INVENTEURS DU MONDE NUMÉRIQUE

Energy-Efficient Overlay for Data Transfers in Private Networks

Anne-Cécile Orgerie & Laurent Lefèvre
annececile.orgerie@ens-lyon.fr, laurent.lefevre@inria.fr

ICON 2011, Singapore, De. 2011

EQUIPE PROJET
RESO / LIP
Ecole Normale
Supérieure de Lyon

ENS DE LYON

Energy : 1st limiting factor for large scale systems like networks

- Wired Networks are not green and must be
- Networks energy flat models/equipments
- Energy efficiency/Energy awareness → new usages
- Green leverages in wired networks
  - Hardawre improvements
  - Shutdown (On/Off, LPI, proxies)
  - Slowdown (ALR)
  - Coordination (EA routing)
- GreenTouch initiative

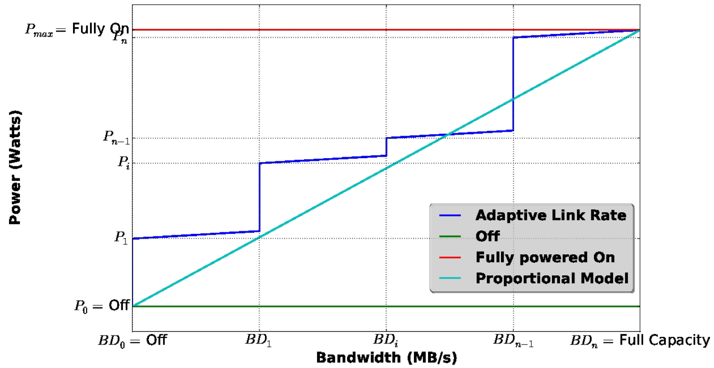

Scenario: bulk data transfers with advance reservations on private networks

Bulk data transfers: large volume of data to transfer, moldable/malleable, with deadline

Advance bandwidth reservations: bandwidth provisioned for the transfer (no resource competition)

Dedicated wired networks: private, managed by one entity, secure

Advance Reservation to avoid overprovisioning and to allow usage scheduling

Assumptions

Network devices

- are able to adapt their transmission rate to the load (Adaptive Link Rate),
- can be entirely switched on and off,
- can turned on and off their ports independently.

Routing is symmetric.

Energy profiles of all the devices are known.

Private network

Deadline, volume of data, source and destination

Not store-and-forward approach for data transfers (end-to-end reservations)

Scheduling on all the agendas of the path

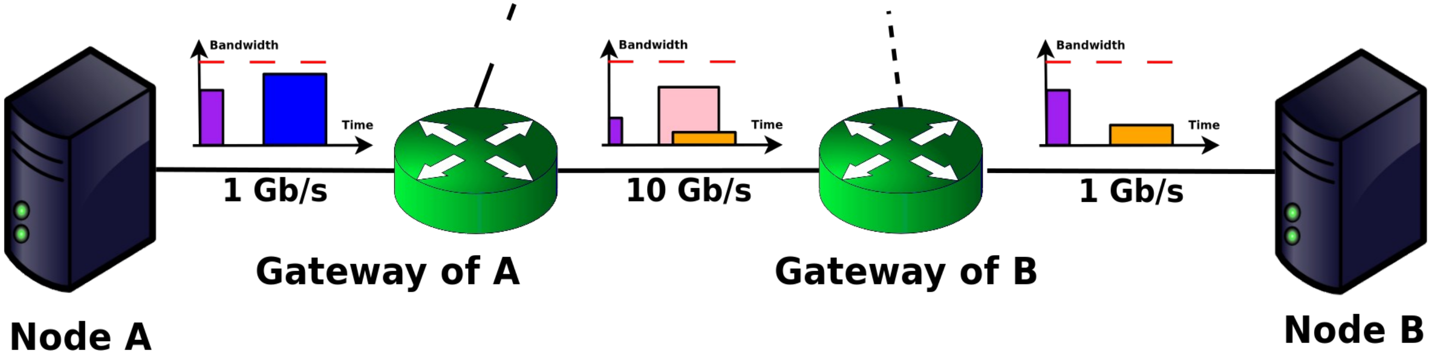

HERMES : High-level Energy-awaRe Model for bandwidth reservation in End-to-end networkS

