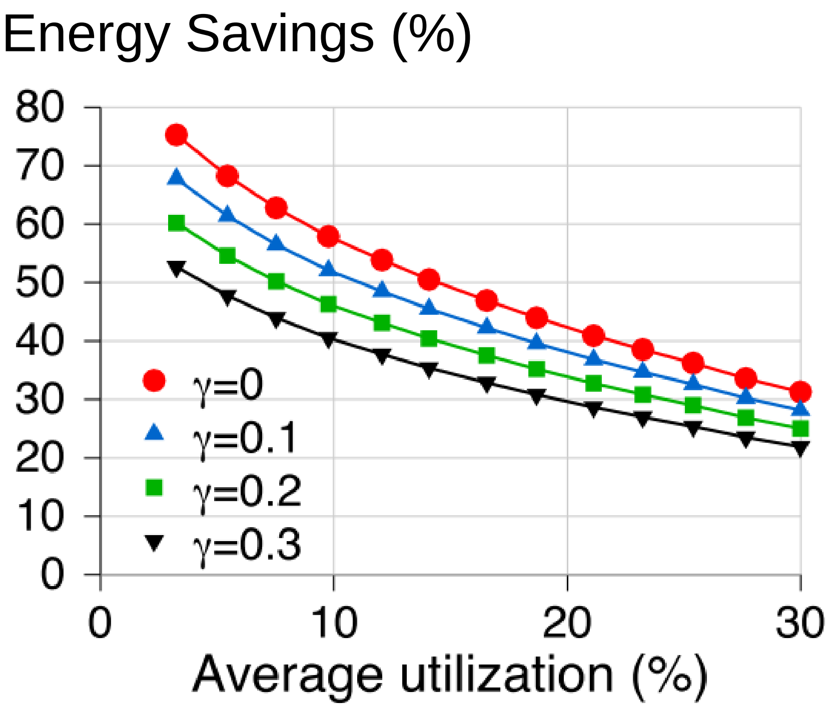

The theory says we can save tens of energy % in ISP networks.

Academia NSDI 2008

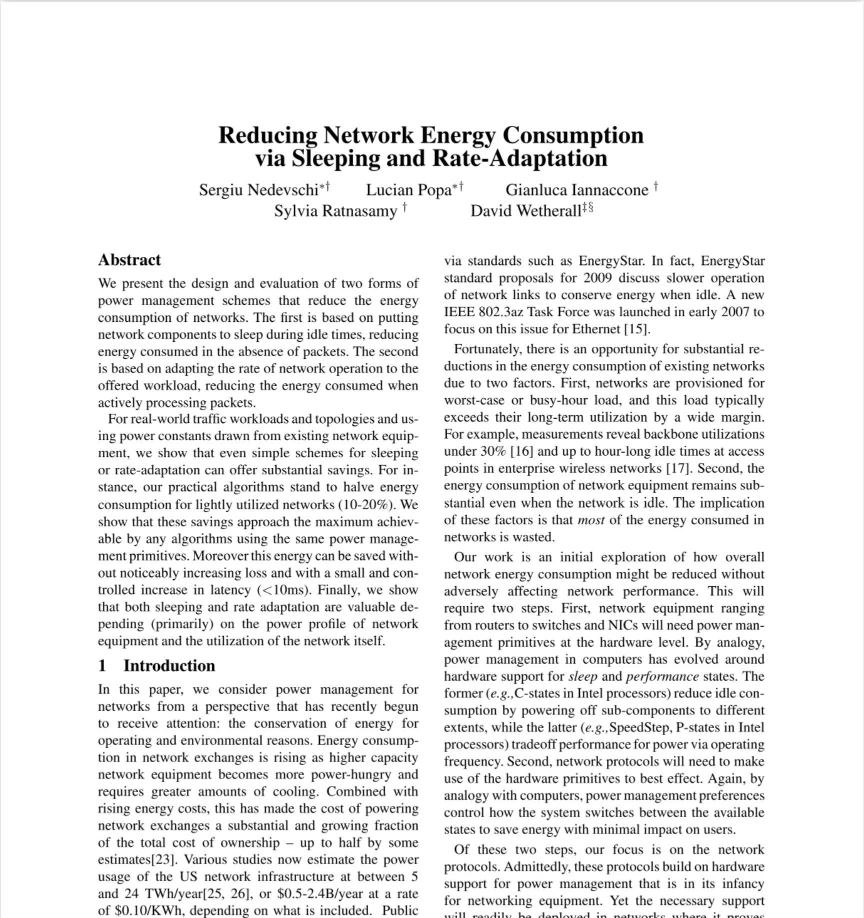

Reducing Network Energy Consumption via Sleeping and Rate-Adaptation

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

How? Buffer-and-Burst

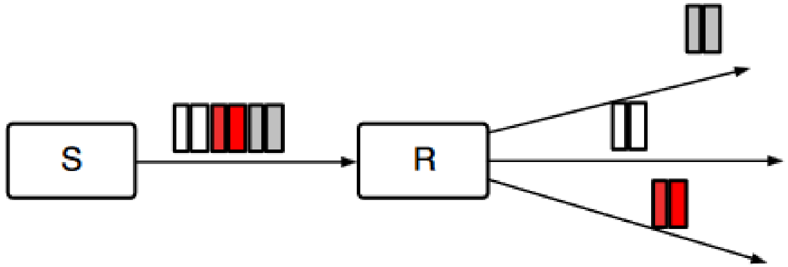

Assumes

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

Theory

How? Buffer-and-Burst

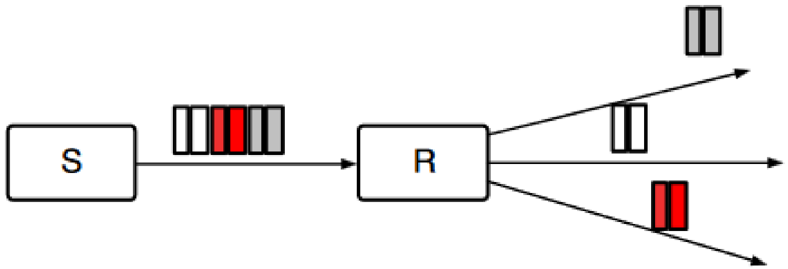

Assumes

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

VS

Practice

Wake-up delay (s)

