Energy-Efficiency in Communication

Jeremias Hartz

Arbeitsbereich Wissenschaftliches Rechnen
Fachbereich Informatik
Fakultät für Mathematik, Informatik und Naturwissenschaften
Universität Hamburg

2014-11-26

Outline

Energy

- Energy measured in Joules
- 3,600 kJ = 1 kWh $\rightarrow$ electric energy
- Power measured in Watts
- Average energy consumption in 2013 of
  - My Laptop*: $\underline{206 kWh}$
  - German household per capita: $\underline{1,770 kWh}$
  - All U.S. data centers: $\underline{91,000,000,000 kWh}$
    $\Rightarrow$ 63.4% of the energy consumption of all German homes that year

*Thinkpad X; 7h active, 12h idle, 4h sleep, 1h off per day

Communication

- Energy needed for the communication process vs communication speed
- Many connected components $\Rightarrow$ difficult to measure

Figure: 1hour videoconference between Switzerland and Japan: 3.6kWh!?

Cables

Optical Fiber vs Copper

- Fiber:
  - non-flammable
  - further distance and more bandwidth
  - cable itself cheaper than copper
  - generally more expensive because of the installation and electronics
  - overall cheaper compared to copper at high performance
    Example: up to 86% less energy in 10Gbit/s ports

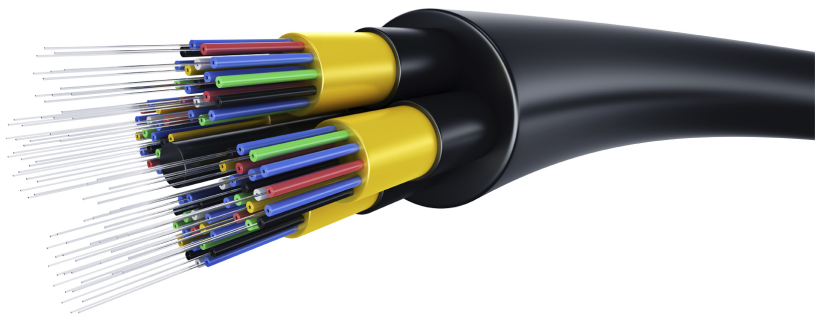

Figure: Optical Fiber Cable, Source: www.ecmag.com

Energy-Efficiency of Ports

100Mbit/s 24-port Switch by TP-Link TL-SF1024D:

- 3.53W / 24 Ports $\rightarrow$ 0.147W per Port

10Gbit/s 8-port 10-Gigabit ProSAFE® Plus Switch:

- 58.8W / 8 Ports $\rightarrow$ 7.35W per Port

$0.147W * 100$ (100Mbit to 10Gbit) $= 14.7W \rightarrow$
Highspeed Port more efficient

Figure: TL-SF1024D (TP-Link) vs XS708E (Netgear)

Common Protocols

- Ethernet for local area networks (LAN) (computer to computer)
  - Up to 40Gbit/s
- IP for internet communication in general
- TCP for connection oriented data transport over ip (safe $\rightarrow$ handshake)
- UDP for connectionless fast data transport over ip (unsafe)
- HTTP(S) super efficient for websites (uses 1 port)
- (S)FTP for file transmission (uses 2 ports)

IB, FC, SCSI

- Infiniband and Fiber Channel for storage area networks (SAN)
- Small Computer System Interface (SCSI)
  - Connection standard between computers and storage devices
- Mostly used in data centers through fiber optic cables
- Transfer as stream not in packages $\rightarrow$ faster
- iSCSI for connection from outside the SAN

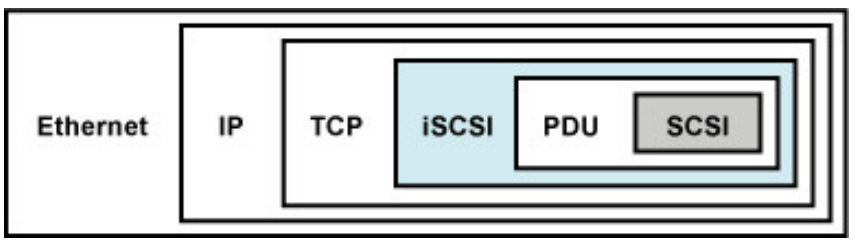

Figure: iSCSI, Source: www.elektronik-kompendium.de

Networks Energy-Efficiency Differences

- LAN Components – slower, cheaper, less energy
- WAN Components – faster, expensive, more energy

\+ Single devices in LAN often more energy-efficient for their tasks and throughput

– Their energy-efficiency not directly scalable; disproportionate increase of performance vs. energy (cooling, etc.)

