

EENDRA

Energy efficiency for network devices using rate adaptation

DoE Review Meeting, San Francisco CA



Steven Fortune

March 2011

EENDRA - Energy efficiency for network devices through rate adaptation

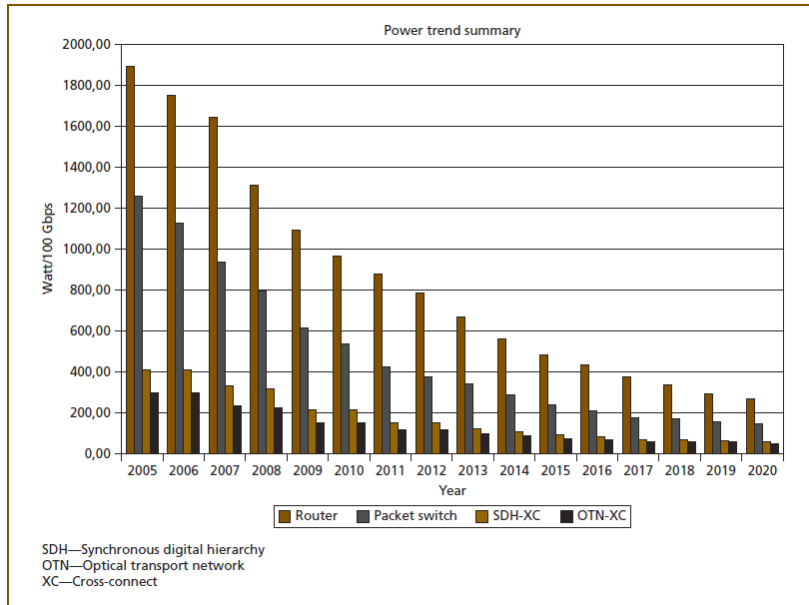
- Bell Laboratories, Alcatel-Lucent
- Steven Fortune, PI
 - also involved Andrea Francini, Lisa Zhang, Spyridon Antonakopoulos, Marco Ricca, BL
- April 1, 2010 - March 31, 2011
- Concept definition study

Project objective

EENDRA

- Develop a scientific understand of methods to reduce energy consumption in data networks using rate adaptation
 - Energy profiling of network elements
 - Energy optimization: the single element case
 - Energy optimization: the multi-element case
 - TCP performance in energy-optimized networks