Measured on Cisco Nexus 9300

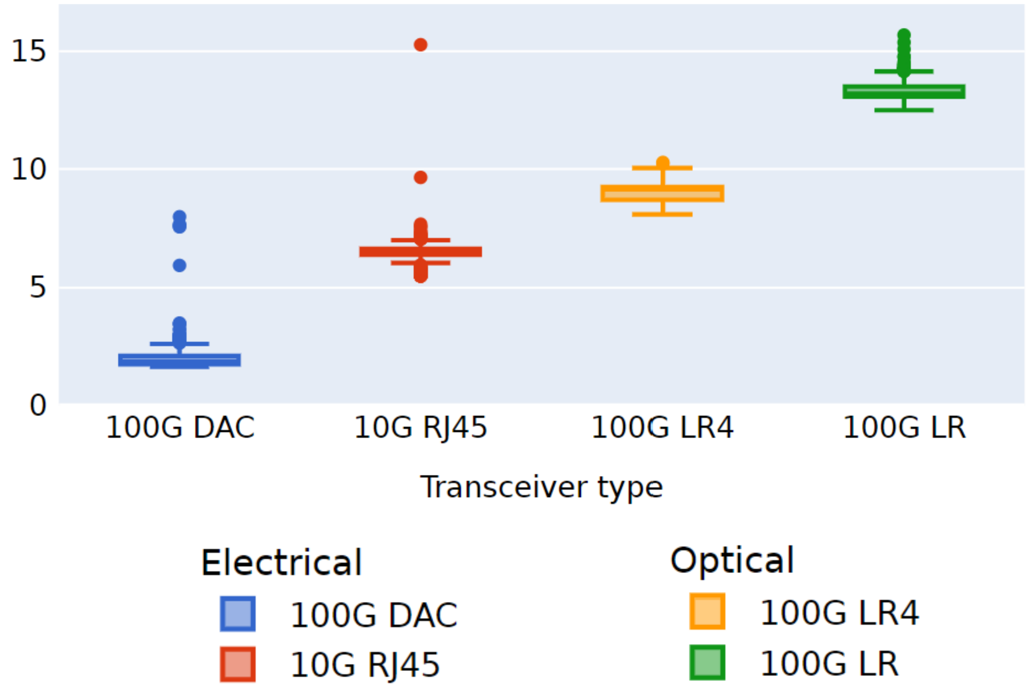

Theory

How? Buffer-and-Burst

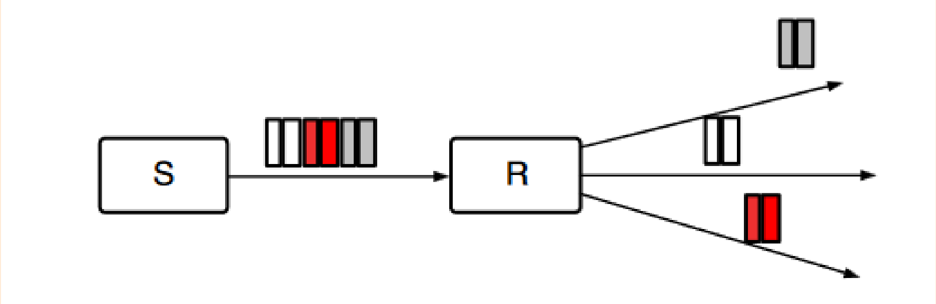

Assumes

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

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

Wake-up delay (s)

Measured on Cisco Nexus 9300

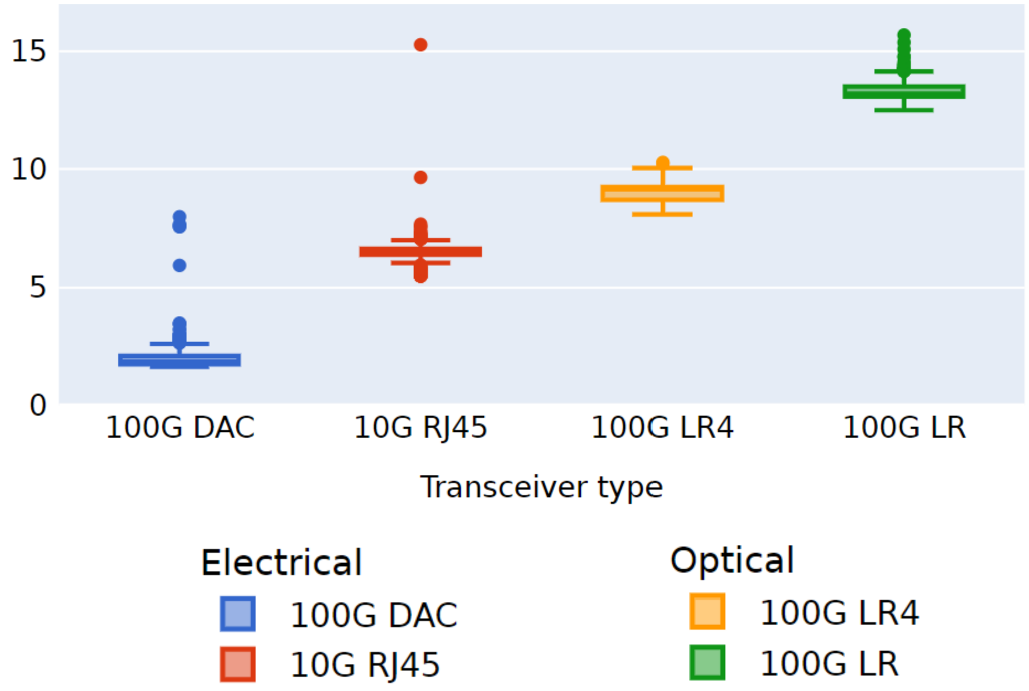

How? Buffer-and-Burst

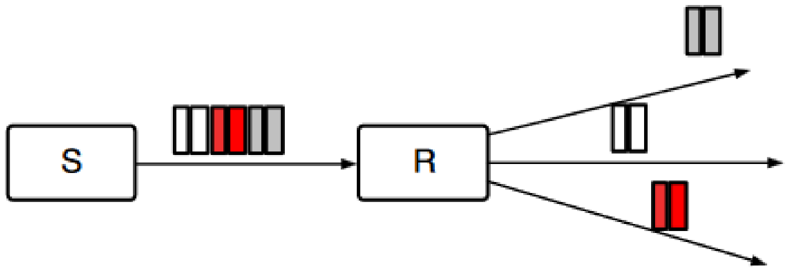

Assumes

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

We can still "sleep" at longer timescales.