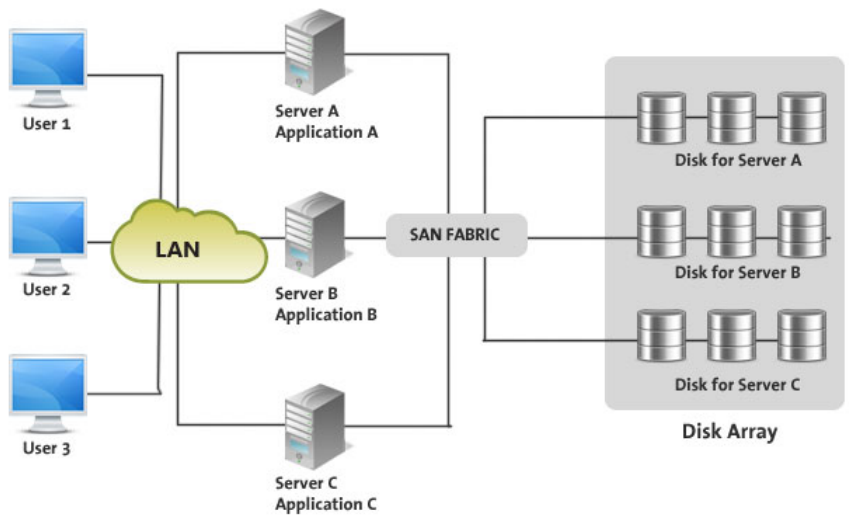

Figure: Networks, Source: www.cloudinfinit.com

Data Center Analysis

- *tpmC* $\rightarrow$ transaction rate per minute defined by Transaction Processing Performance Council (TPC)
- Transaction Performance Example from 2001-2008
  - 40.000 to 520,000 tpmC (Best in 2014: 8.5Million)
  - \$18 to \$1.71 per tpmC (Best in 2014: \$0.19)
  - This gives us: $\underline{\$720,000 \rightarrow \$889,200 \Rightarrow +\$169,200}$

