Energy-Efficient Algorithms

Susanne Albers

Chair for Efficient Algorithms

Motivation

Energy is scarce and/or expensive resource.

- Limited availability: Portable, battery-operated devices; sensor networks.

Motivation

- Electricity cost: substantial strain for computing and data centers
  Google: 1 billion $ per year

  “What matters most [...] at Google is not speed, but power, low power because data centers can consume as much electricity as a city.” Eric Schmidt, *NYT* 2002

  Google uses enough electricity to power 200.000 homes. *NYT* 2011

  Low-cost power major criterion where to build data centers. Apple, Facebook, Google

Motivation

- Thermal problems: Most of the energy is converted into heat.

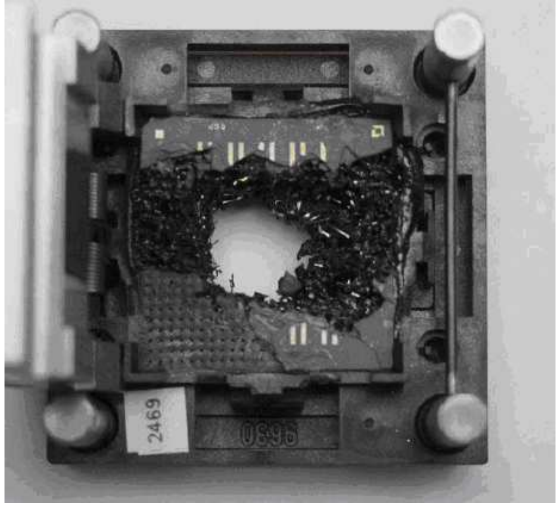

Energy-efficient algorithms

Topics

- Power-down mechanisms: Transition system into sleep state when idle
- Dynamic speed scaling: Microprocessors can run at variable speed
  Intel XScale, Intel Speed Step, AMD PowerNow
- Networks: Optimize transmission energy

Power-down strategies

2-state system

- Active state: r energy units per time unit.
- Sleep state: 0 energy units per time unit.
- Wake-up operation: W energy units.
- When active period starts, system must be in / moved to active state.