Average link utilization in the Switch network [%]

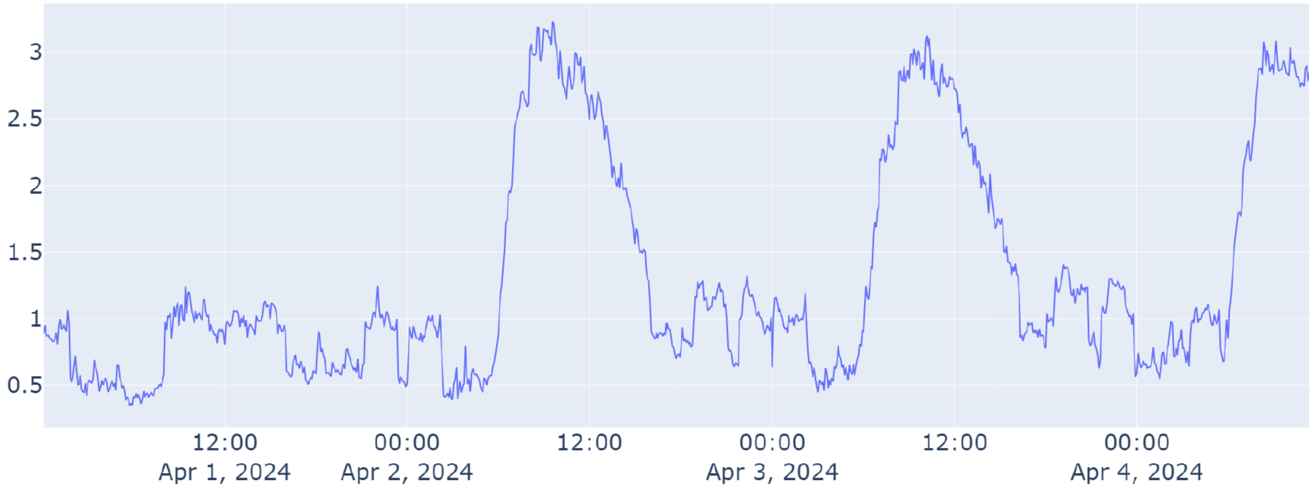

That is not very different from other routine network management tasks.

On Switch, Hypnos can turn off ~1/3 of the links.

Hypnos – Greek god of sleep

- Presented at HotCarbon'24
- Paper and presentation
- Evaluation code on GitHub
- Hardware implementation for an earlier prototype

How much energy can we really save?

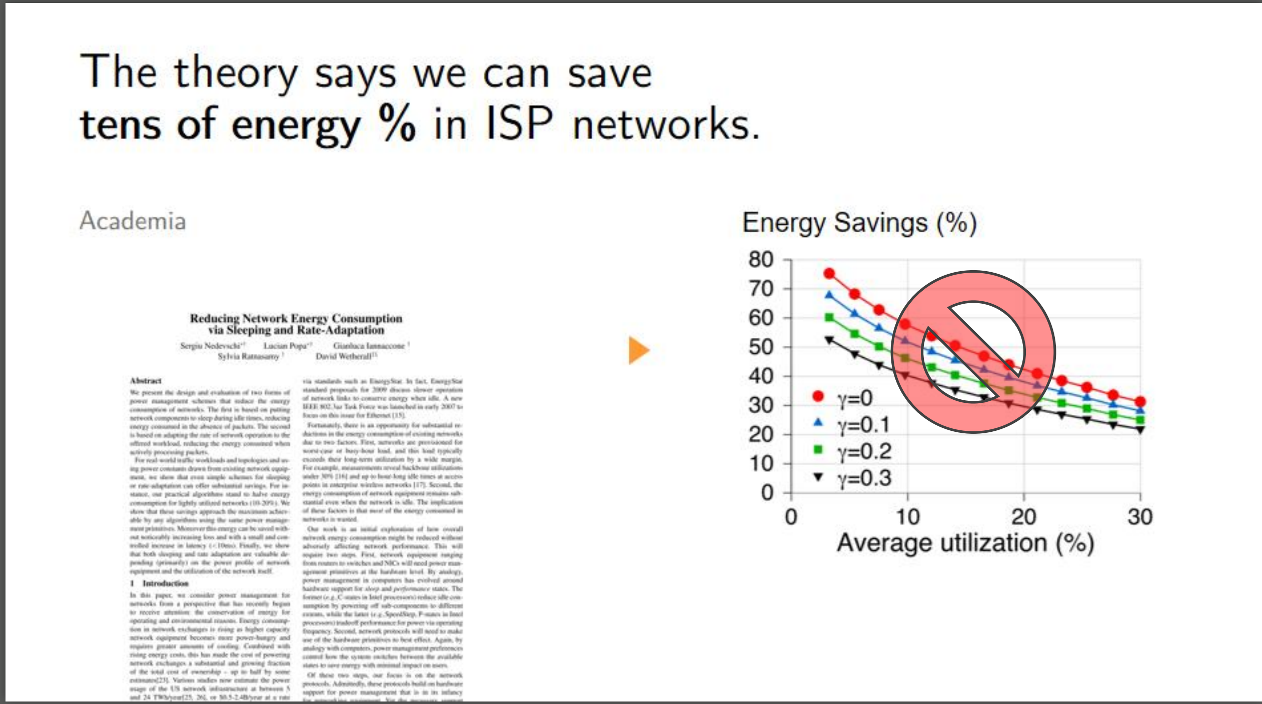

With Hypnos?

