Tomorrow's Internet must sleep more and grow old

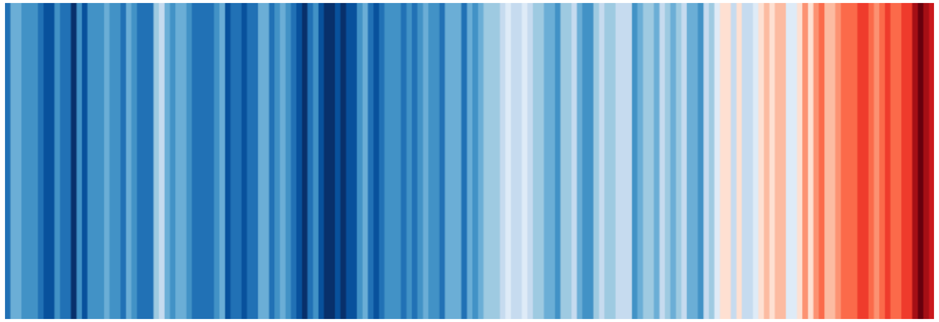

Romain Jacob Laurent Vanbever

1st HotCarbon Workshop

July 10, 2022

Let us distinguish two classes of "efficiencies."

Application efficiency	More "data" per Byte
Network efficiency	Fewer Joules per Byte

- Both matter
- Independent

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Embodied costs

Use network devices longer to
balance their manufacturing cost

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Embodied costs

Use network devices longer to
balance their manufacturing cost

- Grow old

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Embodied costs

Use network devices longer to
balance their manufacturing cost

- Grow old

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.

SIGCOMM'03, August 25–29, 2003, Karlsruhe, Germany.
Copyright 2003 ACM 1-58113-735-4/03/0008 ...$5.00.

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.

Gupta and Singh

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

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

[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.

Gupta and Singh

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.