- Switching off unused nodes
- Distributed network management
- Energy-efficient scheduling with reservation aggregation
- Usage prediction to avoid on/off cycles
- Minimization of the management messages
- Usage of DTN (Disruptive-Tolerant Network) for network management purpose

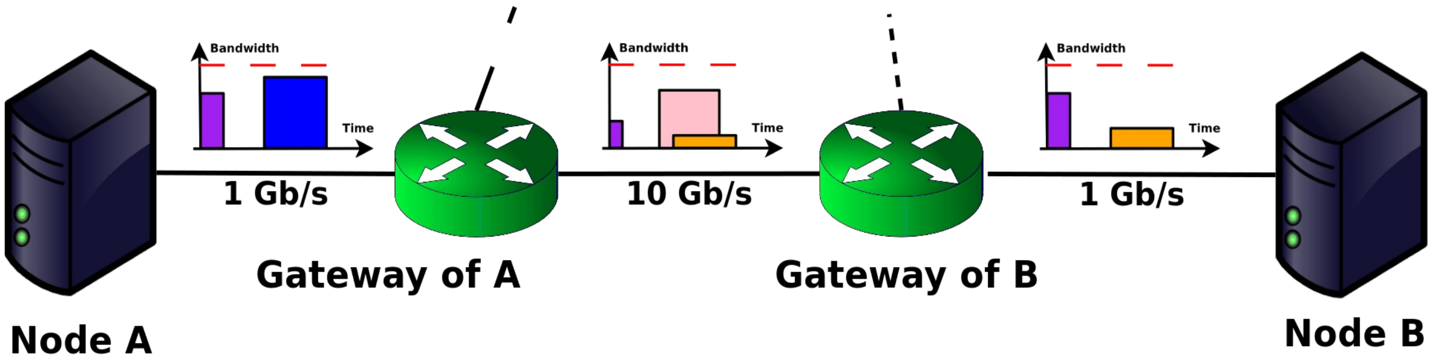

Energy-Efficient Bandwidth Reservation for Bulk Data Transfers in Dedicated Wired Networks", Anne-Cécile Orgerie, Laurent Lefèvre, and Isabelle Guérin-Lassous, The Journal of SuperComputing, Special issue on Green Networks, 2011

Energy-Efficient Framework for Networks of Large-Scale Distributed Systems, Anne-Cécile Orgerie and Laurent Lefèvre. ISPA 2011 : The 9th IEEE International Symposium on Parallel and Distributed Processing with Applications, Busan, Korea, May 2011