How much energy can we really save? Very little. (<1%)

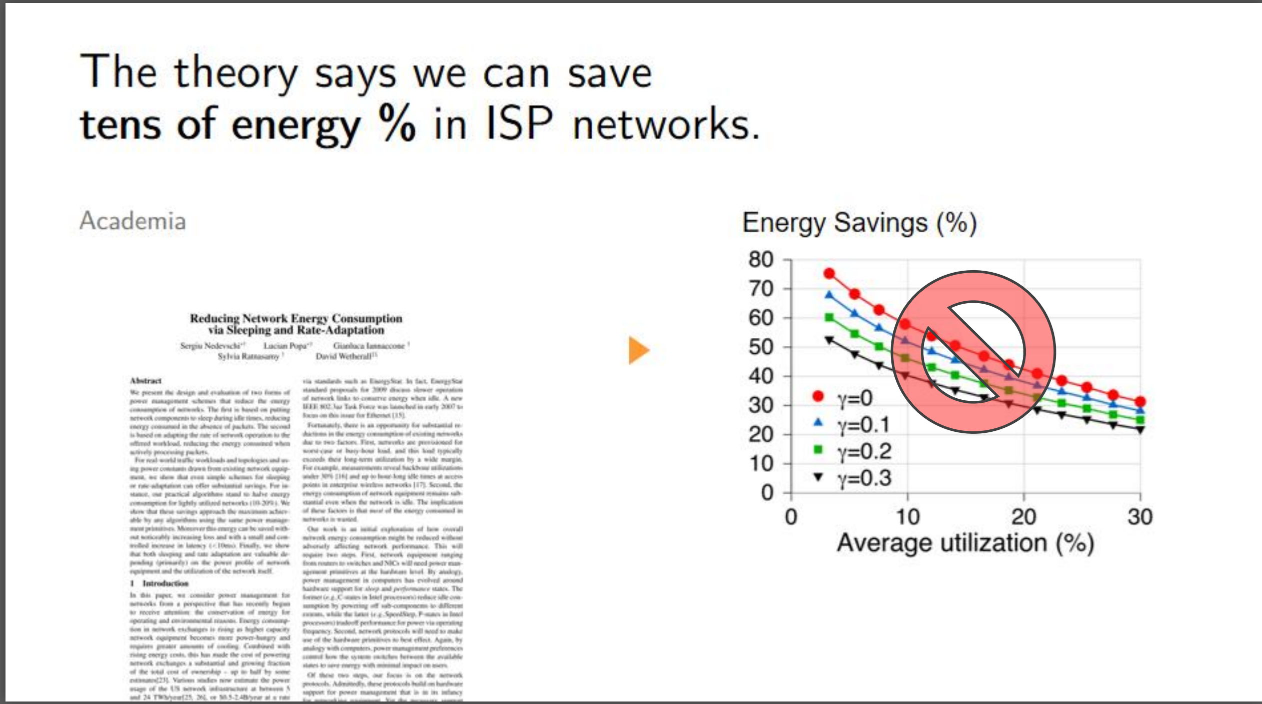

With Hypnos?

How much energy can we really save? Very little. (<1%)

Hyperloop tubes require loads of power to create vacuum, whether used or not.

Because "turned off" does **not** always mean "powered off."

EdgeCore Wedge 100BF-32X

Router

Transceiver

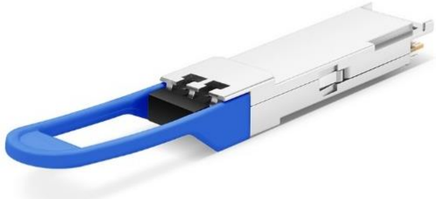

QSFP28 LR

EdgeCore Wedge 100BF-32X

Router  110 W “empty”

Transceiver 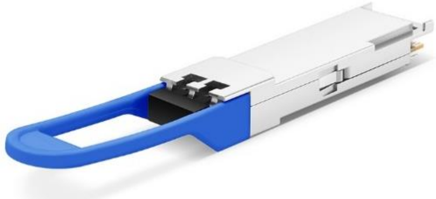 3-4 W per piece

QSFP28 LR

We are doing a terrible job at power proportionality today.

EdgeCore Wedge 100BF-32X

Router  110 W "empty"

Transceiver 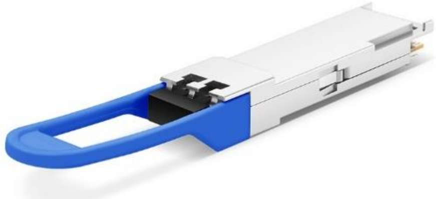 3-4 W per piece

QSFP28 LR

Power increases at plug-in
even if the port is "off"!

We are doing a terrible job at power proportionality today.

Similar problems with internal router components
that remain powered on while being configured "off."

- Thankfully, hardware vendors are starting to notice...
- ... and it is actually easy to fix!

https://community.juniper.net/blogs/nicolas-fevrier/2023/09/11/saving-power-on-acx7000-series

We are doing a terrible job at power proportionality today. It’s a bug.

Similar problems with internal router components that remain powered on while being configured “off.”

- Thankfully, hardware vendors are starting to notice...
- ... and it is actually easy to fix!

> Until this gap is addressed, a simple workaround exists: configuring the empty ports in “unused” mode.

In case you wonder
The usual config is “shutdown.”

https://community.juniper.net/blogs/nicolas-fevrier/2023/09/11/saving-power-on-acx7000-series

Don't fill the pipes!

Why we must keep the Internet empty.

Once the network is up, traffic costs almost nothing.