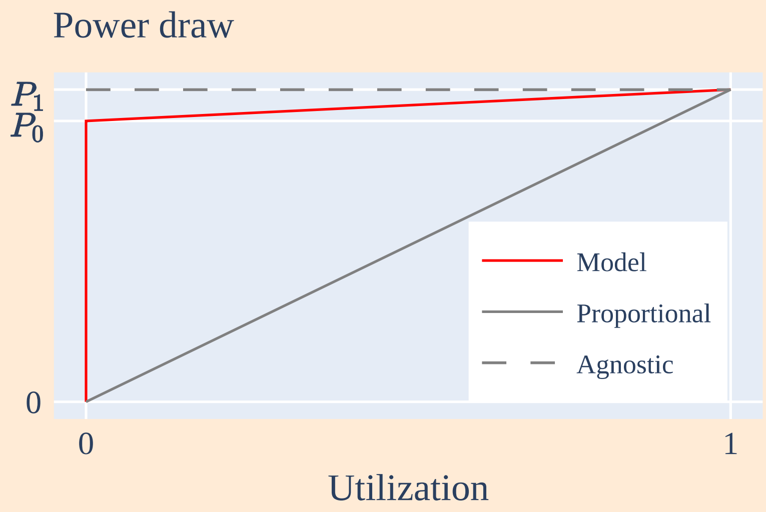

1. Network devices are always "on."

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

3. Network devices are under-utilized.

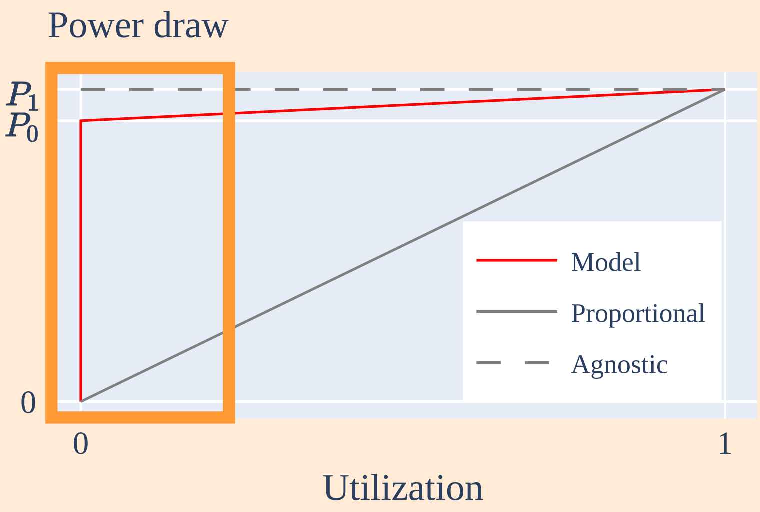

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

1. Network devices are always "on."

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

3. Network devices are under-utilized.

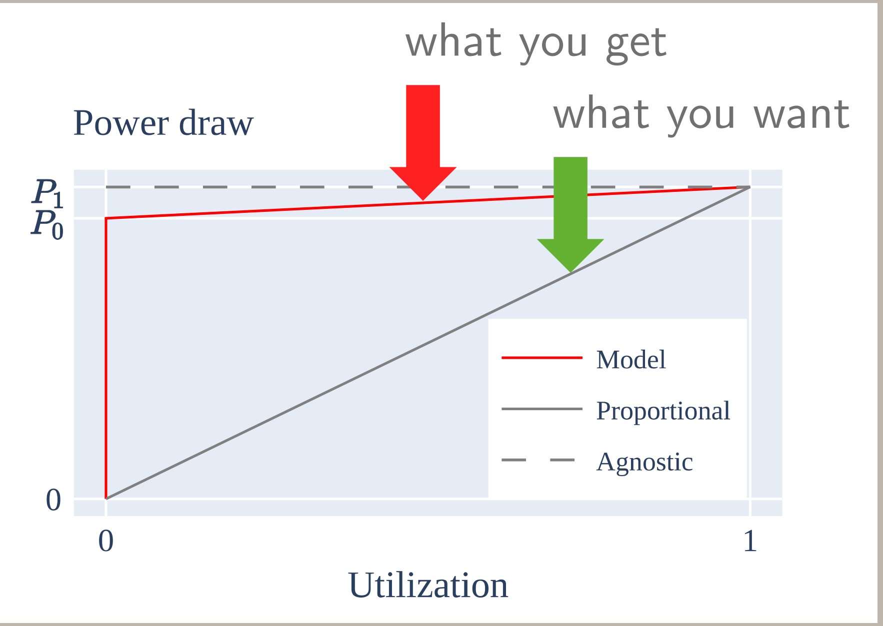

ISP overprovision networks to support

- Peak traffic
- Fault tolerance

Networking researchers investigated powering devices off to save energy.

One example

ElasticTree NSDI 2010

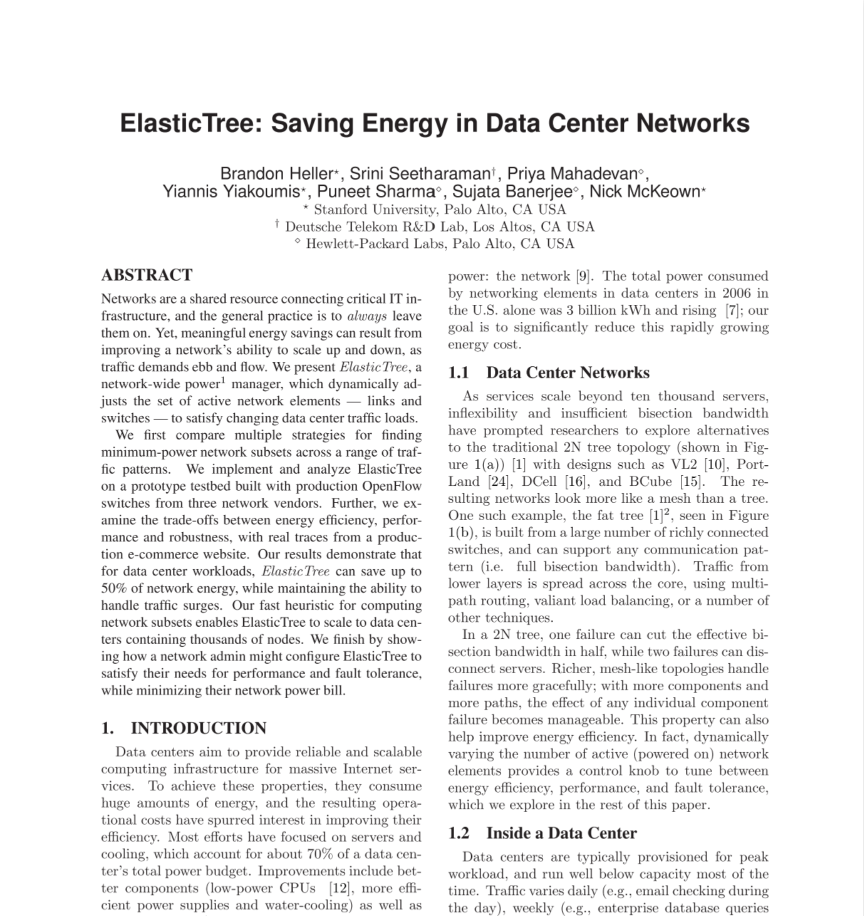

ElasticTree: Saving Energy in Data Center Networks

Brandon Heller*, Srini Seetharaman†, Priya Mahadevan◦,
Yiannis Yiakoumis*, Puneet Sharma◦, Sujata Banerjee◦, Nick McKeown*
* Stanford University, Palo Alto, CA USA
† Deutsche Telekom R&D Lab, Los Altos, CA USA
◦ Hewlett-Packard Labs, Palo Alto, CA USA

ABSTRACT

Networks are a shared resource connecting critical IT infrastructure, and the general practice is to *always* leave them on. Yet, meaningful energy savings can result from improving a network's ability to scale up and down, as traffic demands ebb and flow. We present *ElasticTree*, a network-wide power[1] manager, which dynamically adjusts the set of active network elements — links and switches — to satisfy changing data center traffic loads.

We first compare multiple strategies for finding minimum-power network subsets across a range of traffic patterns. We implement and analyze ElasticTree on a prototype testbed built with production OpenFlow switches from three network vendors. Further, we examine the trade-offs between energy efficiency, performance and robustness, with real traces from a production e-commerce website. Our results demonstrate that for data center workloads, *ElasticTree* can save up to 50% of network energy, while maintaining the ability to handle traffic surges. Our fast heuristic for computing network subsets enables ElasticTree to scale to data centers containing thousands of nodes. We finish by showing how a network admin might configure ElasticTree to satisfy their needs for performance and fault tolerance, while minimizing their network power bill.