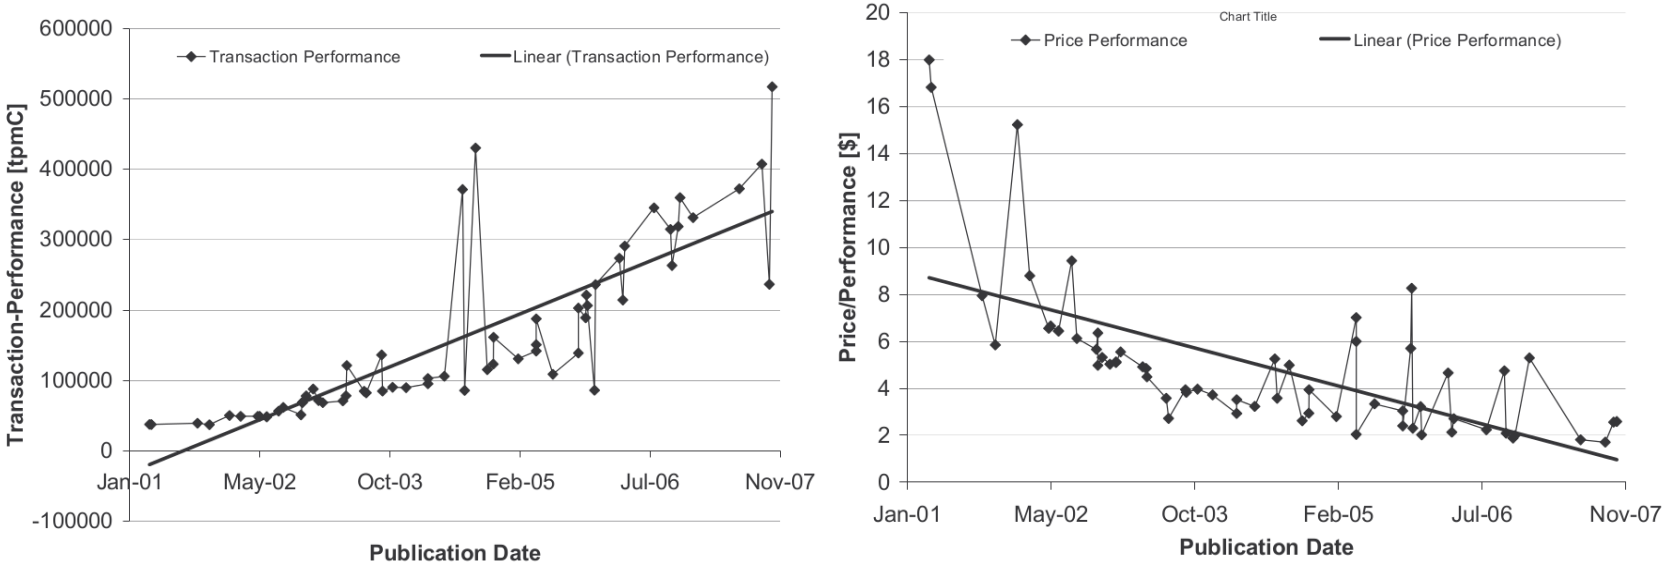

Figure: TPC-C Results,Source: Meikel Poess, Raghu Nambiar

Energy Efficiency Methods

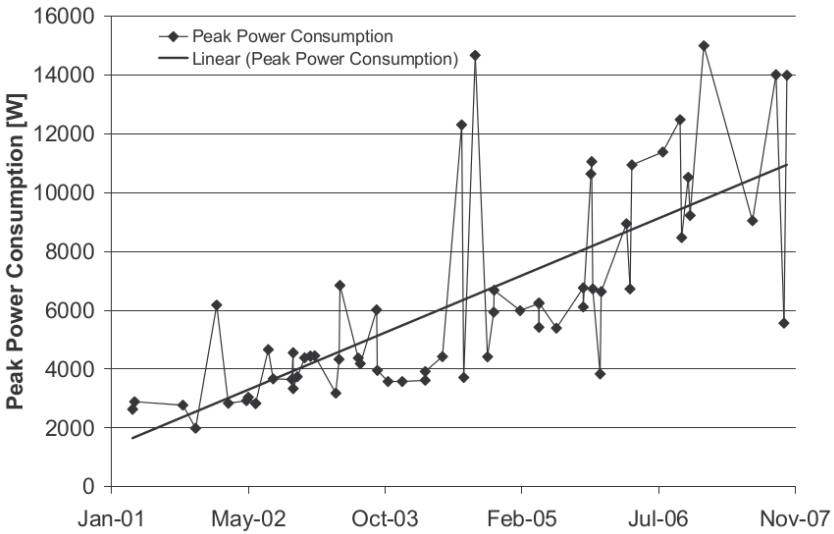

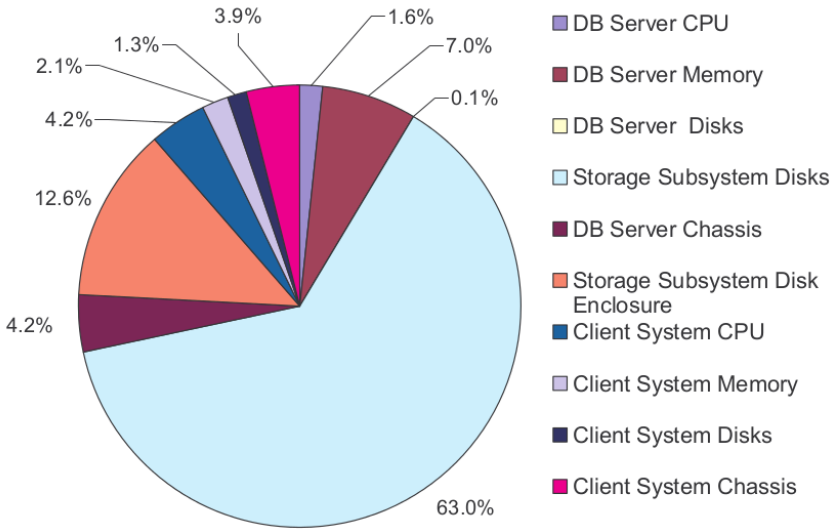

Figure: Power Consumption, Source: Meikel Poess, Raghu Nambiar

Improvements through:

- Newest storage disks
  E.g.: Serial attached SCSI (SAS) disk drives
- Bigger RAM - avoids temporarily
  stored information on disk drives