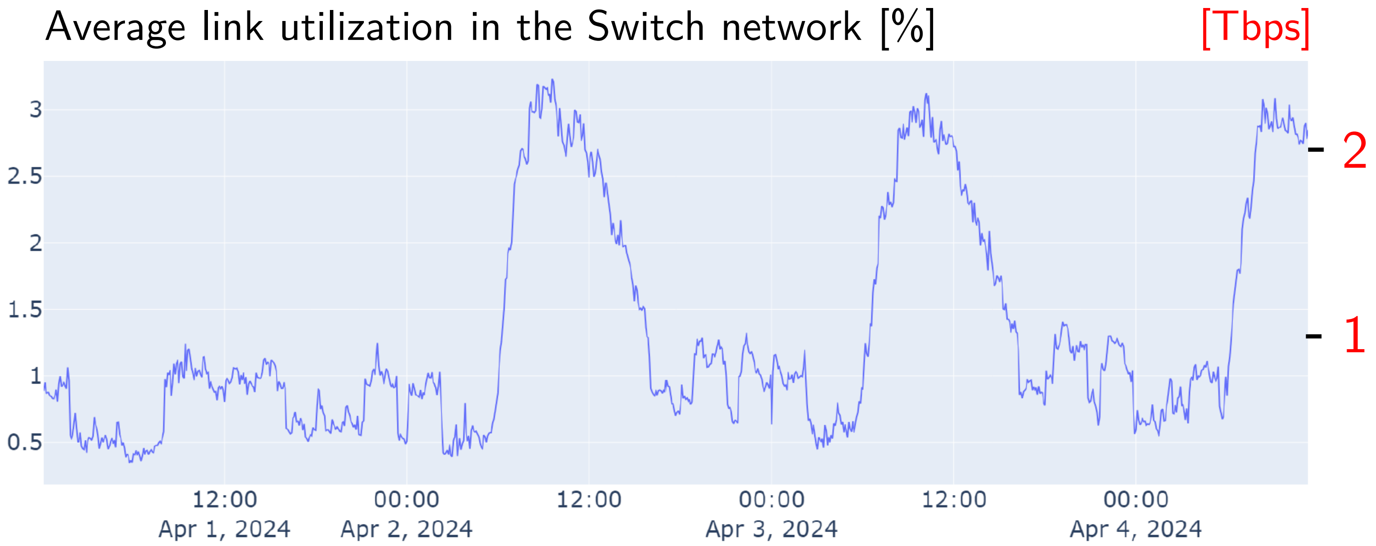

Once the network is up, traffic costs almost nothing.

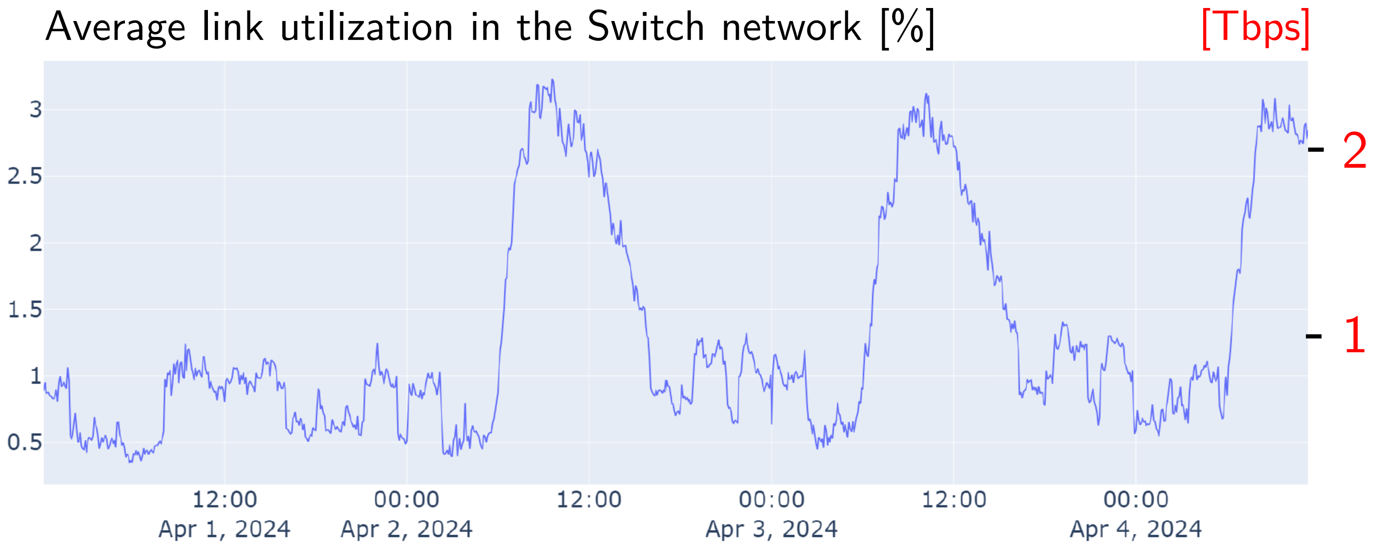

The power to send that traffic is

~6W

That's 0.02% of the total network power.

Increasing utilization improves efficiency

- Run more often at high utilization

- Take low-utilization power down

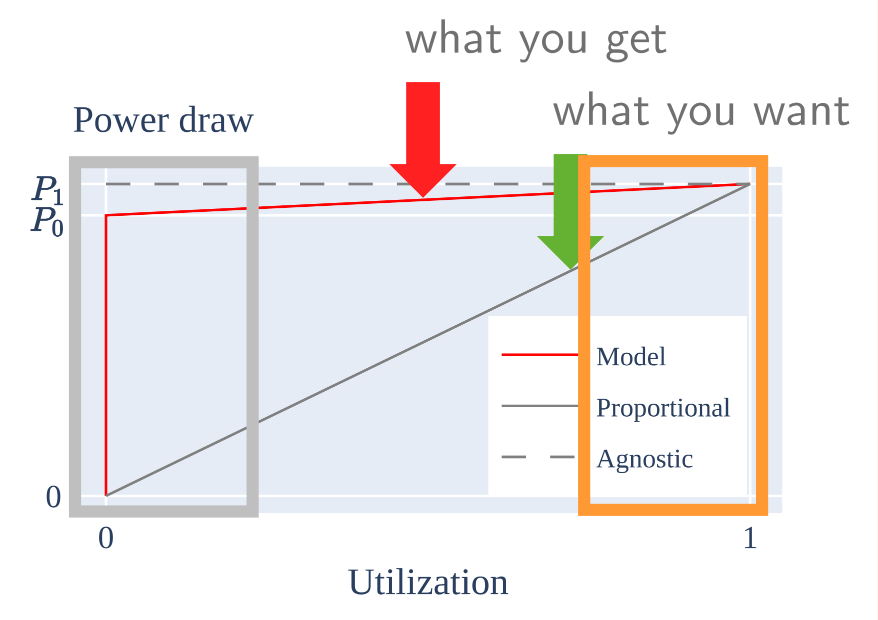

Increasing utilization improves efficiency in the short term.