1. INTRODUCTION

Data centers aim to provide reliable and scalable computing infrastructure for massive Internet services. To achieve these properties, they consume huge amounts of energy, and the resulting operational costs have spurred interest in improving their efficiency. Most efforts have focused on servers and cooling, which account for about 70% of a data center's total power budget. Improvements include better components (low-power CPUs [12], more efficient power supplies and water-cooling) as well as power: the network [9]. The total power consumed by networking elements in data centers in 2006 in the U.S. alone was 3 billion kWh and rising [7]; our goal is to significantly reduce this rapidly growing energy cost.

1.1 Data Center Networks

As services scale beyond ten thousand servers, inflexibility and insufficient bisection bandwidth have prompted researchers to explore alternatives to the traditional 2N tree topology (shown in Figure 1(a)) [1] with designs such as VL2 [10], PortLand [24], DCell [16], and BCube [15]. The resulting networks look more like a mesh than a tree. One such example, the fat tree [1][2], seen in Figure 1(b), is built from a large number of richly connected switches, and can support any communication pattern (i.e. full bisection bandwidth). Traffic from lower layers is spread across the core, using multipath routing, valiant load balancing, or a number of other techniques.

In a 2N tree, one failure can cut the effective bisection bandwidth in half, while two failures can disconnect servers. Richer, mesh-like topologies handle failures more gracefully; with more components and more paths, the effect of any individual component failure becomes manageable. This property can also help improve energy efficiency. In fact, dynamically varying the number of active (powered on) network elements provides a control knob to tune between energy efficiency, performance, and fault tolerance, which we explore in the rest of this paper.

1.2 Inside a Data Center

Data centers are typically provisioned for peak workload, and run well below capacity most of the time. Traffic varies daily (e.g., email checking during the day), weekly (e.g., enterprise database queries

Proofs of concept show the potential saving opportunities

Networking researchers investigated powering devices off to save energy.

One example

ElasticTree NSDI 2010

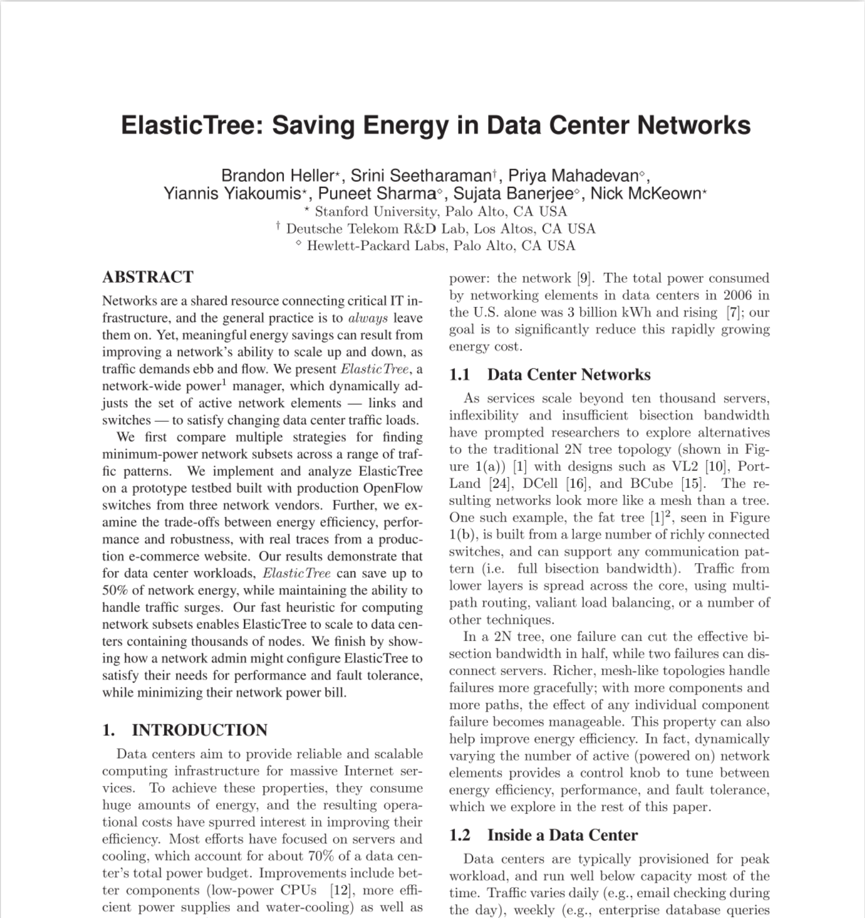

ElasticTree: Saving Energy in Data Center Networks

Brandon Heller*, Srini Seetharaman†, Priya Mahadevan◦,
Yiannis Yiakoumis*, Puneet Sharma◦, Sujata Banerjee◦, Nick McKeown*
* Stanford University, Palo Alto, CA USA
† Deutsche Telekom R&D Lab, Los Altos, CA USA
◦ Hewlett-Packard Labs, Palo Alto, CA USA

ABSTRACT

Networks are a shared resource connecting critical IT infrastructure, and the general practice is to *always* leave them on. Yet, meaningful energy savings can result from improving a network's ability to scale up and down, as traffic demands ebb and flow. We present *ElasticTree*, a network-wide power[1] manager, which dynamically adjusts the set of active network elements — links and switches — to satisfy changing data center traffic loads.