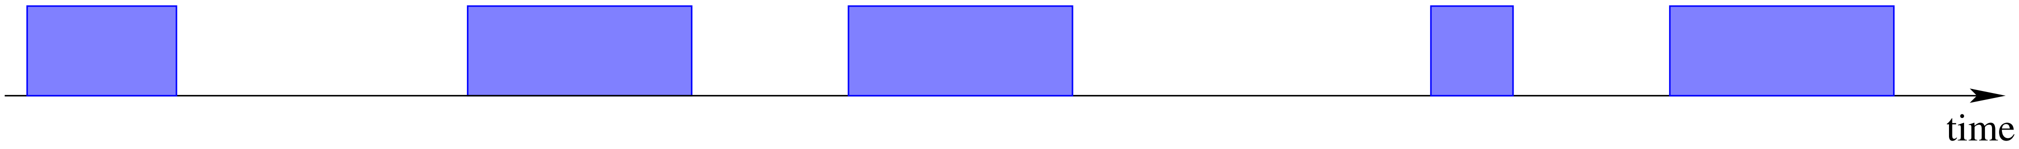

Dynamic speed scaling

Variable-speed microprocessors

Intel XScale, Intel Speed Step, AMD PowerNow

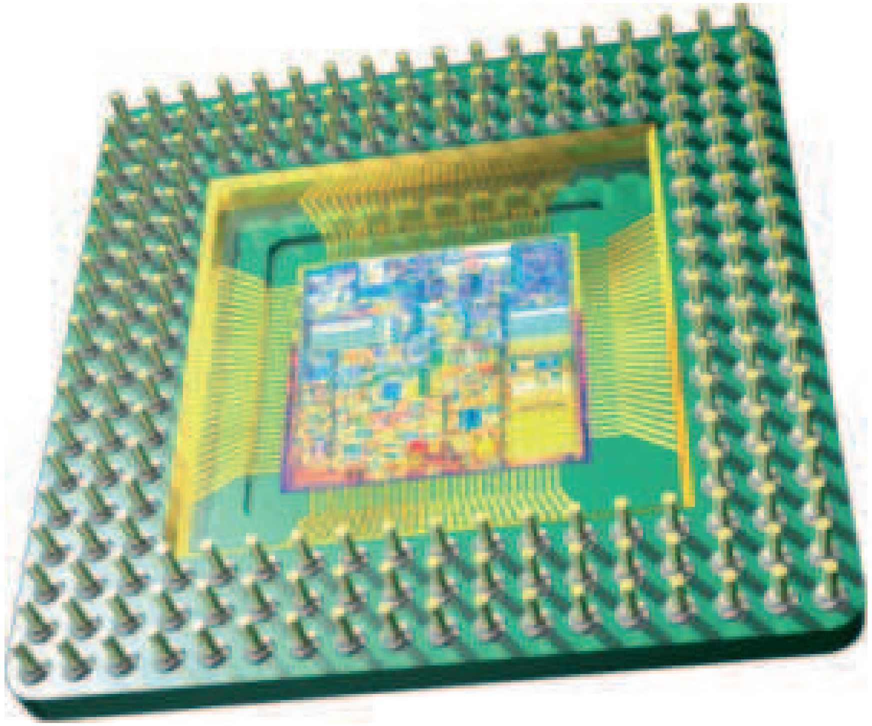

The higher the speed, the higher the energy consumption

Speed s

Power required

$P(s) = s^{\alpha} \quad \alpha > 1$

$P(s)$ = general convex function

Scheduling with deadlines

1 processor

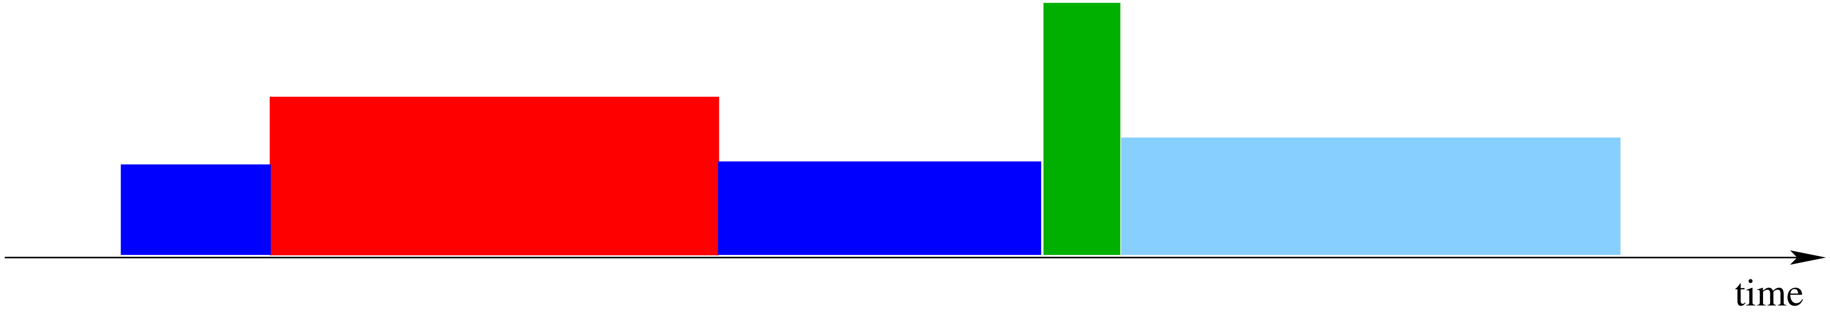

- Speed s Power consumption $P(s) = s^\alpha$ $\alpha > 1$
- $\sigma = J_1, \ldots, J_n$
  J_i: a_i = arrival time
  b_i = deadline
  v_i = processing volume $t = v_i / s$
- Preemption allowed
- Construct feasible schedule minimizing total energy consumption.