- Run more often at high utilization

- Take low-utilization power down

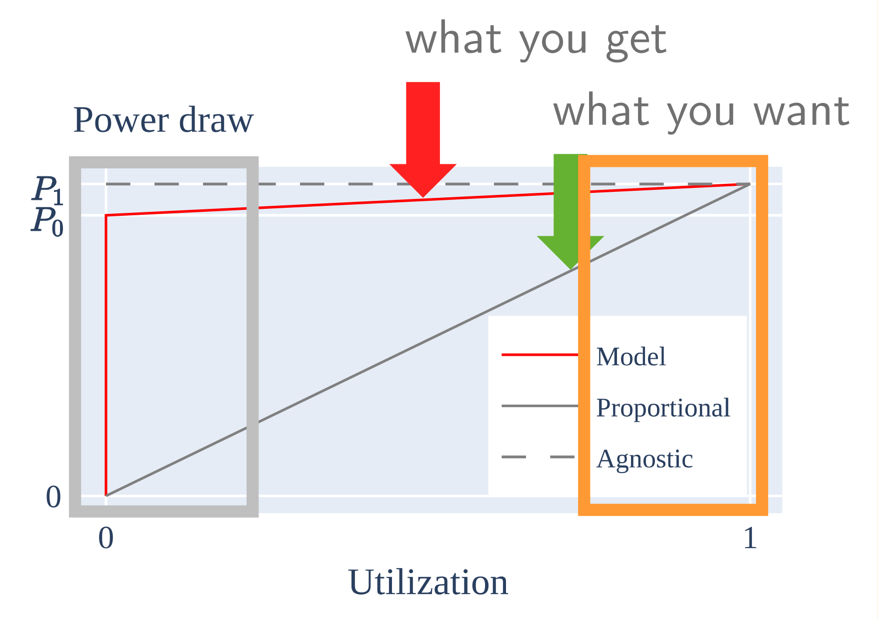

Increasing utilization improves efficiency in the short term.