We first compare multiple strategies for finding minimum-power network subsets across a range of traffic patterns. We implement and analyze ElasticTree on a prototype testbed built with production OpenFlow switches from three network vendors. Further, we examine the trade-offs between energy efficiency, performance and robustness, with real traces from a production e-commerce website. Our results demonstrate that for data center workloads, *ElasticTree* can save up to 50% of network energy, while maintaining the ability to handle traffic surges. Our fast heuristic for computing network subsets enables ElasticTree to scale to data centers containing thousands of nodes. We finish by showing how a network admin might configure ElasticTree to satisfy their needs for performance and fault tolerance, while minimizing their network power bill.

1. INTRODUCTION

Data centers aim to provide reliable and scalable computing infrastructure for massive Internet services. To achieve these properties, they consume huge amounts of energy, and the resulting operational costs have spurred interest in improving their efficiency. Most efforts have focused on servers and cooling, which account for about 70% of a data center's total power budget. Improvements include better components (low-power CPUs [12], more efficient power supplies and water-cooling) as well as power: the network [9]. The total power consumed by networking elements in data centers in 2006 in the U.S. alone was 3 billion kWh and rising [7]; our goal is to significantly reduce this rapidly growing energy cost.

1.1 Data Center Networks

As services scale beyond ten thousand servers, inflexibility and insufficient bisection bandwidth have prompted researchers to explore alternatives to the traditional 2N tree topology (shown in Figure 1(a)) [1] with designs such as VL2 [10], PortLand [24], DCell [16], and BCube [15]. The resulting networks look more like a mesh than a tree. One such example, the fat tree [1][2], seen in Figure 1(b), is built from a large number of richly connected switches, and can support any communication pattern (i.e. full bisection bandwidth). Traffic from lower layers is spread across the core, using multipath routing, valiant load balancing, or a number of other techniques.

In a 2N tree, one failure can cut the effective bisection bandwidth in half, while two failures can disconnect servers. Richer, mesh-like topologies handle failures more gracefully; with more components and more paths, the effect of any individual component failure becomes manageable. This property can also help improve energy efficiency. In fact, dynamically varying the number of active (powered on) network elements provides a control knob to tune between energy efficiency, performance, and fault tolerance, which we explore in the rest of this paper.

1.2 Inside a Data Center

Data centers are typically provisioned for peak workload, and run well below capacity most of the time. Traffic varies daily (e.g., email checking during the day), weekly (e.g., enterprise database queries

Proofs of concept show the potential saving opportunities

But...

- Convergence issues
  when turning on/off devices
- Management issues
  with complexity and scalability
- Start-up delay issues
  between 30s and 3min

Other fields deal with energy shortage forever!