HERMES Gateways & Routers

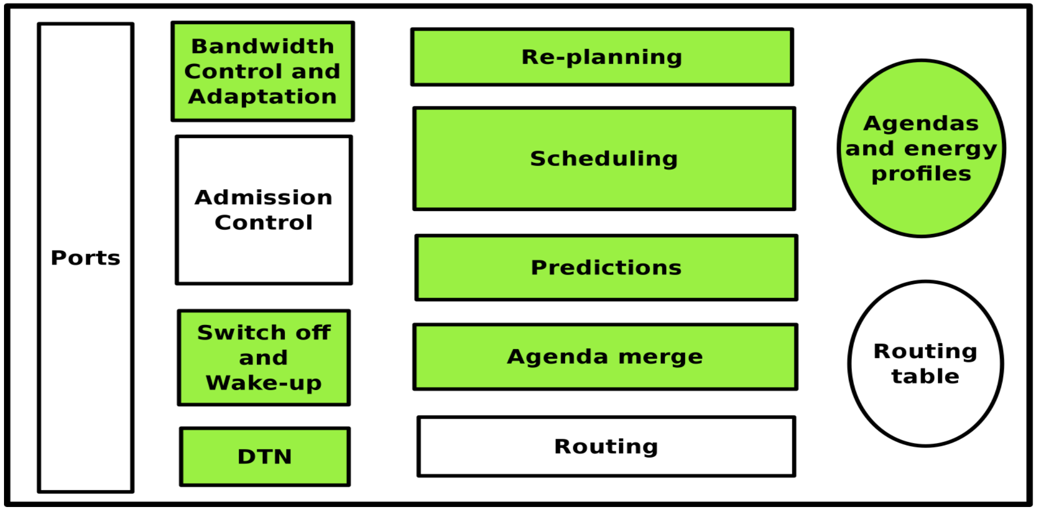

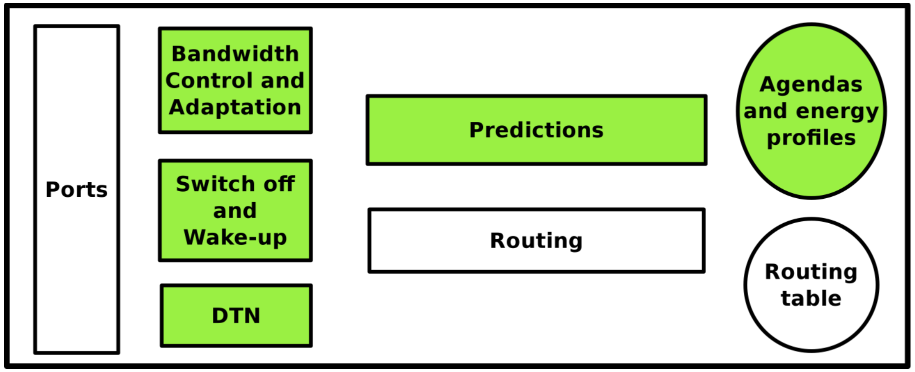

Hermes : Reservation process

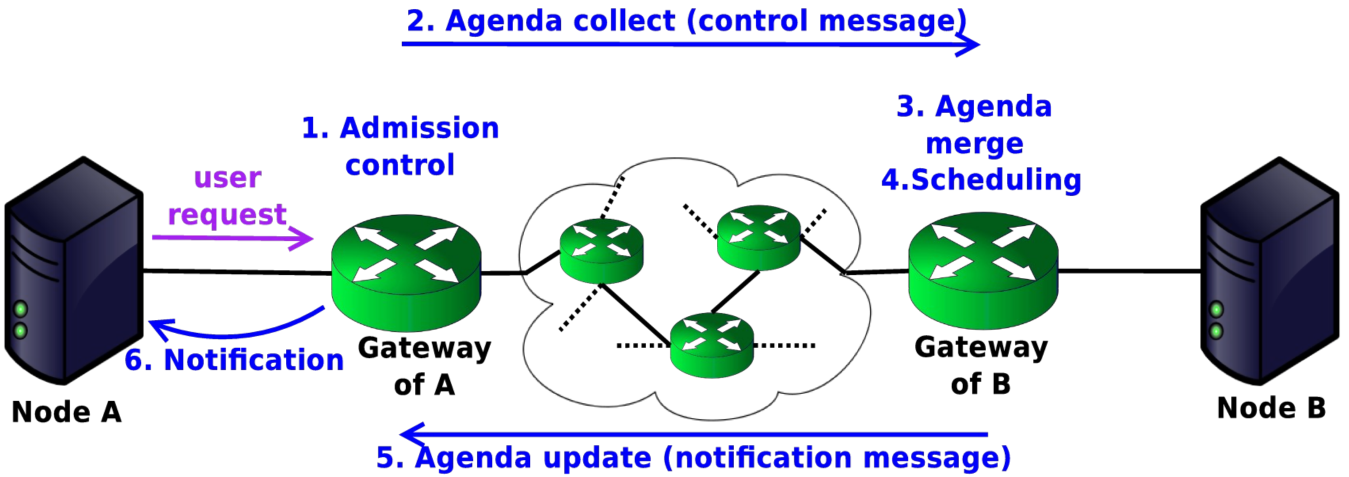

Hermes

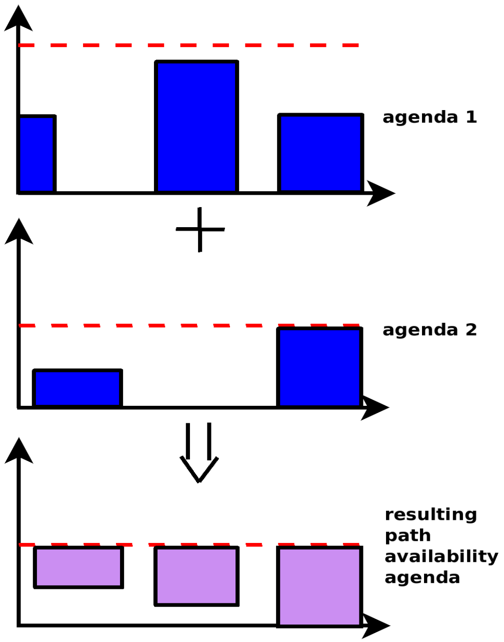

Agenda fusion

ABR scheduling

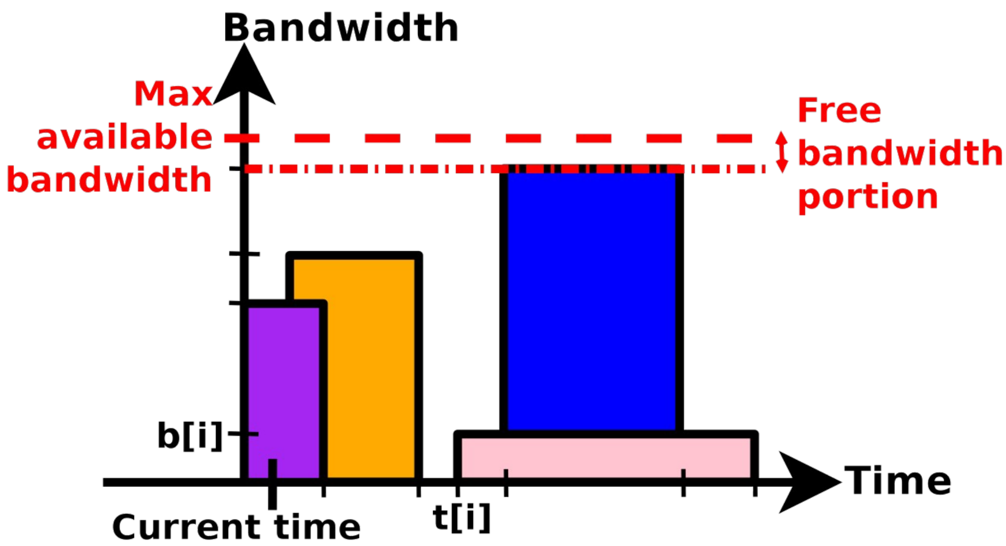

HERMES advanced functionalities

Locking :

- Supporting multiple/long paths reservations
- Focus on network interfaces

Re-planning capabilities

- On-line → off-line
- Depending on reservations types
- User feedback
- Difficult in decentralized approach

Simulation framework

- BoNeS (Bookable Network Simulator)
- Written in Python (6,000 lines)
- Generates random network with the Molloy & Reed method or uses configuration file
- Generates traffic according to statistical laws:
- - submission times (log-normal distribution)
- - data volumes (negative exponential)
- - sources and destinations (equiprobability)
- - deadlines (Poisson distribution)

ECOFEN energy model

Energy Consumption mOdel For End-to-end Networks:

Part of the GreenTouch project (SEASON)