Yao, Demers, Shenker FOCS'95

Data aggregation

Tree topology

Data (packets) to be sent to sink

Pakets may be aggregated

Energy: 1 per sending operation

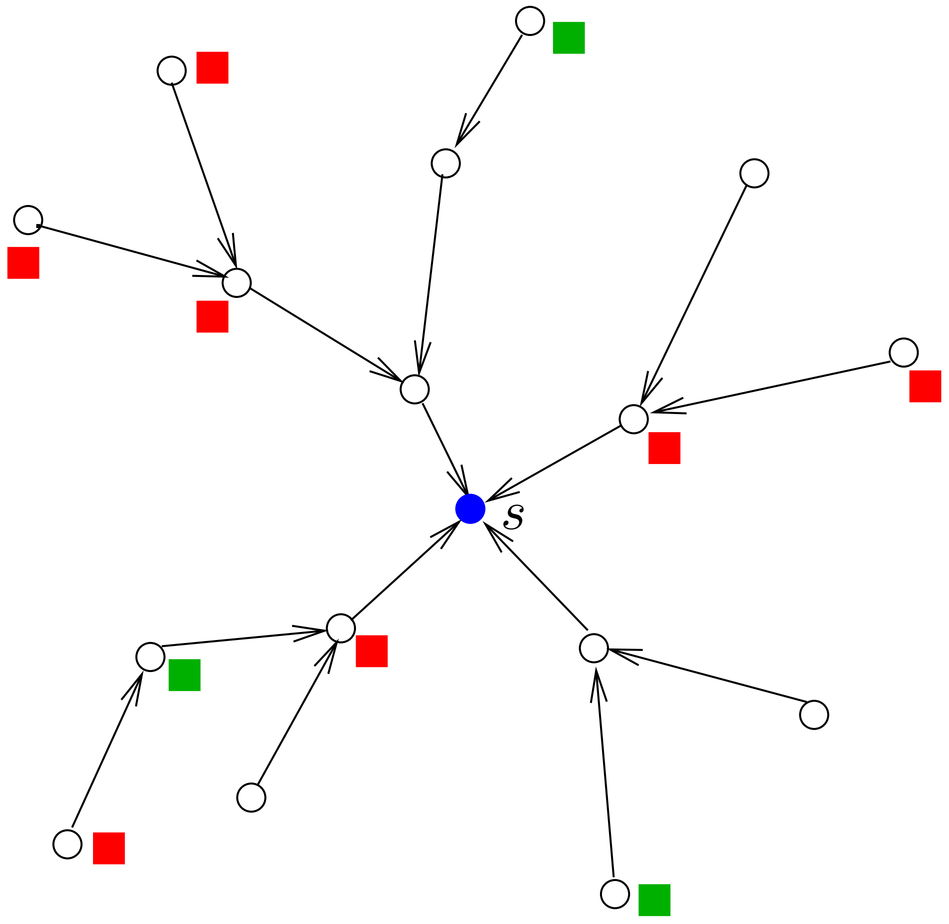

Routing

$G = (V, E)$

D demands (s_i, t_i)