- Run more often at high utilization

- Take low-utilization power down

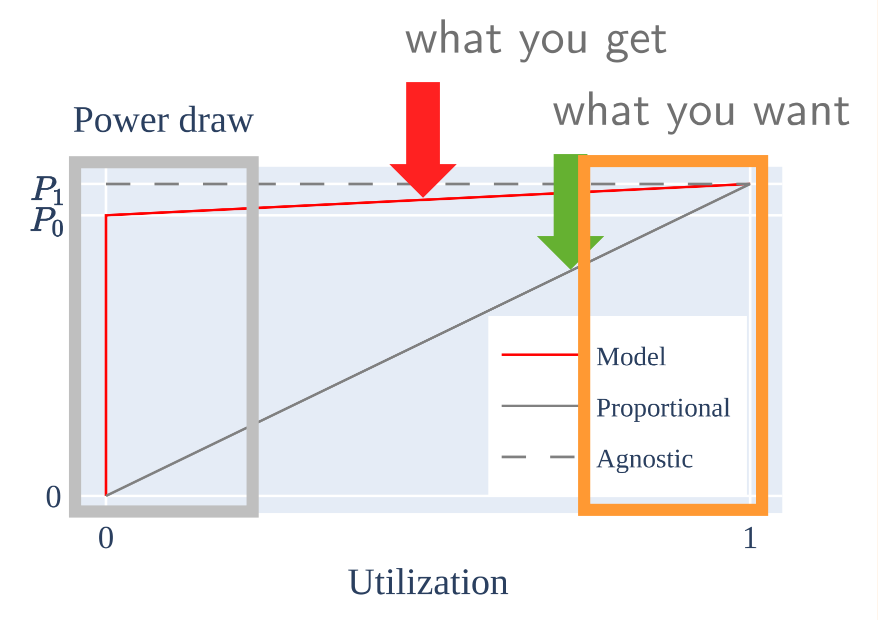

Do not forget

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

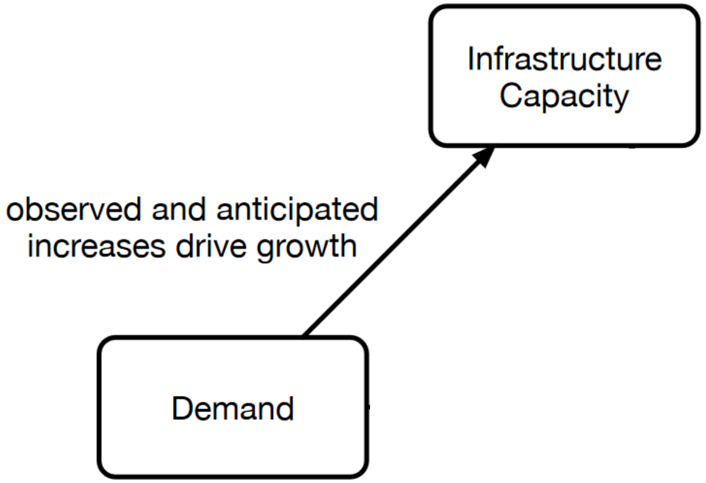

https://doi.org/10.1145/2858036.2858378

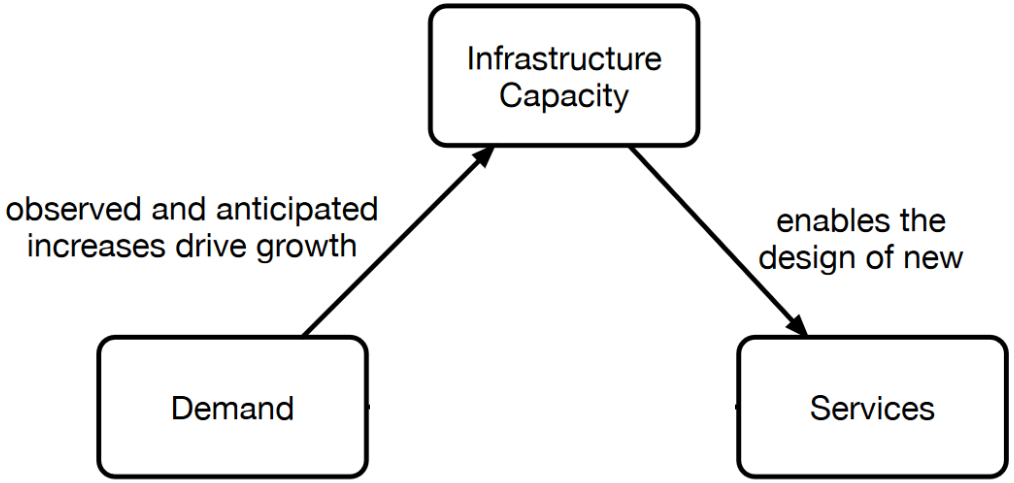

https://doi.org/10.1145/2858036.2858378

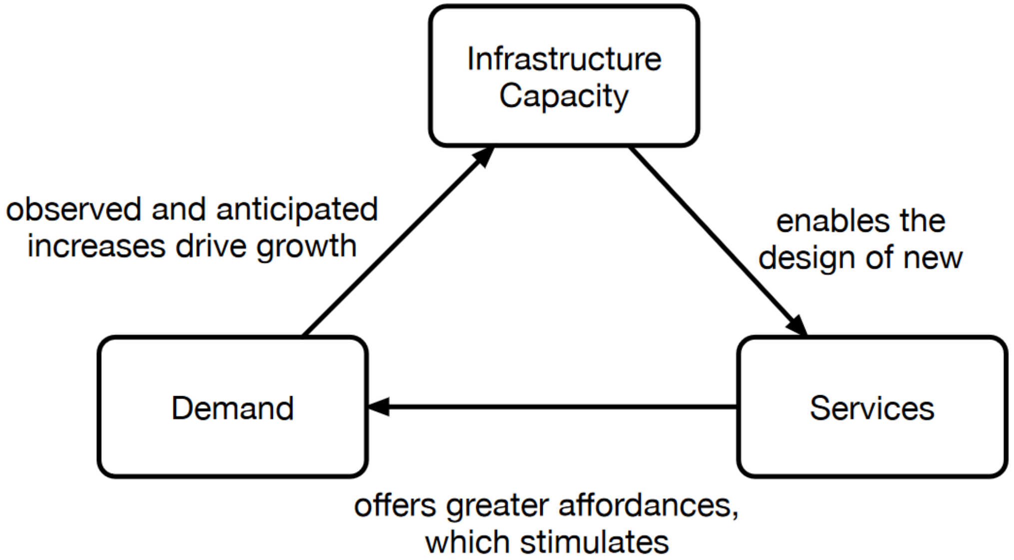

https://doi.org/10.1145/2858036.2858378

There is a feedback loop that stimulates network capacity increase

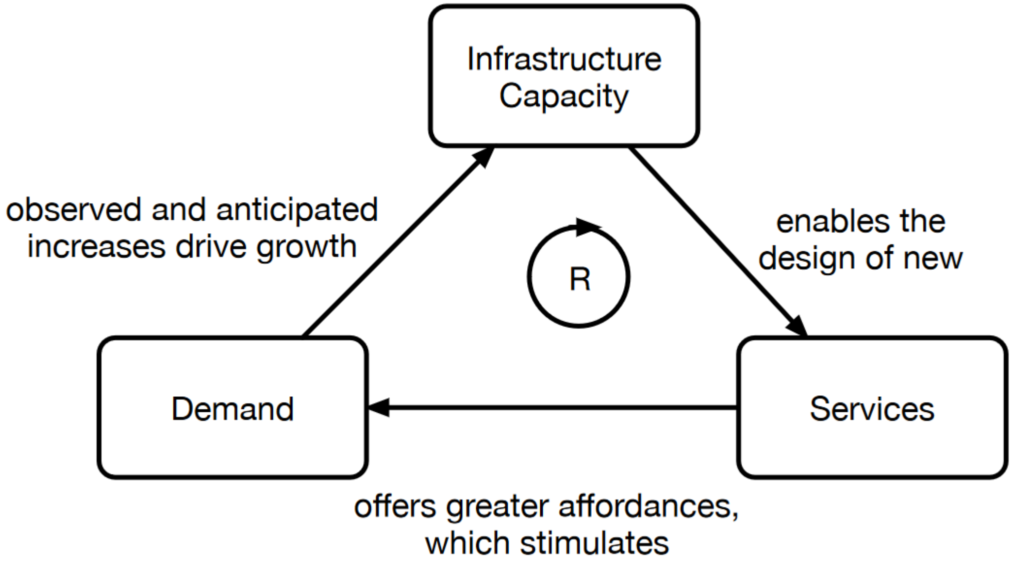

https://doi.org/10.1145/2858036.2858378

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

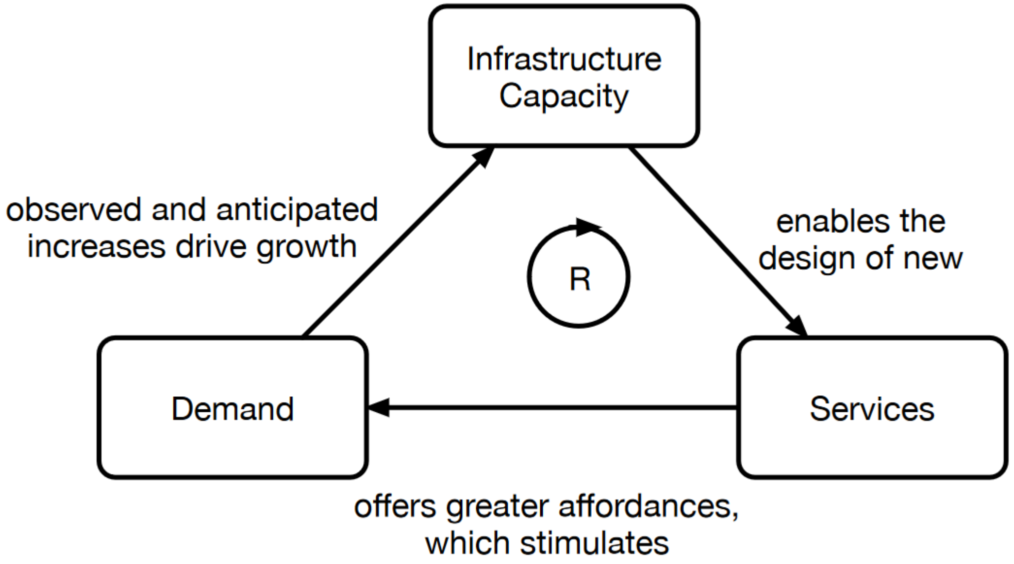

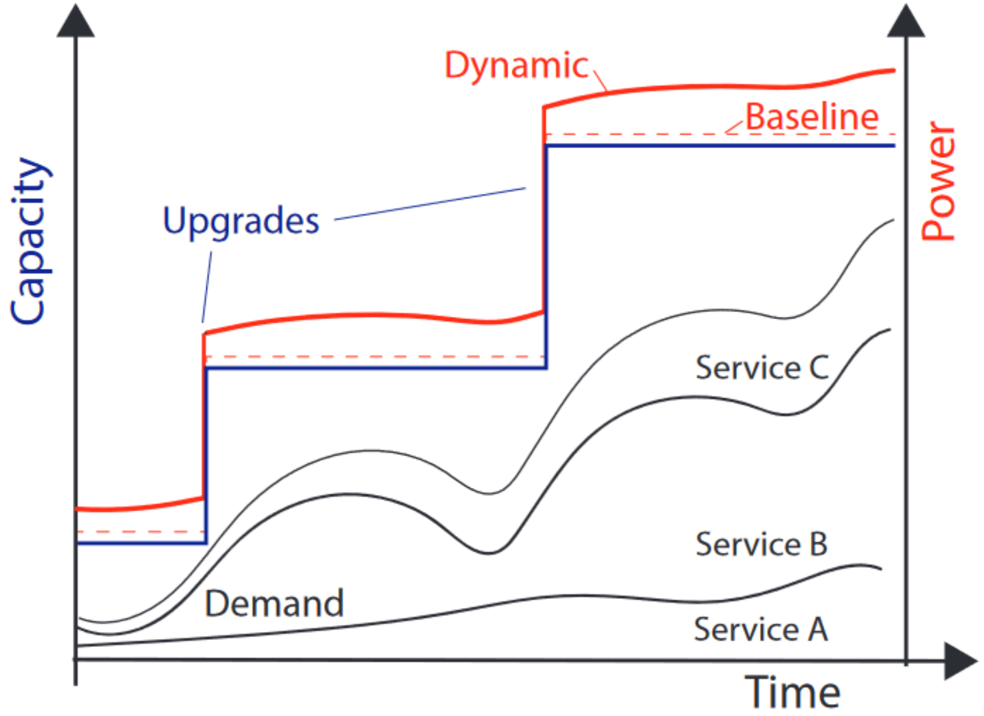

https://doi.org/10.1145/2858036.2858378

https://research-information.bris.ac.uk/en/publications/rethinking-allocation-in-high-baseload-systems-a-demand-proportio

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

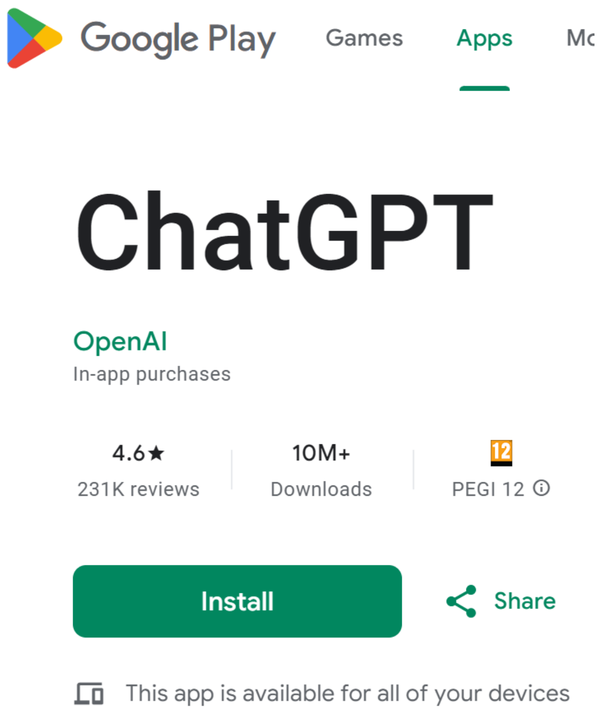

We must embrace some digital sobriety.

Everything has a cost.

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

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.

That does not have to imply worse UX!

- Software bloat can be huge.
- Applications can be redesigned without assuming permanent high-speed connectivity.

That's where GreenIO folks come in!

… and what's with the Zoo?

> What is the most surprising word that you will use in your speech?

Zoo

We have a data transparency problem for networking hardware.

- Vendors hardly provide any useful power info.
- Embedded power sensors are not trustworthy.
- There is no standard for power data reporting.

TL;DR

We need more power measurements and a systematic data collection.

We created a public database for power data: Network Power Zoo

The database contains

- Datasheet information
- PSU readings
- External measurements
- Power models

Would you share your network's data?

We work on tools to make it easy

Welcome to the networkpowerzoo.ethz.ch

1 The Internet peak traffic increases. That is driving network growth, which results in emptier networks on average.

1 The Internet peak traffic increases.
That is driving network growth, which results in emptier networks on average.

2 Networking hardware is far from power proportionality;
It is very inefficient at low utilization.

- It can and must be improved.

1 The Internet peak traffic increases.
That is driving network growth, which
results in emptier networks on average.

2 Networking hardware is far
from power proportionality;
It is very inefficient at low utilization.

- It can and must be improved.

3 Even if direct cost of traffic is minuscule,
sending more traffic has an important
systemic cost. Cf. Point 1.

- Don't send if you don't need to.
- Stay away from the peak!

The Internet is getting emptier; that's a sustainability problem that we can fix.

An Hyperloop tube segment.

Romain Jacob
jacobr@ethz.ch