NS3 simulator

Chassis

Linecards

Ports

Bandwidth

Time

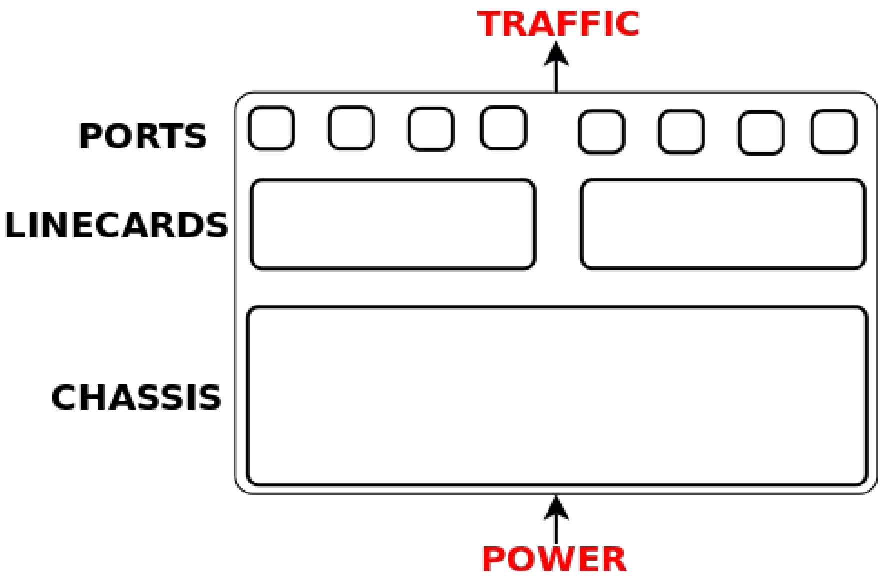

ECOFEN energy models for interfaces

off: the interface is off;

- max-power: the interface always uses the maximal power;
- active/sleep: the energy consumption of the interface depends on the state of the interface (active or sleep);
- bandwidth-dependent: the energy consumption of the interface depends on the bandwidth usage (linear dependence);
- rate-dependent: the energy consumption of the interface depends on the negotiated rate and on the band- width usage;
- complete: the energy consumption of the interface depends on the negotiated rate, on the bandwidth usage (number of sent Bytes) and on the size of the packets.

ECOFEN energy models for chassis

off: the equipment is off;

- max-power: the equipment always uses the maximal power;
- smart-linecards: the equipment can turn off unused linecards;
- dvs: the equipment can turn off unused linecards and use DVS techniques for its processors depending on the traffic load.

The energy models used for the equipments and their interfaces can be chosen independently.