l_e = load on e

$f(l) = c + l^\alpha$

$\min \sum_e f(l_e)$

Andrews, Antonakopoulos,
Zhang FOCS'10

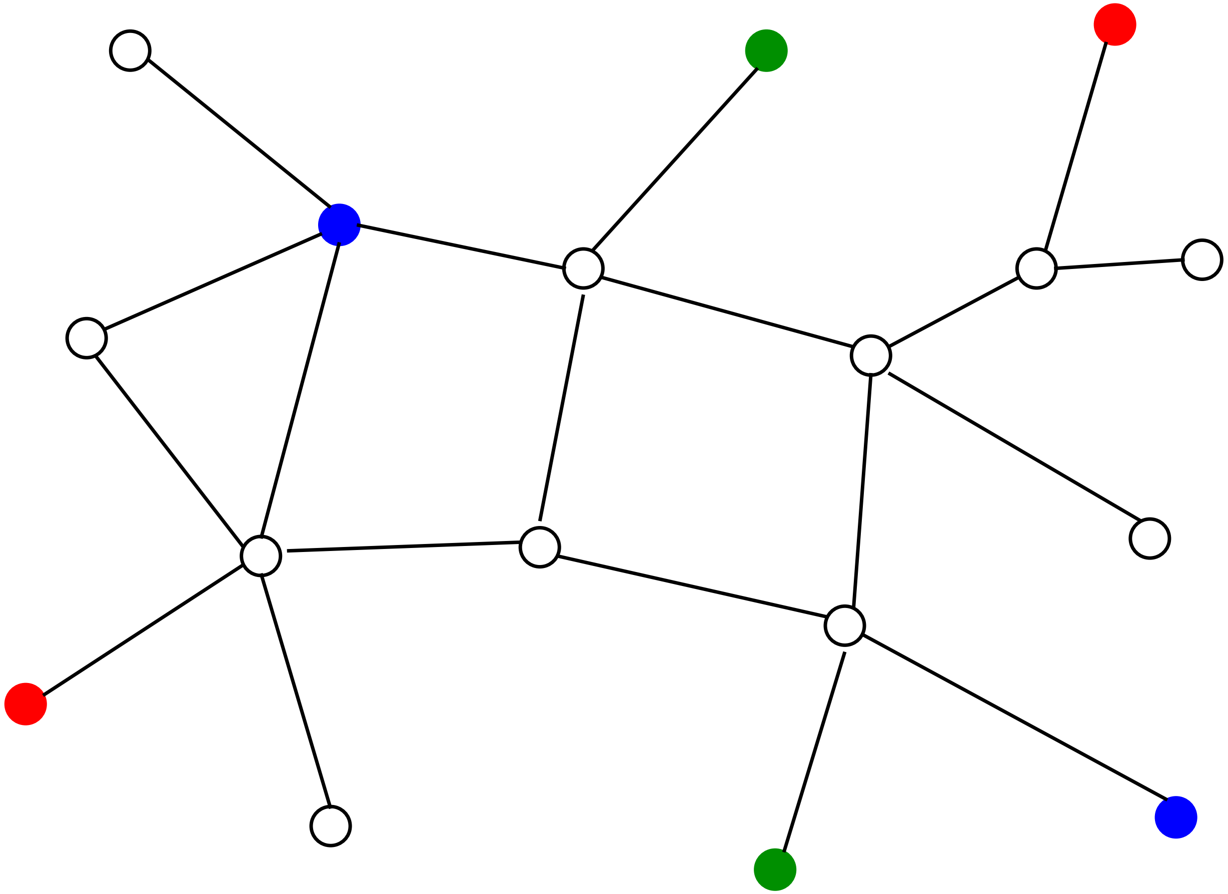

Routing

$G = (V, E)$

D demands (s_i, t_i)