Improvement with Routers and Switches

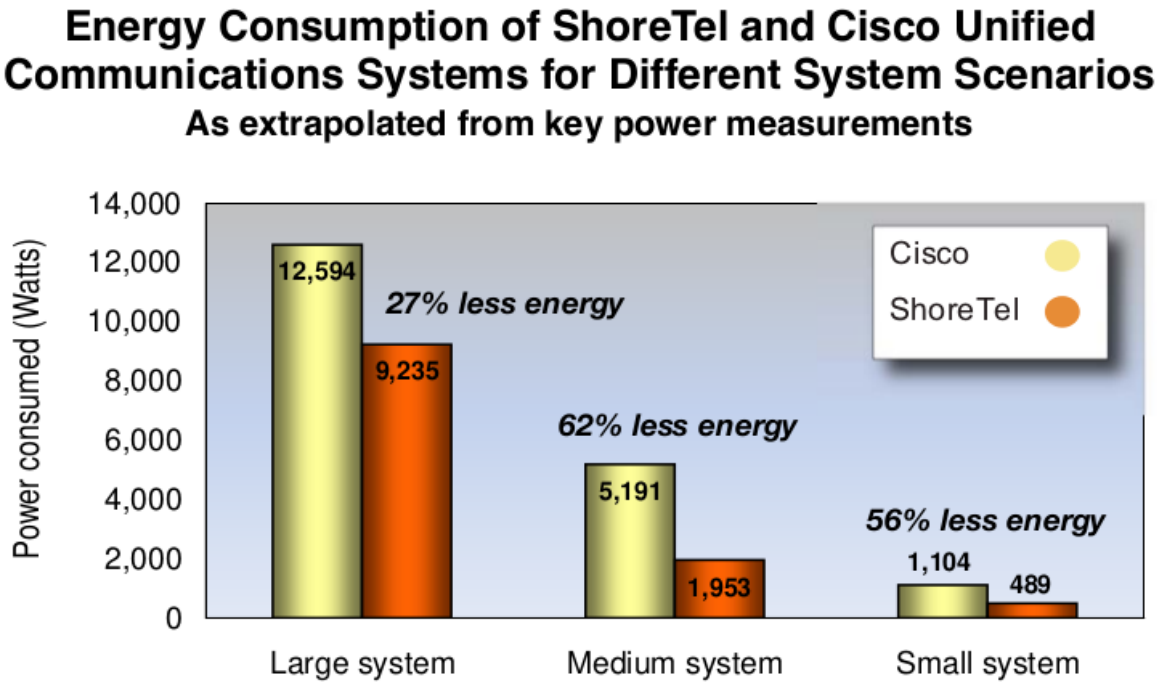

Figure: www.nbcsolutions.net

- Test: 1,500 vs 350 vs 65 Users
- Most energy saved by voice over ip servers and switches
- IBM servers better then Cisco servers in this example