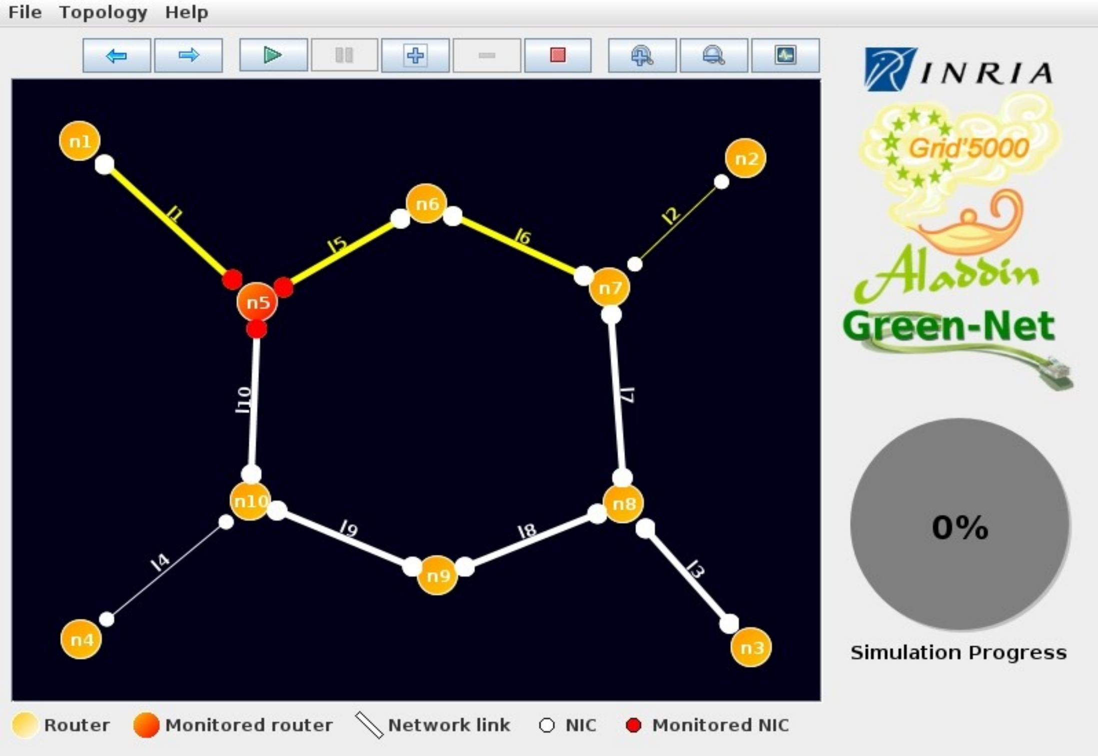

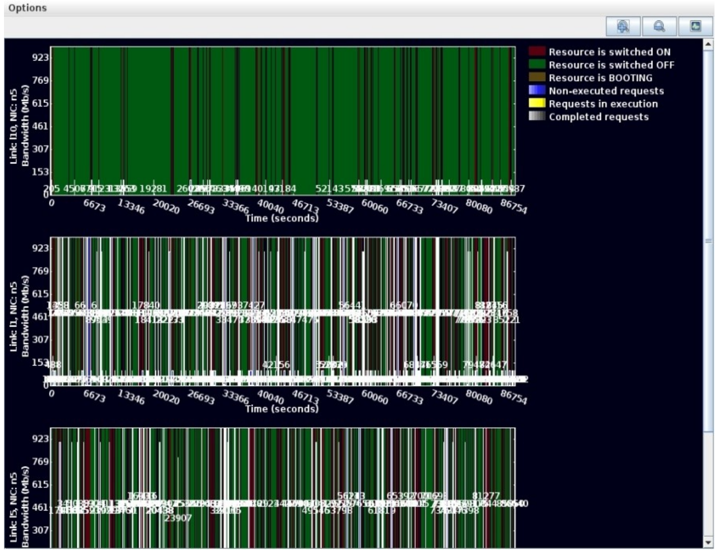

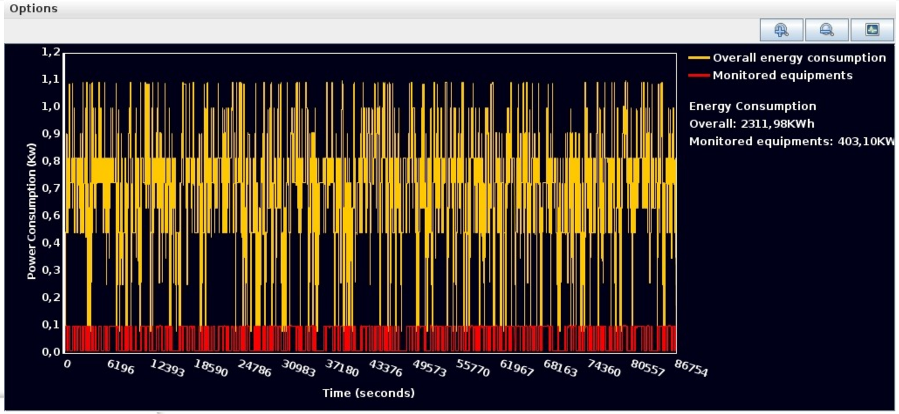

10 SuperComputing demo, Marcos Dias de

Comparison of scheduling

- **First**: the reservation is scheduled at the earliest possible place;
- **First green**: the reservation is aggregated with the first possible reservation already accepted;
- **Last**: the reservation is scheduled at the latest possible place;
- **Last green**: the reservation is aggregated with the latest possible reservation already accepted;
- **Green**: HERMES scheduling;
- **No-off**: first scheduling without any energy management.
- → always before deadline

Hermes results : random network

- Network simulated: 500 nodes, 2 462 links.
- Random Network (Molloy & Reed method)
- All the nodes can be sources and destinations.
- Time to boot: 30 s.; time to shutdown: 1 s.
- 1 Gbps per port routers

Component	State	Power
Chassis	ON	150 W
	OFF	10 W
Port	1 Gbps	5 W
	100 Mbps	3 W
	idle, 10 Mbps	1 W

- 30% workload : Energy consumption in Wh

Scheduling	No off	First	First green	Last	Last green	Green
Average	412 306	205 270	203 844	204 949	196 260	203 342
Standard deviation	2 685	2 477	1 938	2 375	2 695	2 145
Accepted volume (Tb)	2 148	2 148	2 128	2 014	1 853	2 149
Cost in Wh per Tb	191.92	95.55	95.78	101.74	105.92	94.60