l_e = load on e

$f(l) = c + l^\alpha$

$\min \sum_e f(l_e)$

Andrews, Antonakopoulos,
Zhang FOCS'10

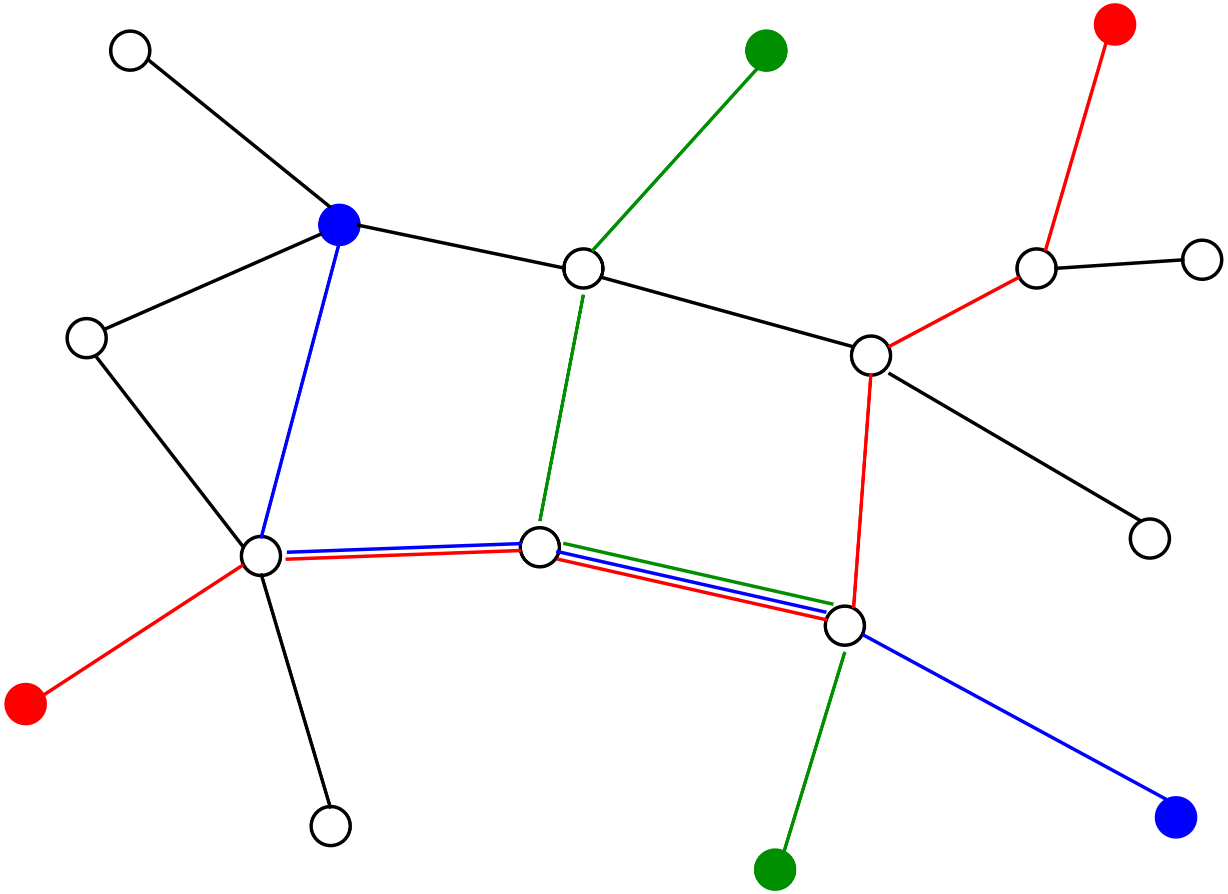

Routing

$G = (V, E)$

D demands (s_i, t_i)