Offloading

Code Offload for Mobile Devices

- Type of connection is relevant

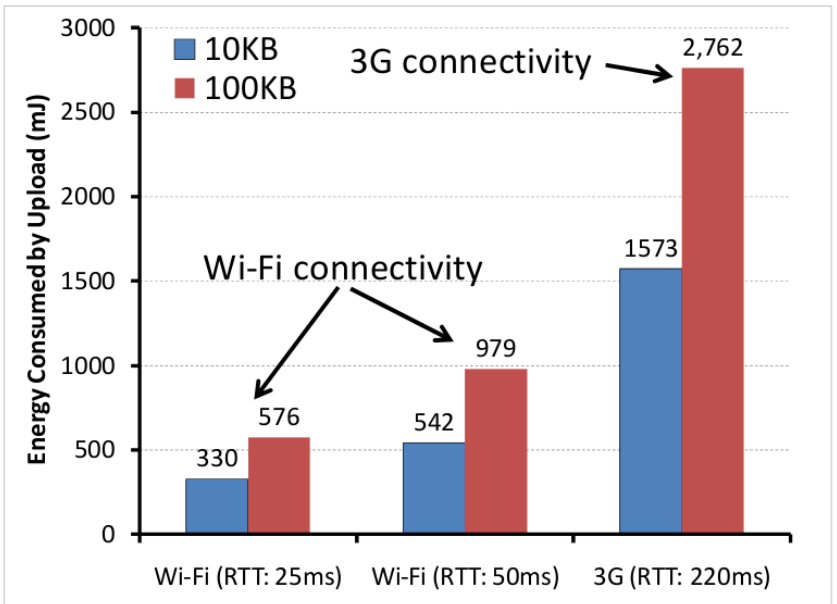

Figure: The Energy Consumption of Wi-Fi vs 3G

Offloading

Code Offload for Mobile Devices

- Offloading generally more efficient for big data usage

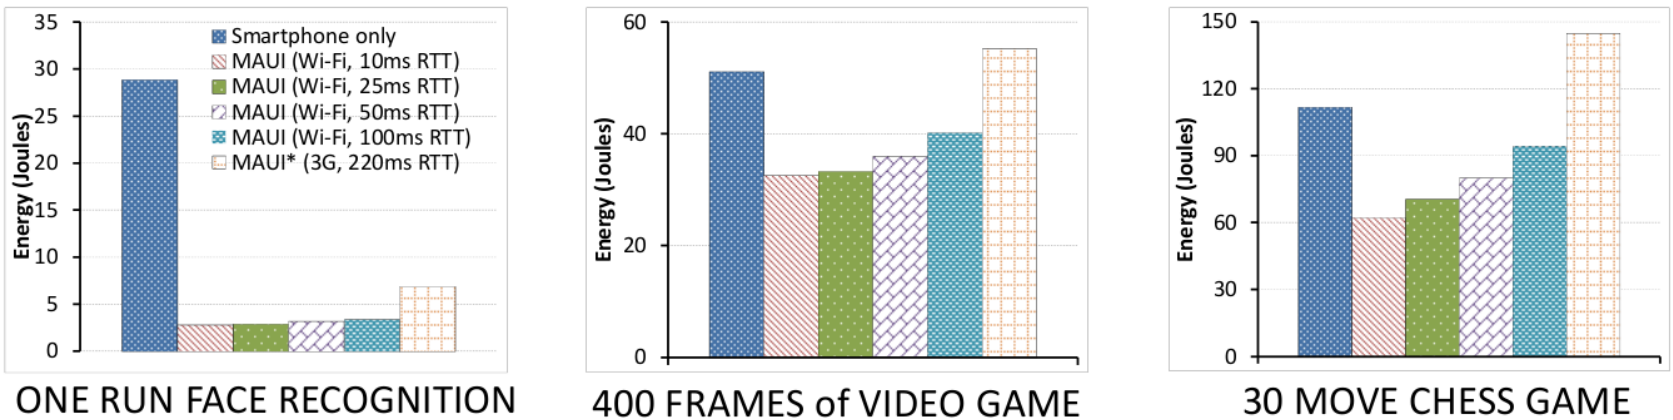

Figure: A comparison of MAUI's energy consumption

Offloading

- TCP/IP Offload Engine (TOE)
  in network interface card or host bus adapter
- Kernel doesn't handle TCP/IP Stack anymore
- Only useful with high speed networks $> 100 MBits/s$