- Average occupancy per link
- Compared to current case (no-off), HERMES could save 51%, 46% and 43% of the energy consumed depending on the workload

Workload	No off	First	First green	Last	Last green	Green
31%	191.92	95.55	95.78	101.74	105.92	94.60
46%	149.84	81.61	81.95	87.74	92.40	80.63
61%	130.45	74.73	74.91	80.09	84.63	73.79

Scenario : Fedwire network

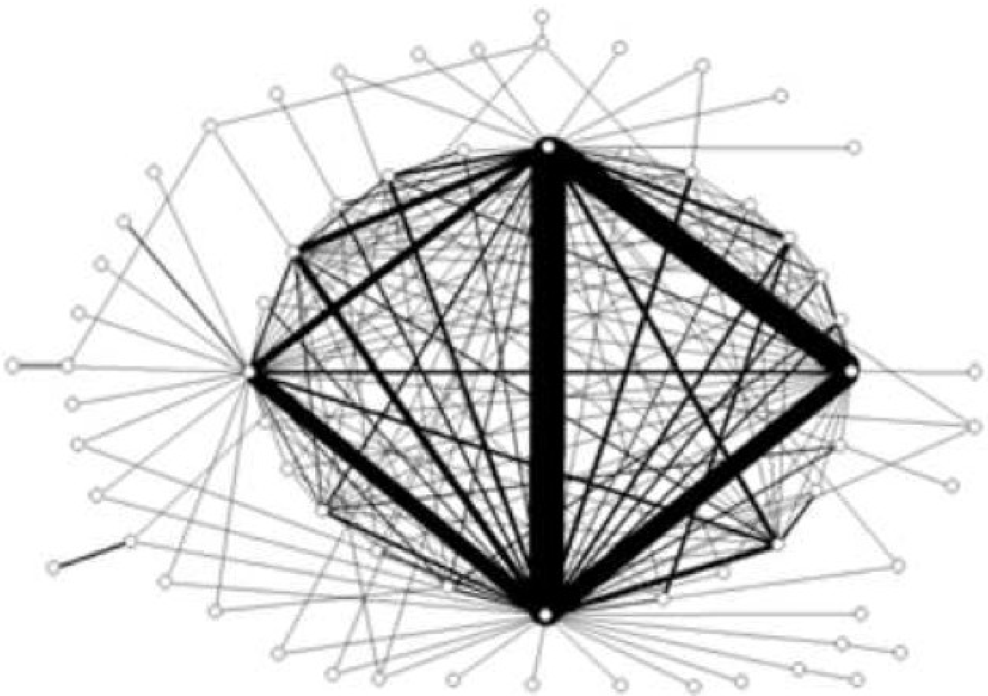

75% of daily amount of transferred data.
66 nodes and 181 links, (25 nodes -> densely connect network)

From K. Soramki, M. Bech, J. Arnold, R. Glass, and W. Beyeler, “The Topology of Interbank Payment Flows,” *Physica A: Statistical Mechanics and Its Applications*, vol. 379, no. 1, pp. 317–333, 2007.

Results

Scheduling	First	First green	Last	Last green	Green	No off
Average (Wh)	606 542	604 531	606 819	605 199	605 743	1 050 544
Standard deviation	7 106	6 994	7 012	7 010	7 090	1 511
Accepted requests	98.6%	98.3%	97.9%	97.5%	98.6%	98.6%
Cost in Wh per Tb	4 113	4 119	4 163	4 176	4 106	7 123

TABLE I
ENERGY CONSUMPTION FOR THE CORE OF THE FEDWIRE INTERBANK PAYMENT NETWORK.