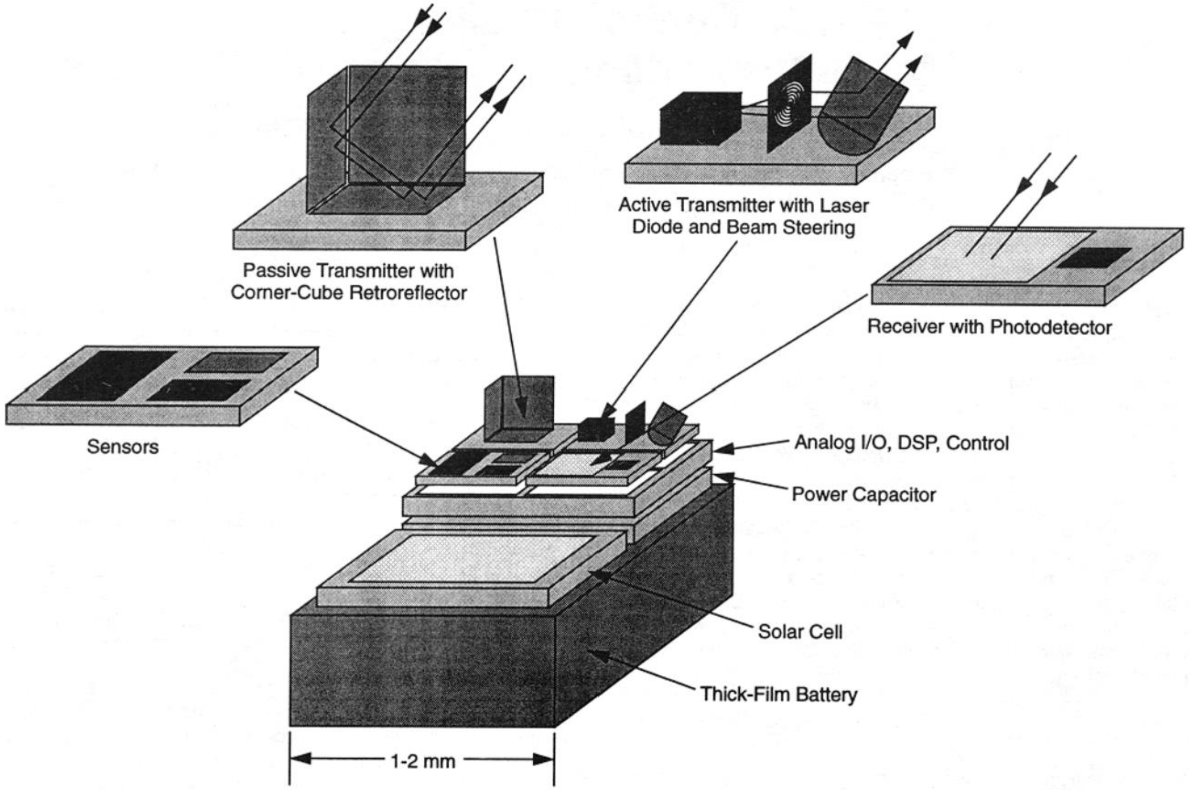

The "smart dust" – 1999, Berkeley

Other fields deal with energy shortage forever!

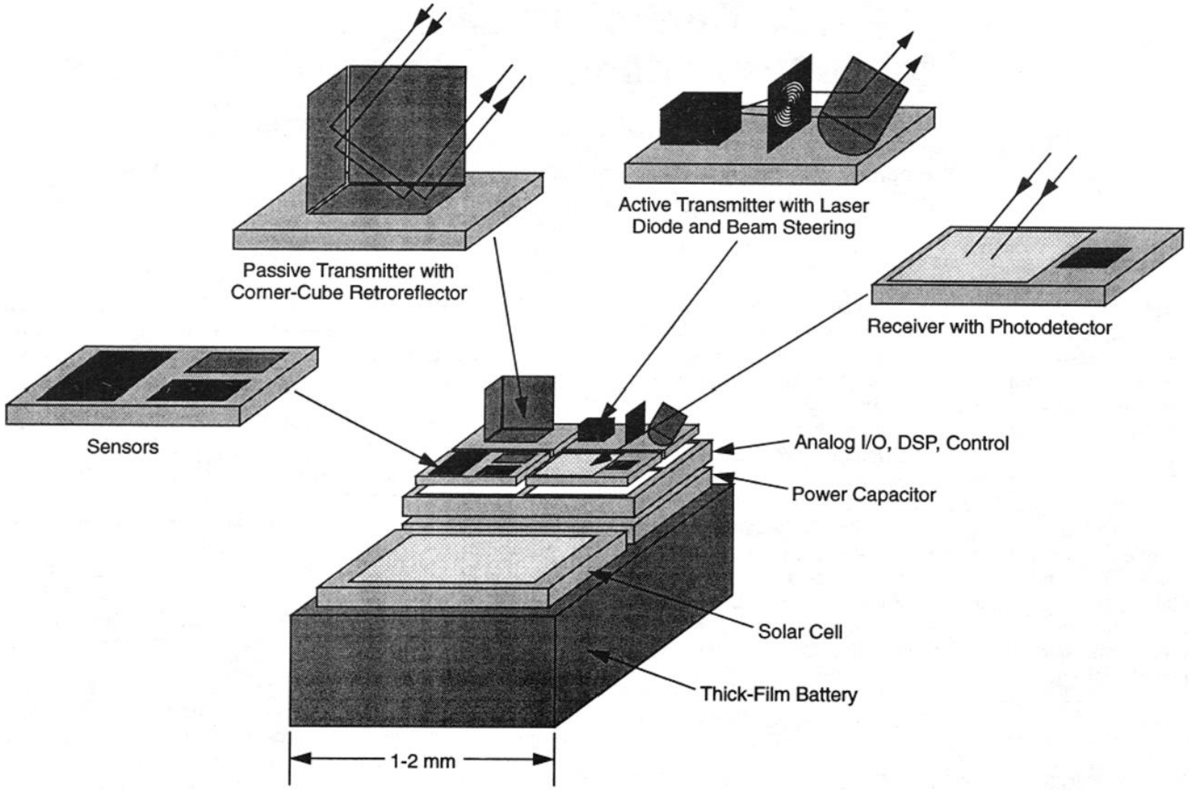

The "smart dust" – 1999, Berkeley

The Internet of Things – 2022

There are 2+ decades of research dedicated to exchanging information with minimum energy!