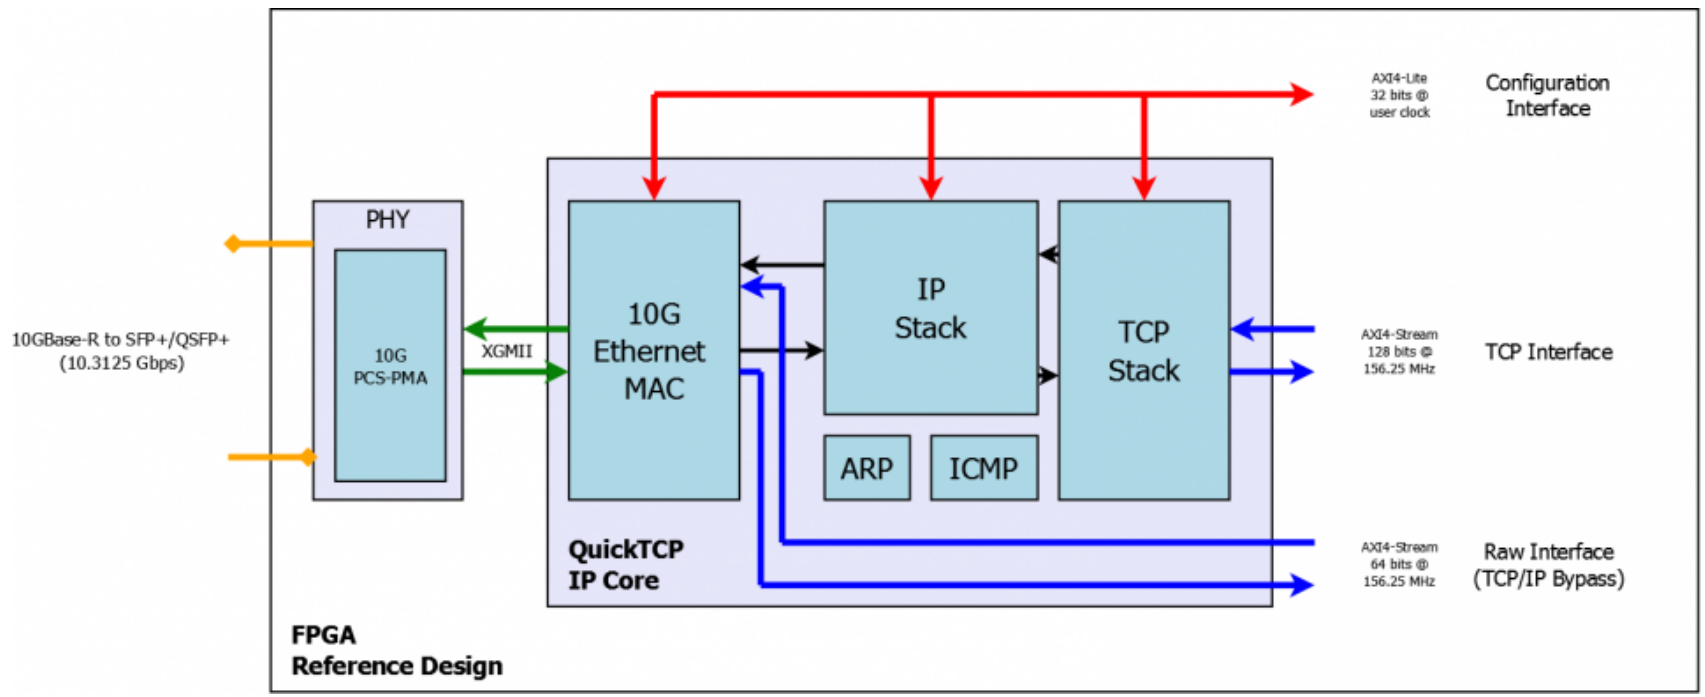

Figure: TOE for 10-gigabit Ethernet, Source: www.plda.com

Remote Direct Memory Access

- Network adapter writes/reads data directly onto/from the RAM
- Kernel bypass
- Works over a Host Bus Adapter (HBA)
- Protocols:
  - IB →Infiniband
  - RoCE → RDMA over Converged Ethernet
  - iWARP → Internet Wide Area RDMA Protocol