20% average load per node
Green (Hermes) accepts as many requests a First but less energy consumption

Scenario : enterprise network

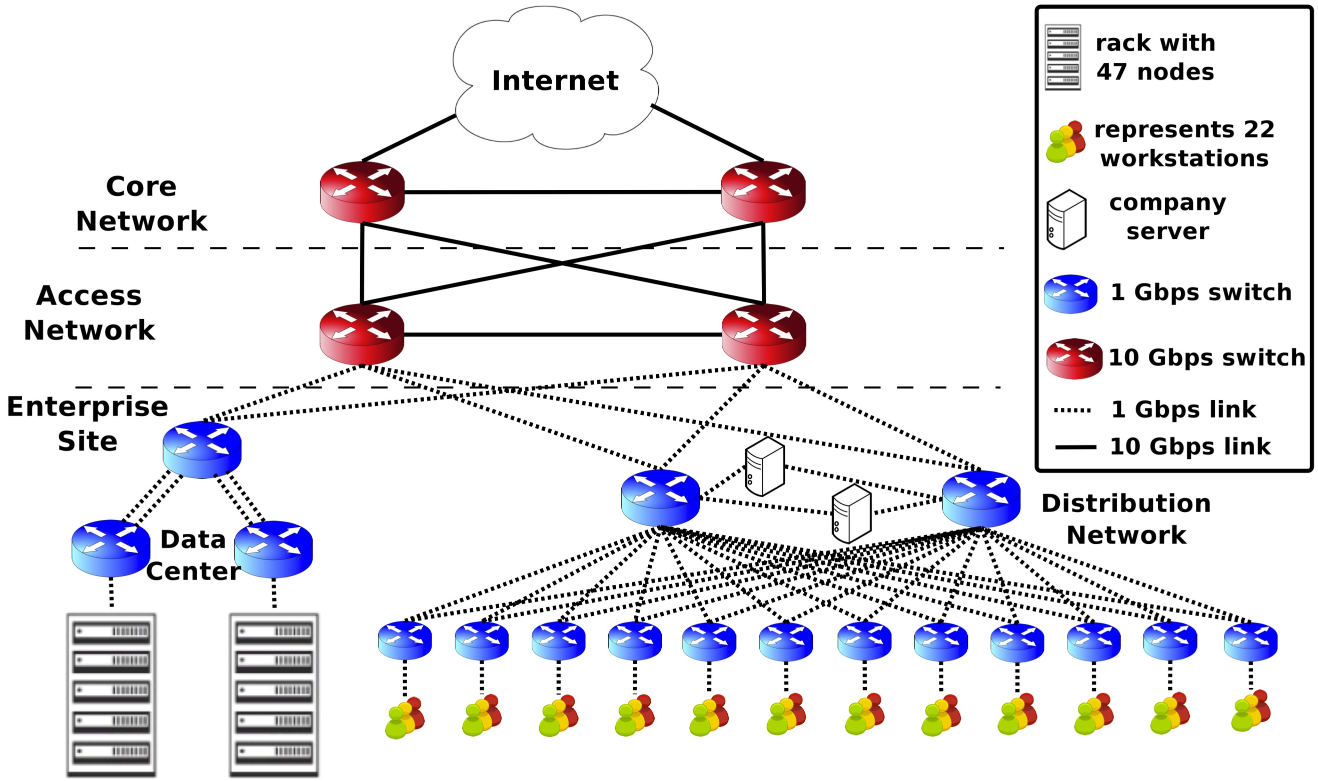

21 routers, 360 end-hosts and 404 links

Inspired from P. Oppenheimer, *Top-Down Network Design*. Cisco Press, 2010.

Results

Scheduling	First	First green	Last	Last green	Green	No off
Average (Wh)	80 987	80 990	80 255	78 611	80 335	103 466
Standard deviation	392	405	316	333	307	563
Accepted volume (Tb)	19.402	19.306	19.098	17.748	19.400	19.402
Cost in Wh per Tb	4 174	4 195	4 202	4 429	4 141	5 333

TABLE II
ENERGY CONSUMPTION FOR THE ENTERPRISE NETWORK.

25% average load per node
Green (Hermes) : less volume transported (compared to first/nooff) but more energy efficient

Contributions and Perspectives

- Energy aspects change the way we design applications, protocols, services and policies (i.e. load balancing is not always the best solution)
- Challenge : design energy proportional equipments and frameworks (network usage)
- HERMES : Complete and energy-efficient bandwidth reservation framework for data transfers including scheduling, prediction and on/off algorithms – Hermes focused on dedicated networks
- Validation of HERMES through simulations
- Perspective : dealing with elephants/mice and dogs

Questions ?

Laurent.lefevre@inria.fr

INVENTEURS DU MONDE NUMÉRIQUE