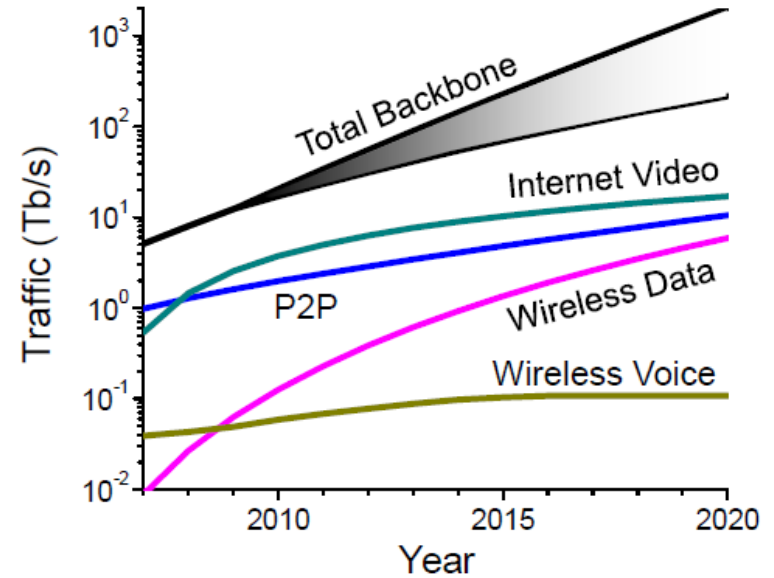
Technology and Usage Trends



Projected port power consumption

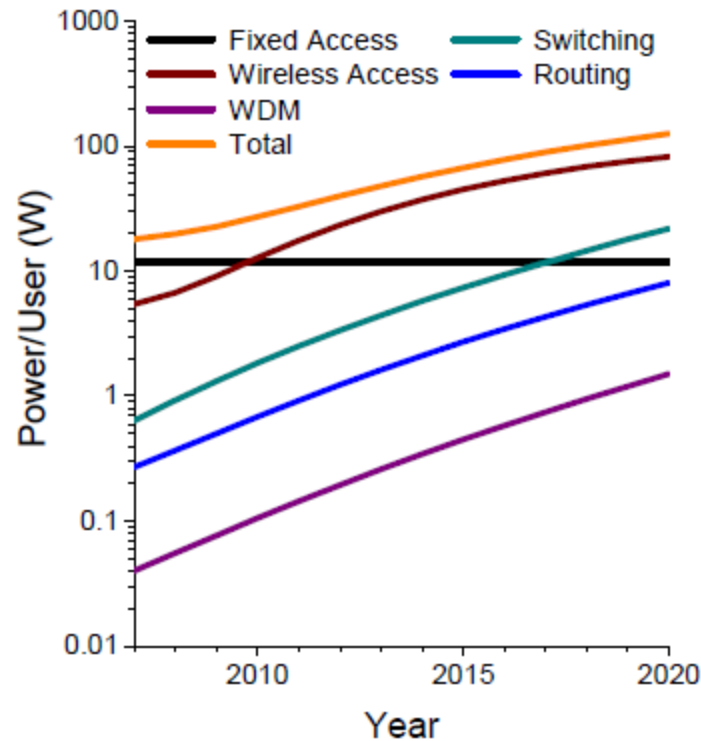
Tamm et al BLTJ V14(4) 2010

Bell Labs Traffic Projections (North America)



Data from: RHK, McKinsey-JPMorgan, AT&T, MINTS, Arbor, ALU, Bell Labs internal Green Touch study, 2010

Composite trend



Bell Labs internal Green Touch study, 2010

wireless access dominates network energy requirements

Opportunity for IP routers:

Rate adaptivity - Energy proportional to load

Average link utilization estimates

enterprise ~2%

[Fraleigh et al, Sprint TR Nov 2001]


core ~30%

[TeleGeography Research Feb 2005]

resilience, traffic fluctuations, over-provisioning

Current router energy use is largely independent of traffic

Goal: router energy use should correlate with traffic



Rate adaptation - single router

Packet processing

determine priority, destination, etc.
significant energy use, done by cpu

Packets arrive on microsecond time scale

100% load: every us
10% load: every 10us

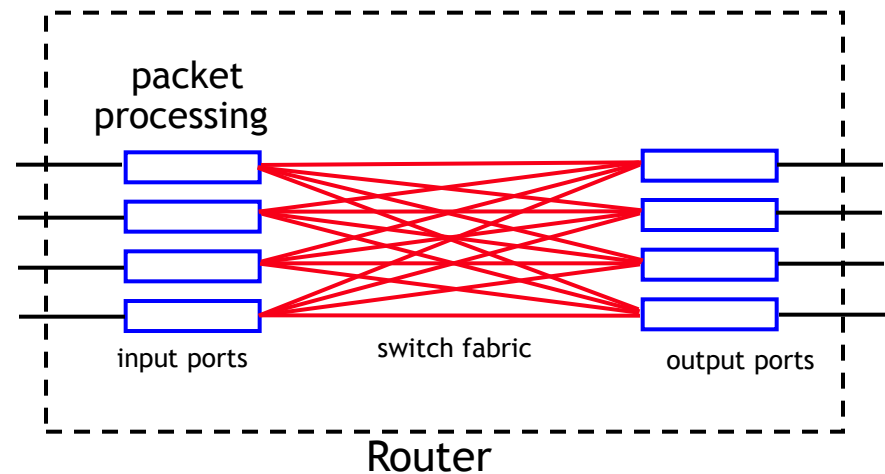
CPU energy control: many options

sleep state, frequency scaling,
dynamic voltage scaling

Current CPU state switching ~ 250 us

Issues:


control algorithms for state selection and queue
management
delay minimization



EENDRA results - Rate adaptation - single router

Single stage of packet processing


sophisticated speed scaling beneficial
relatively complex control algorithm



single-stage power-traffic curve

Multiple stages (typical case)

on/off operation indistinguishable from speed scaling
simpler control algorithm



packet processing pipeline

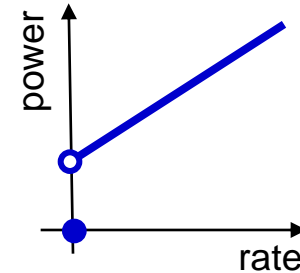
A. Francini, Selection of a Packet-Timescale Rate Adaptation Scheme for Data-Path Network Hardware, 2011, in preparation

Rate adaptation - network-wide optimization


Perfect rate-adaptation unlikely

Alternative approach

network-wide power control/rerouting in response to traffic variations



achievable power curve has nontrivial startup cost



required network as a function of traffic

EENDRA results - rate adaptation/routing study

Study:


