DSL-Lab: a Low-power Lightweight Platform to Experiment on Domestic Broadband Internet

Gilles Fedak, Jean-Patrick Gelas, Thomas Herault, Victor Iniesta, Derrick Kondo, Laurent Lefèvre, Paul Malécot, Lucas Nussbaum, Ala Rezmerita and Olivier Richard

INRIA, Université de Lyon, Université Paris XI,
Université Joseph Fourier Grenoble, LORIA / Nancy-Université

Work supported by ANR under grant JC05_55975

Introduction

Development of Internet access at home over the last decade :

- High speed broadband (xDSL, cable)
  20 to 100 Mbps in many countries
- Led to the development of many services and applications
  IPTV, VOIP, P2P, desktop computing, network games

However, those network have special characteristics :

- Performance
  - Asymetric download / upload bandwidth
  - Rather high latency
  - Performance depending on length and quality of line
- Often shared medium
  - Through different kind of networks : WiFi, Ethernet
  - NAT, firewall

Experimenting on the broadband Internet

Need to experiment on the broadband Internet :

- Understand performance characteristics
- Build better applications

Existing experimental platforms :

- Grid'5000, DAS-3 :
  - HPC clusters and very fast networks (10 GbE)
  - Add an emulator (calibration ?) to study broadband Internet ?
- PlanetLab :
  - Hundred of nodes distributed on the Internet
  - Almost always hosted by universities, on high-speed networks
- SatelliteLab (PlanetLab extension) :
  - Uses nodes on broadband Internet to forward traffic
  - Application not running directly on broadband nodes

Contribution : DSL-Lab

- Experimental platform to study the broadband Internet
- 40 nodes located in users' home, using their Internet connection
- Available for researchers to run experiments
- Large variety of network connections (different ISPs, NAT, firewall)

This talk :

Description of DSL-Lab design

Report on early experiments

DSL-Lab Design

Hardware

Requirements :

- Small and low power
- Silent
- x86 architecture

Neo CI852A-4RN10 barebone

- Mini-ITX
- Intel Celeron M 1 GHz
- 512 MB RAM
- 2 GB SSD drive
- No fan
- 4 Gb Ethernet ports