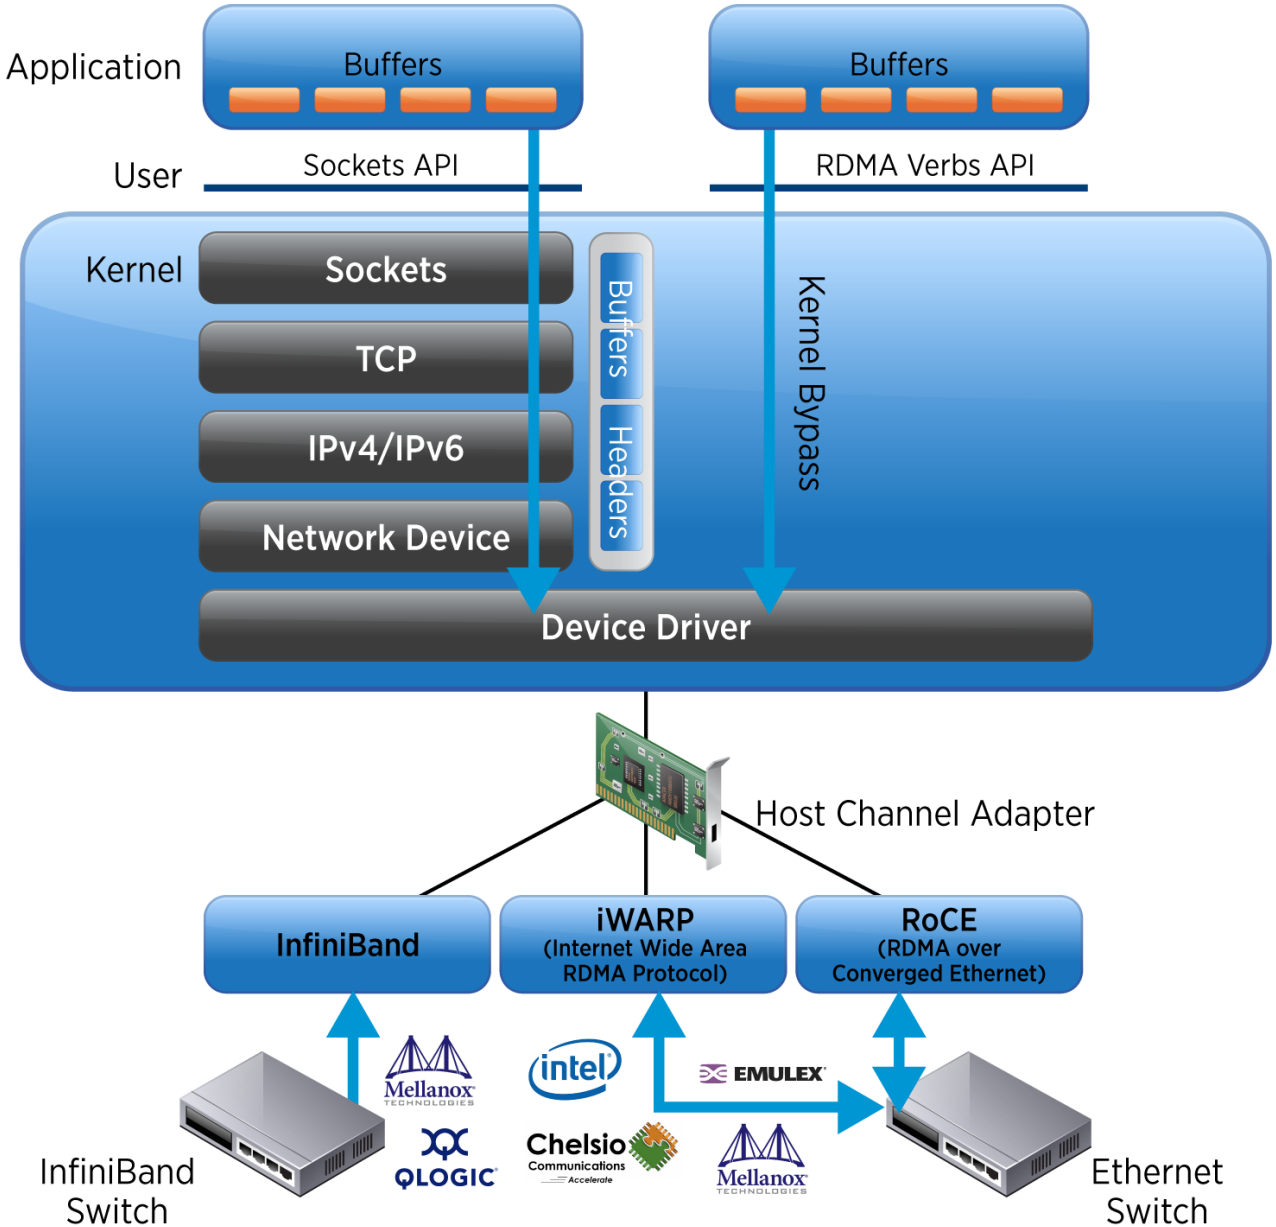

Figure: RDMA Overview, Source: labs.vmware.com

Infiniband vs RDMA over Converged Ethernet

- API is identical across technologies
  $\rightarrow$ Application on RoCE works also for Infiniband.
- IB better overall - faster, less energy consuming
- RoCE preferably used - cheaper