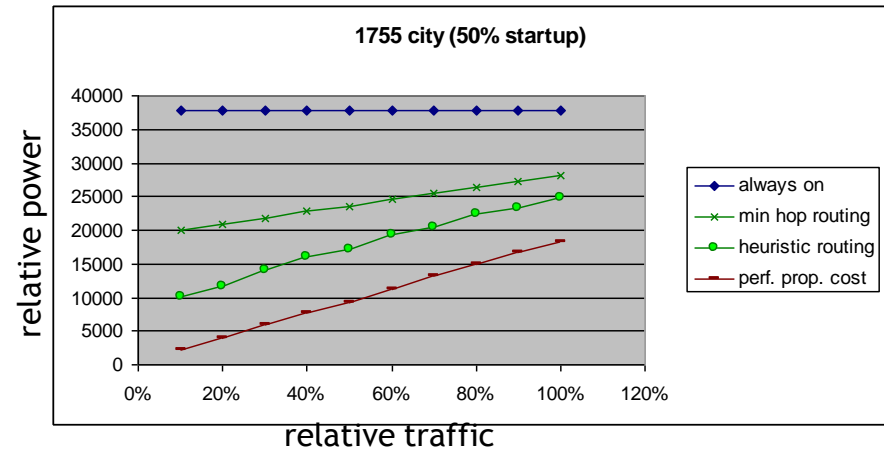
- compare benefits of router-level rate-adaptation (local) and network-wide routing (global)
- real networks, synthetic traffic

Results:

- Strong dependence on startup-cost
- Primary benefit from rate adaptation
- Network-wide routing additional benefit
- Benefits consistent across many scenarios

Caveats:

- Computationally hard problems
- Hardware implementation currently unavailable
- Many constraints on network management
 - protocols, SLAs, delay/congestion, ...
- requires fast restoration



Antonakopoulos et al, Power-aware routing with rate-adaptive network elements, IEEE GreenCom 2010

EENDRA results - Optimization algorithms for minimum power routing

Classic algorithmic problem:

given traffic matrix, network:

route traffic minimizing cost

New problem


given traffic matrix, network, router power profile:

route traffic minimizing power

Results:


Algorithm that gives approximately minimal power profile

Close to 'best possible' (with theoretical computer science complexity assumptions)



	A	B	C	...
A	0	5	3	...
B	5	0	1	...
C	3	1	0	...
⋮	⋮	⋮	⋮	⋮

traffic matrix



Antonakopoulos et al, Minimum-Cost Network Design with (Dis)economies of Scale, FOCS 2010

EENDRA results - Network wide scheduling with rate adaptivity

Classic scheduling problem:

given traffic routes and arrival rates, network:

schedule traffic achieving network-wide stability

New problem:

given traffic routes and arrival rates, network, router power profile:

schedule traffic and set router speed, achieving stability and minimizing total power

Results:

Combine new rate-adaptation protocols with existing scheduling protocols

Trade stability and small end-to-end delay against small power consumption

Antonakopoulos et al, Energy-aware Scheduling
Protocols for Network Stability, IEEE INFOCOM 2011

EENDRA results - TCP interaction/rate adaptation interaction

Transmission Control Protocol (TCP) - Kahn et al, 1970s

- flow control at end points
 - maximize bandwidth utilization, minimize congestion
 - subtle but well-established control algorithm
- potential interaction between rate-adaptation and TCP
 - rate adaptation can introduce additional delay and packet losses
 - interaction with TCP can reduce network rates or increase power consumption
- original intent: Joint simulations of TCP and rate adaptation
 - result: negligible effect on TCP throughput; additional delay may affect UDP
- led to new buffer management protocol : Periodic Early Discard (PED)
 - refinement of RED (random early drop/detection) - slower control loop eliminates oscillations
 - early simulations suggest order-of-magnitude decrease in router buffer size
 - buffers are significant contribution to router power use

A. Francini, "Beyond RED: Periodic Early Detection for On-Chip Buffer Memories in Network Elements, 2011, submitted

Conclusions

Project

- DoE funding is complete 3/31
- Funded two researchers for a year
- Identified potential for ~50% savings in IP router energy, beyond components reductions (e.g. CMOS)
 - requires re-architecting routers, augmenting network management
- Active continuation of research within Bell Labs
 - many engineering, component, algorithmic, management issues
 - eventual reduced operational expense for telecom service providers

By 2015, demonstrate architectures and technologies that yield a **1000-fold improvement in network energy efficiency (using 2020 traffic projections) over 2010 network energy efficiency**

- Global consortium with **experts** from across industry and world's top institutions working together in matrix of **open** innovation
- **Pre-competitive** research that emphasizes clean slate architectures and **out-of-the-box** thinking
- **Broad support** and funding
- **Ambitious goal** in finite period of time