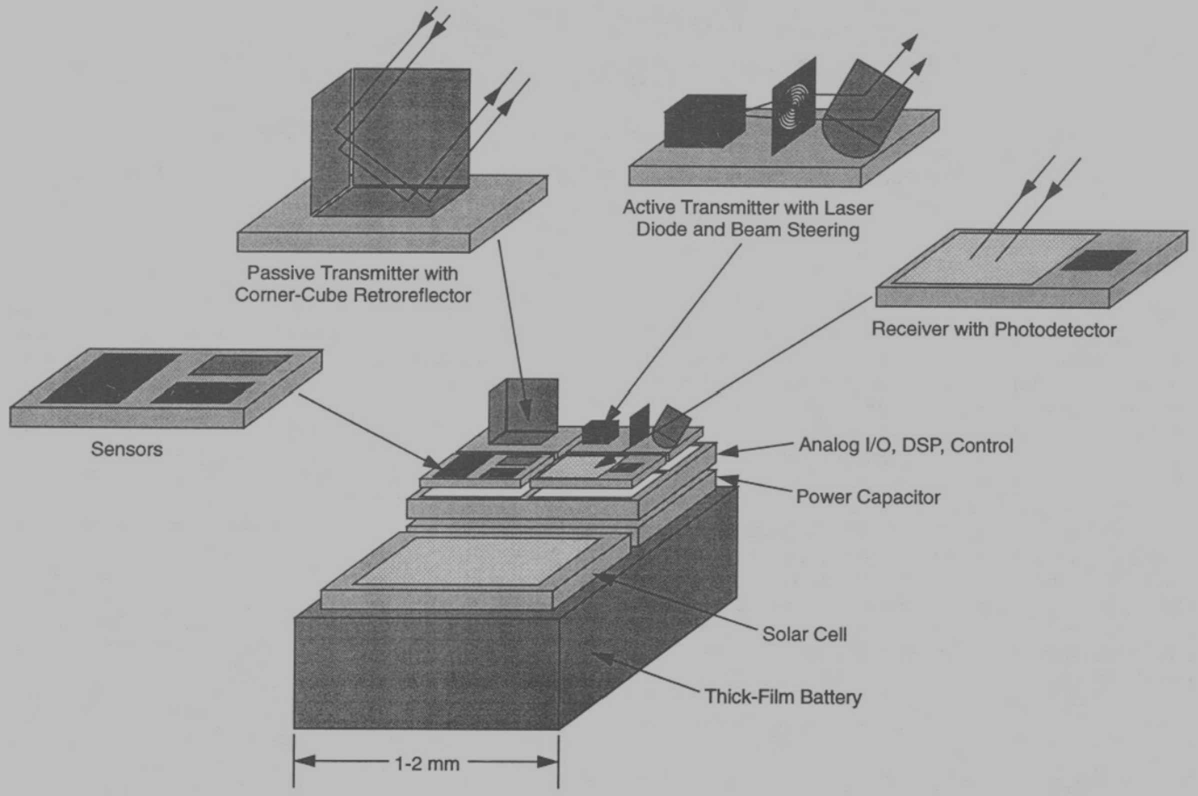

The "smart dust" – 1999, Berkeley

The Internet of Things – 2022

There are 2+ decades of research dedicated to exchanging information with minimum energy!

In low-power wireless networking energy efficiency boils down to turning the radios off as long as possible.

▶ State-of-the-art is batteryless networking!

Different routing strategies

- Centralized
- Distributed
- Hybrid / Mixed

> Networked Embedded Systems researchers have been designing and deploying networks that successfully route packets using "off-by-default" principles.

1. Can computer networking find inspiration from the low-power networking literature?

1. Can computer networking find inspiration from the low-power networking literature? Yes

1. Can computer networking find inspiration from the low-power networking literature? Yes

2. If we redesign routing with energy efficiency as primary objective, what would it look like?

1. Can computer networking find inspiration from the low-power networking literature? — Yes

2. If we redesign routing with energy efficiency as primary objective, what would it look like? — Not clear — but likely with a more decentralized management plane

1. Can computer networking find inspiration from the low-power networking literature? — Yes

2. If we redesign routing with energy efficiency as primary objective, what would it look like? — Not clear — but likely with a more decentralized management plane

3. How much energy savings can be obtained without degrading the QoS?

1.	Can computer networking find inspiration from the low-power networking literature?	Yes	
2.	If we redesign routing with energy efficiency as primary objective, what would it look like?	Not clear	but likely with a more decentralized management plane
3.	How much energy savings can be obtained without degrading the QoS?	Needs research	

1. Can computer networking find inspiration from the low-power networking literature?

2. If we redesign routing with energy efficiency as primary objective, what would it look like?

3. How much energy savings can be obtained without degrading the QoS?

Is it worth it?

We could save 50% energy
in today's ISP networks.

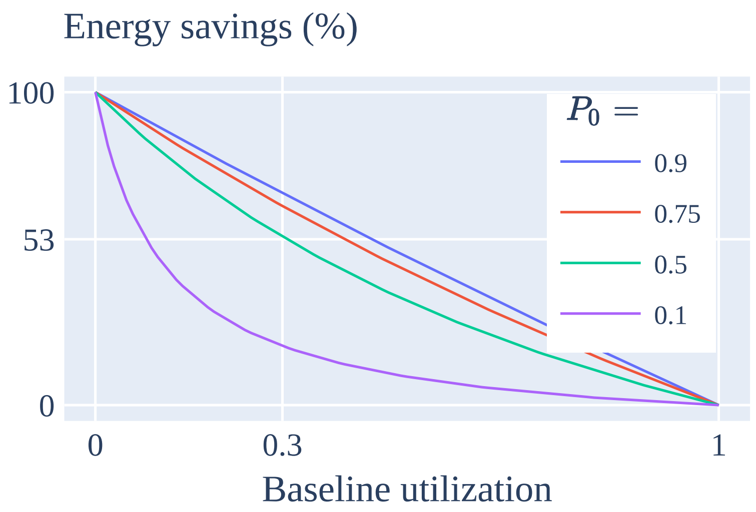

Possible savings are

$\geq$ 50% given

- $P_0/P_1 \geq 0.5$
- Utilization $\leq$ 30%

[in an ideal world]

To harness these benefits, we must speed up the devices "start-up time."