Power consumption

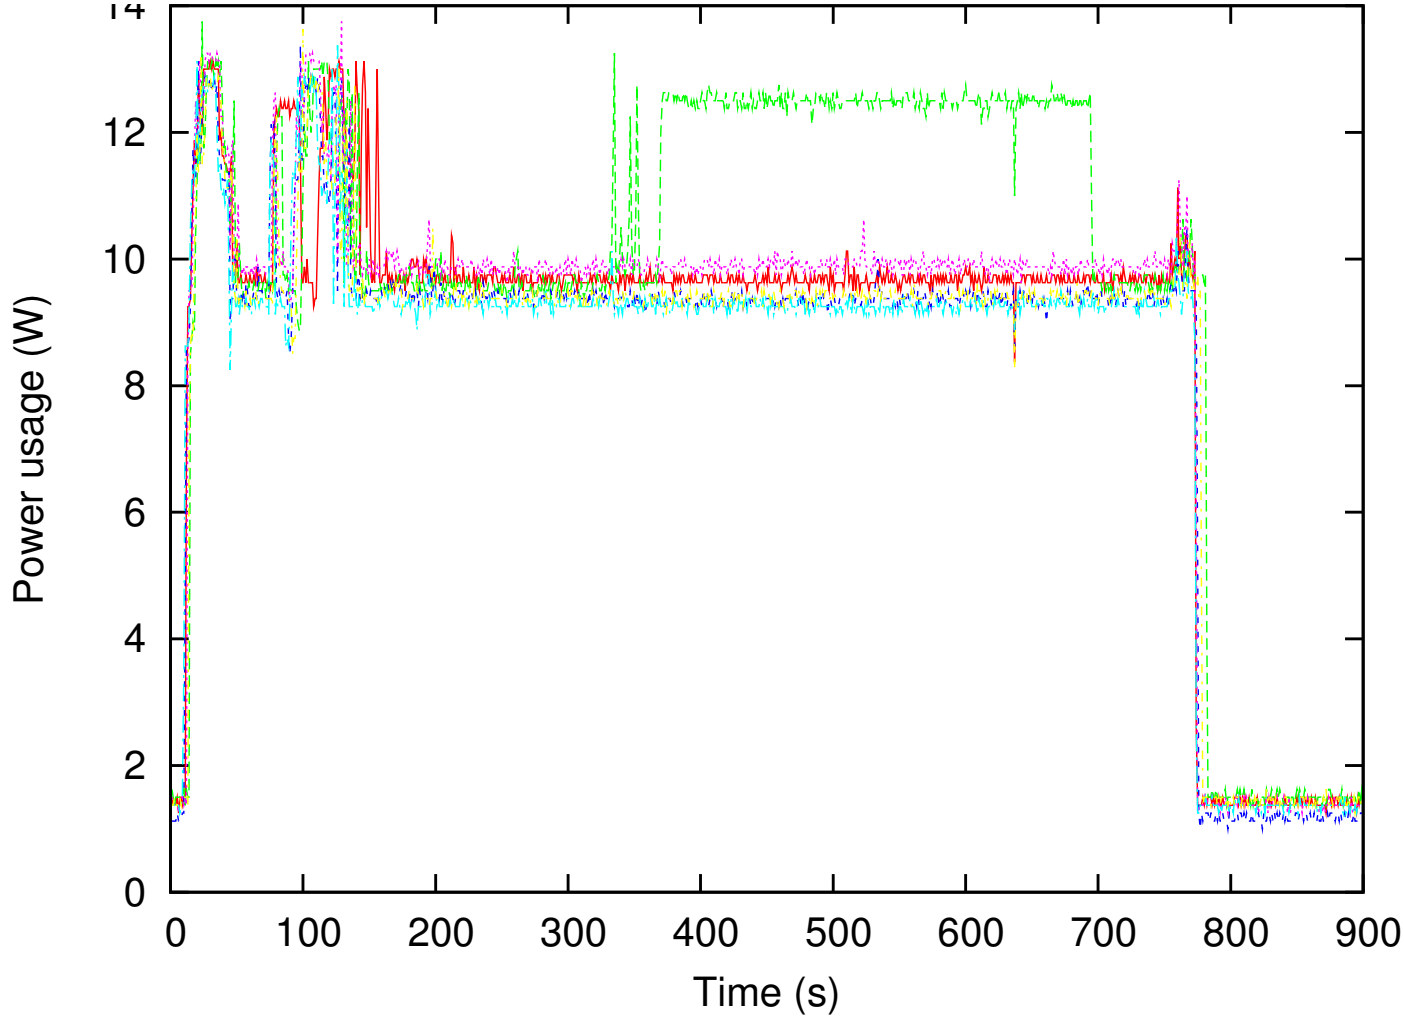

6 nodes booting, doing CPU-intensive tasks and stopping

Off : 1.5 W ; Idle : 10 W ; Full load : 13 W

Remote OS deployment

- Nodes installed in researchers' homes
- No way to access them physically
- Need to repair or upgrade the system remotely
- Virtualization : would require additional resources ; performance

Remote deployment implemented using two disk partitions :

- Small partition (5 MB) with read-only minimal system
- Larger partition for the experimentation system
- Bootloader alternatively boots on each partition
  The experimentation system can be re-installed from the first partition

Connectivity

- Most nodes hosted behind firewalls / NAT
- Not accessible from the Internet
- Port forwarding on the router : too hard to configure

Solution :

- Reverse SSH tunnels from node to a central server
- Provides a "cluster" view of the platform