l_e = load on e

$f(l) = c + l^{\alpha}$

$\min \sum_e f(l_e)$

polylog(n, D)-approximation

Andrews, Antonakopoulos,
Zhang FOCS'10

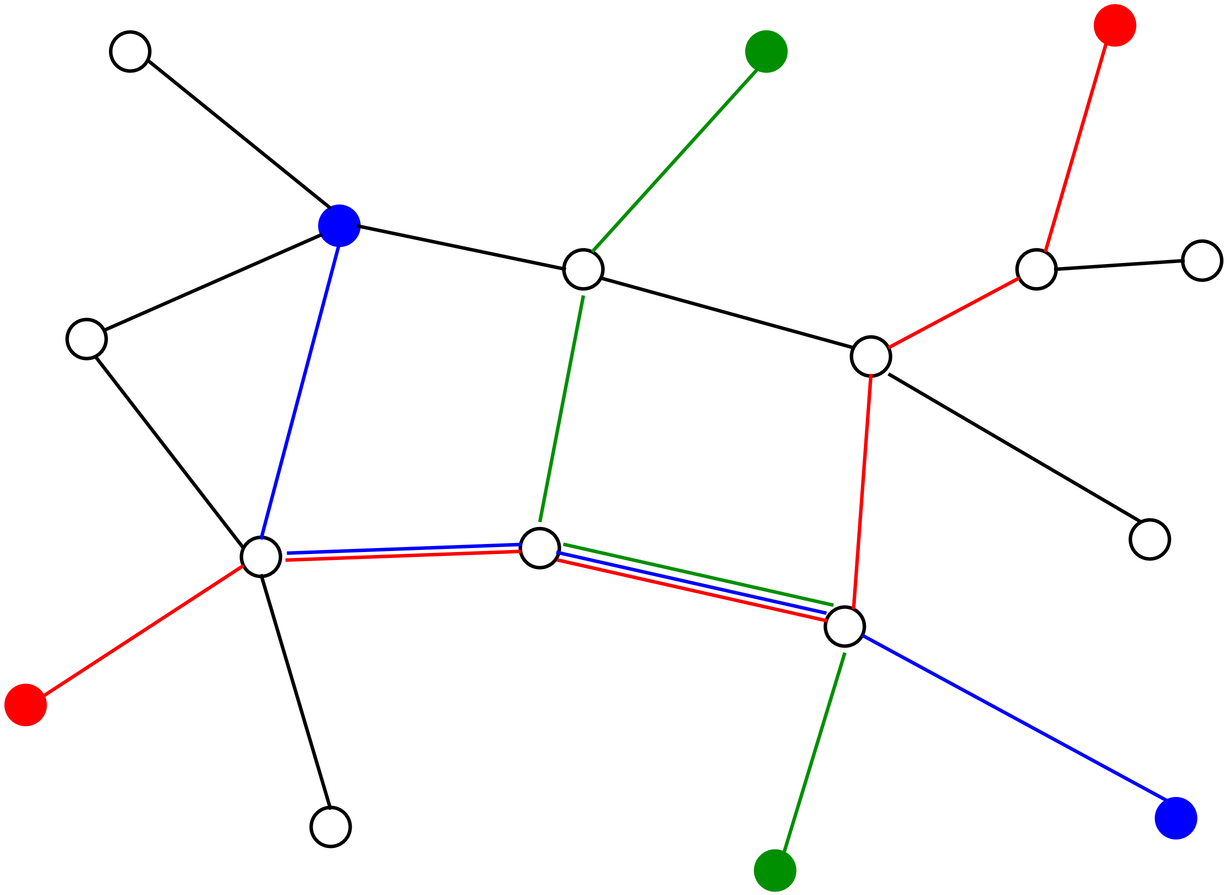

Topics

1. S. Irani, S. K. Shukla and R. K. Gupta. Online strategies for dynamic power management in systems with multiple power-saving states. ACM Trans. Embedded Comput. Syst. 2(3): 325-346, 2003.
2. J. Augustine, S. Irani, and C. Swamy. Optimal power-down strategies. SIAM J. Comput, 37(5): 1499-1516, 2008.
3. P. Baptiste. Scheduling unit tasks to minimize the number of idle periods: A polynomial time algorithm for offline dynamic power management. Proc. 17th Annual ACM-SIAM Symposium on Discrete Algorithms, 364-367, 2006.
4. F. Yao, A. Demers and S. Shenker. A scheduling model for reduced CPU energy. Proc. 36th Annual Symposium on Foundations of Computer Science, 374-382, 1995. M. Li, F. F. Yao. An efficient algorithm for computing optimal discrete voltage schedules. MFCS 2005: 652-663.
5. N. Bansal, D. P. Bunde, H., Kirk Pruhs. Average rate speed scaling. LATIN 2008: 240-251.

Topics

6. N. Bansal, T. Kimbrel and K. Pruhs. Speed scaling to manage energy and temperature. J. ACM 54(1): 2007.
7. S. Albers, F. Mller and S. Schmelzer. Speed scaling on parallel processors. Proc. 19th Annual ACM Symposium on Parallel Algorithms and Architectures, 289 - 298, 2007.
8. K. Pruhs, P. Uthaisombut, G. Woeginger. Getting the best response for your erg. ACM Trans. Algorithms, 4(3), 2008: 1-17.
9. N. Bansal, H. Chan, K. Pruhs. Speed scaling with an arbitrary power function. SODA 2009: 693-701.
10. D. P. Bunde. Power-aware scheduling for makespan and flow. J. Scheduling 12(5): 489-500 (2009).

Topics

11. S. Irani, S. Shukla and R. Gupta. Algorithms for power savings. ACM Transactions on Algorithms 3(4), 2007.
12. S. Albers and A. Antoniadis: Race to idle: New algorithms for speed scaling with a sleep state. SODA 2012: 1266-1285
13. L. Becchetti, P. Korteweg, A. Marchetti-Spaccamela, M. Skutella, L. Stougie and A. Vitaletti. Latency-Constrained Aggregation in Sensor Networks. ACM Transactions on Algorithms 6(1), 2009.
14. M. Andrews, S. Antonakopoulos and L. Zhang. Minimum-cost network design with (dis)economies of scale. FOCS 2010: 585-592.