From "low-power"
to "ready-to-forward"

Routing protocols × Networking software × Networking hardware

Application efficiency

Network efficiency

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Embodied costs

Use network devices longer to
balance their manufacturing cost

- Grow old

The embodied costs of ICT devices is surprisingly high compared to operational ones.

52%	for laptops
72%	for smartphones

Producing these devices has a larger carbon footprint than using them... !

The embodied costs of ICT devices is surprisingly high compared to operational ones.

52%	for laptops
72%	for smartphones
10-20%	for servers

Producing these devices has a larger carbon footprint than using them... !

Frugal Computing.
Wim Vanderbauwhede, 06/2021, Online.

The embodied costs of ICT devices is surprisingly high compared to operational ones.

52%	for laptops	Producing these devices has a larger carbon footprint than using them... !
72%	for smartphones	
10-20%	for servers	
unclear	for networking devices	Likely in the same ballpark

Reducing the embodied cost is simple:
Use the 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.

Today Refresh rates are around 3-5 years only.

- Easy to extend

Okay, but

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

Today

Refresh rates are 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

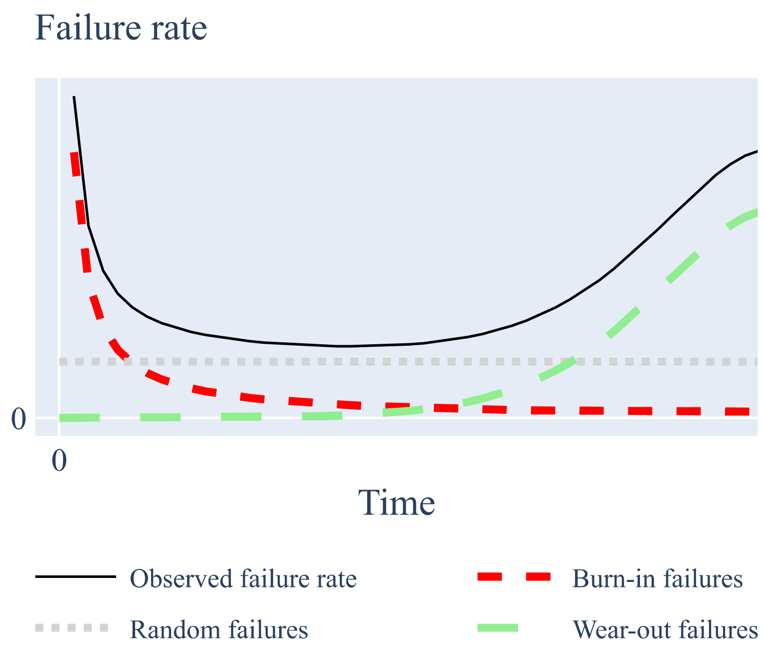

Surprising truth about network hardware failures. CXTEC, 03/2022, Online.

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 network vendors ensure a 5-years window between end-of-sale and end-of-support.
- Some companies specialize in refurbishing network hardware with extensive warranties — sometimes even unlimited!

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 ?

Okay, but

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

See paper

Not necessarily.