Reverse SSH tunnel

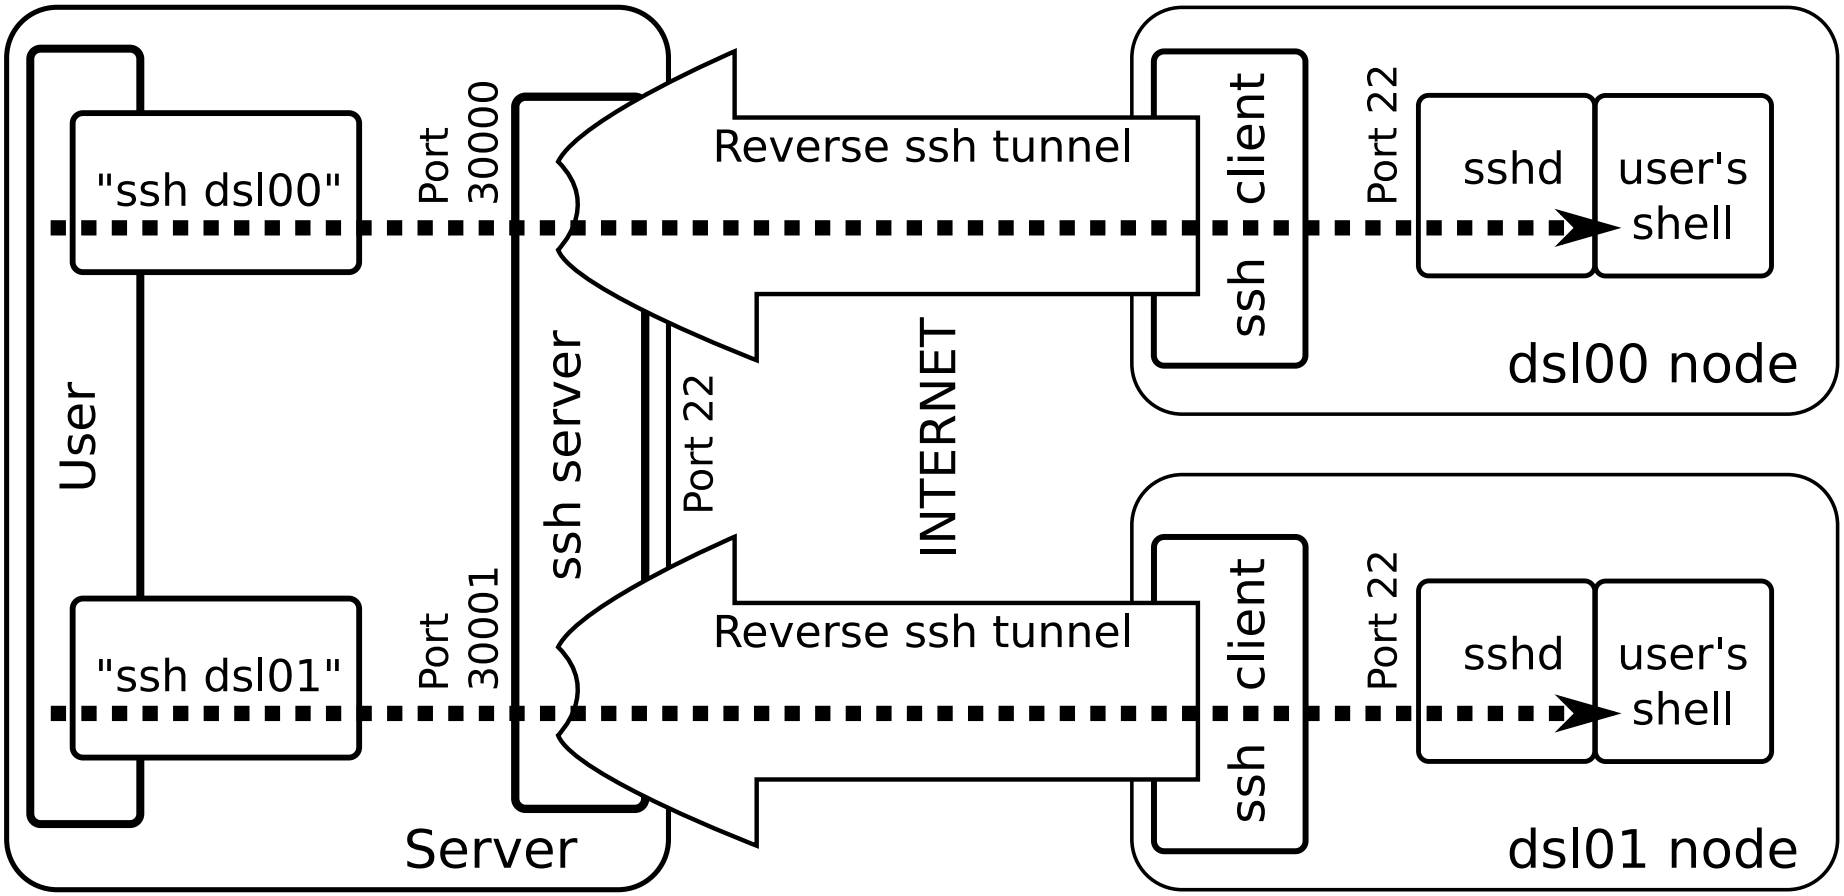

Nodes reservation

- Uses the OAR batch scheduler (also used on Grid'5000)
- Power savings :
  - Node wakes up once a day (ACPI alarm)
  - Checks if tasks submitted
  - Resets the alarm accordingly, and turns off

Early Experiments

Latency

(To central DSL-Lab server)

City	ISP	DSL latency	Tot. latency
Chaville	Free (v.5)	22.18 ms	25.14 ms
Lyon	SFR	37.75	46.47
Versailles	Free (v.5)	4.07	7.14
Lyon	Free	33.5	42.38
Grenoble	Free (v.5)	44.61	54.92
Le Touvet	Free (v.4)	46.85	57.76
Paris	Free (v.5)	33.63	37.23
Saint Gratien	Free	37.86	41.85
Bezons	Free	14.27	17.05
Lyon	SFR	27.29	36.47