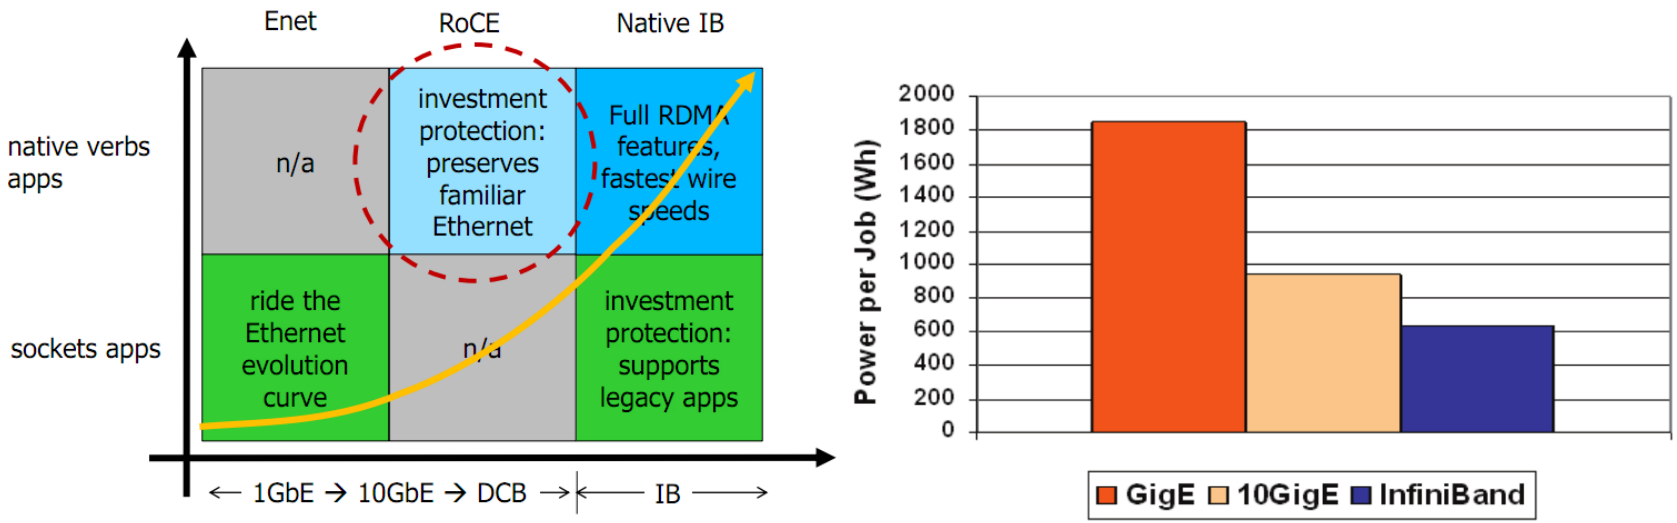

Figure: Trade-offs in choosing an I/O architecture,
Power Consumption of Oil and Gas reservoir simulation application

Conclusion

- Spend more money on better hardware and you'll get more efficiency - just like food: Healthiest, eco-friendly food $\rightarrow$ most expensive
- Only get the speed you really need or you will waste a lot of money
- Future Technologies:
  - 100GBit/s Ethernet
  - New generation of Solid-state drives (SSD)
  - Better data compression
  - Fast CMOS logic / efficient TTL Logic

Literature and Websites

"MAUI: Making Smartphones Last Longer with Code Offload" by Eduardo Cuervo
"Interconnect Analysis: 10GigE and InfiniBand in High Performance Computing" by HPC Advisory Council
"A Performance Study to Guide RDMA Programming Decisions" by Patrick MacArthur, Robert D. Russell
"The Direct Energy Demand of Internet Data Flows" by Vlad C. Coroama*, Lorenz M. Hilty, Ernst Heiri & Frank M. Horn
"Energy Cost, the Key Challenge of Today's Data Centers: A Power Consumption Analysis of TCP-C Results" by Poess and Nambiar in 2008
Issue Paper, August 2014: "Data Center Efficiency Assessment Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers" by the Natural Resources Defense Council
https://www.plda.com/sites/plda.com/files/styles/vt-commerce-image-zoom/public/products/fpga-product/quicktcp-xilinx.png
http://labs.vmware.com/vmtj/toward-a-paravirtual-vrdma-device-for-vmware-esxi-guests
http://www.tpc.org/information/benchmarks.asp
http://www.infoworld.com/article/2627603/infrastructure-storage/fibre-channel-vs–iscsi–the-war-continues.html
http://www.elektronik-kompendium.de/sites/net/0905071.htm
http://www.commprod.com/Solutions/Structured-Cabling/Copper-vs-Fiber-cabling
http://www.cablinginstall.com/articles/print/volume-18/issue-6/features/optical-solutions.html
http://www.lenovo.com/energycalculator/
http://www.stromvergleich.de/durchschnittlicher-stromverbrauch
http://www.nrdc.org/energy/data-center-efficiency-assessment.asp
http://www.netgear.com/business/products/switches/unmanaged-plus/10g-plus-switch.aspx
http://www.amazon.in/TP-Link-TL-SF1024D-24-Port-100Mbps-Switch/dp/B0064GOQ5O
http://www.hpcadvisorycouncil.com/pdf/IB-and-10GigE-in-HPC.pdf