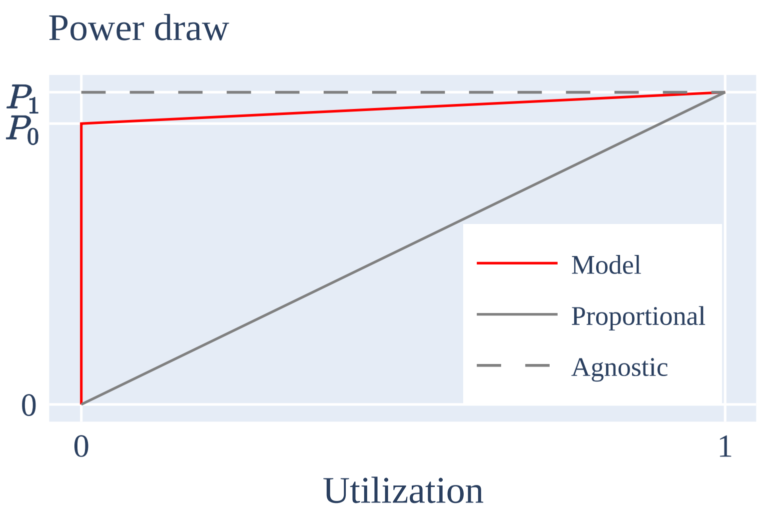

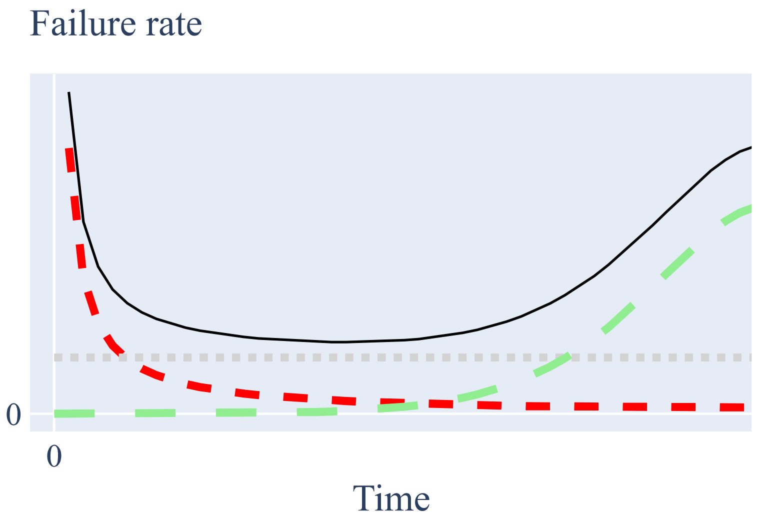

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Embodied costs

Use network devices longer to
balance their manufacturing cost

- Grow old

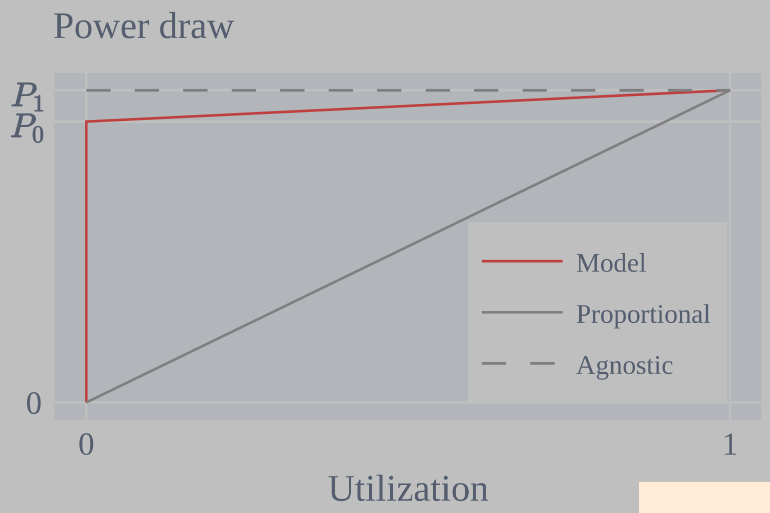

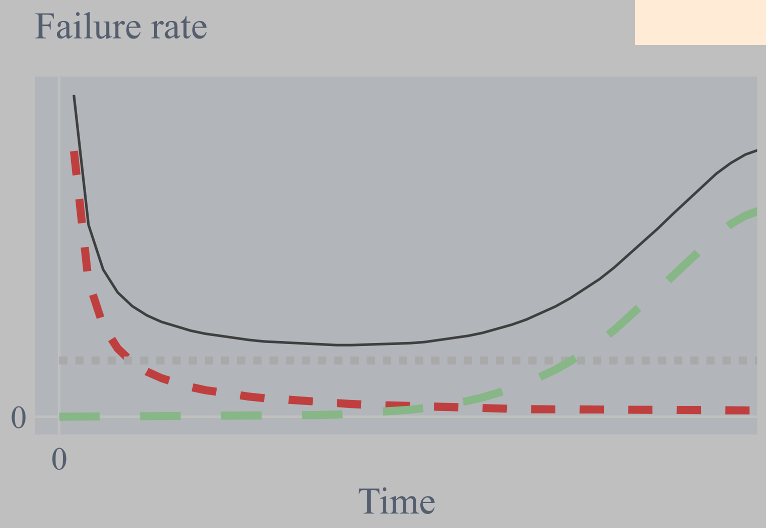

Operational costs

Turn network devices off
to reduce the average power draw

- Sleep more

Help welcome!

Embodied costs

Use network devices longer to
balance their manufacturing cost

- Grow old

Tomorrow's Internet must sleep more and grow old

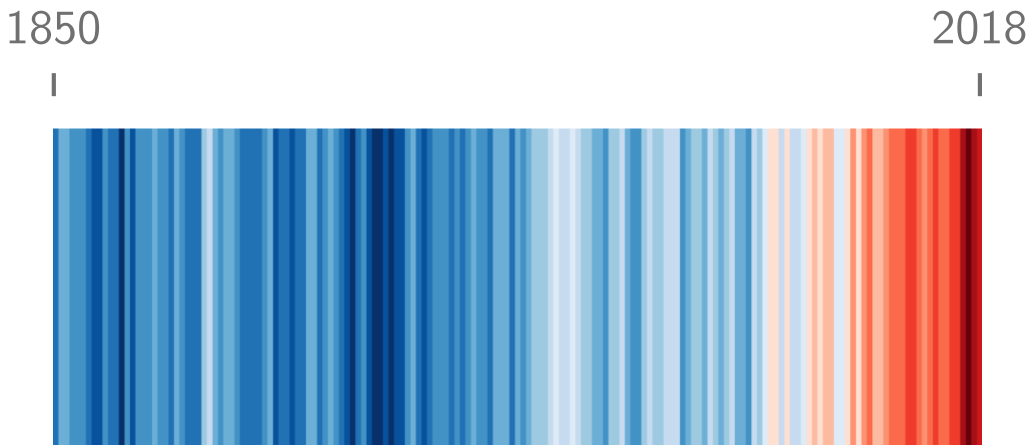

Climate stripes. Ed Hawkins, 2018
portrays the increase of average global temperature

Romain Jacob
ETH Zürich
@RJacobPartner
jacobr@ethz.ch

Laurent Vanbever
ETH Zürich
@lvanbever
lvanbever@ethz.ch