Bandwidth

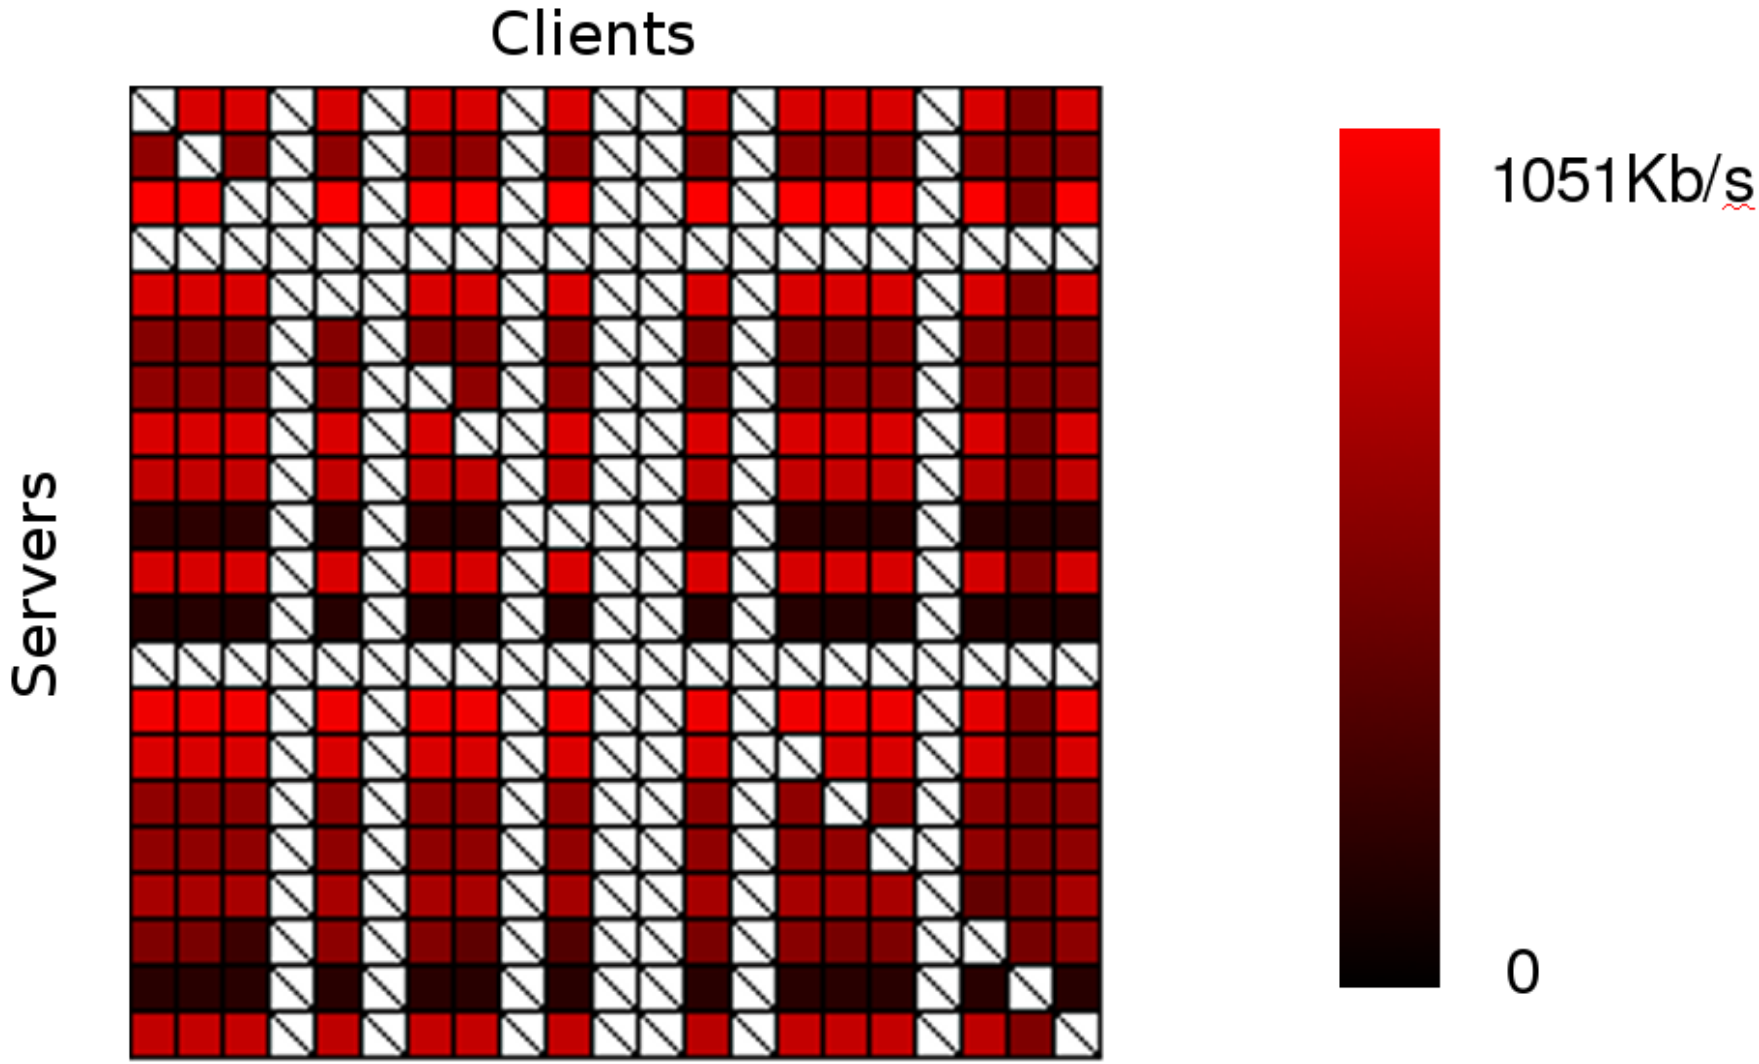

Direct communications between nodes

Private Virtual Cluster (PVC) :

- Transparent execution environment for existing applications
- Enables connecting through NAT and firewalls
  Techniques : UPnP, TCP Hole Punching, Traversing TCP, Proxy

ISP	Nodes	Connection establishment method
Free v4	5	*UPnP* or *Traversing-TCP*
Free v5	16	*UPnP* or Proxy
SFR	4	*UPnP* or *Traversing-TCP*
Orange	2	*UPnP* or *Traversing-TCP*
Comcast	1	*Traversing-TCP*
Noos	1	*UPnP* or *Traversing-TCP*
Lab.	2	Proxy

DSL-Lab used to :

- Validate PVC design and identify shortcomings
- Run experiments (NAS PB, etc) on DSL-Lab with PVC

Desktop grid platforms - BitDew

- BitDew : middleware for data management on desktop grids
  Includes both P2P and GridFTP-like protocols
- Validated on DSL-Lab
- Scenario 1 : data resiliency in presence of host failures
- Scenario 2 : All-to-all data transfers

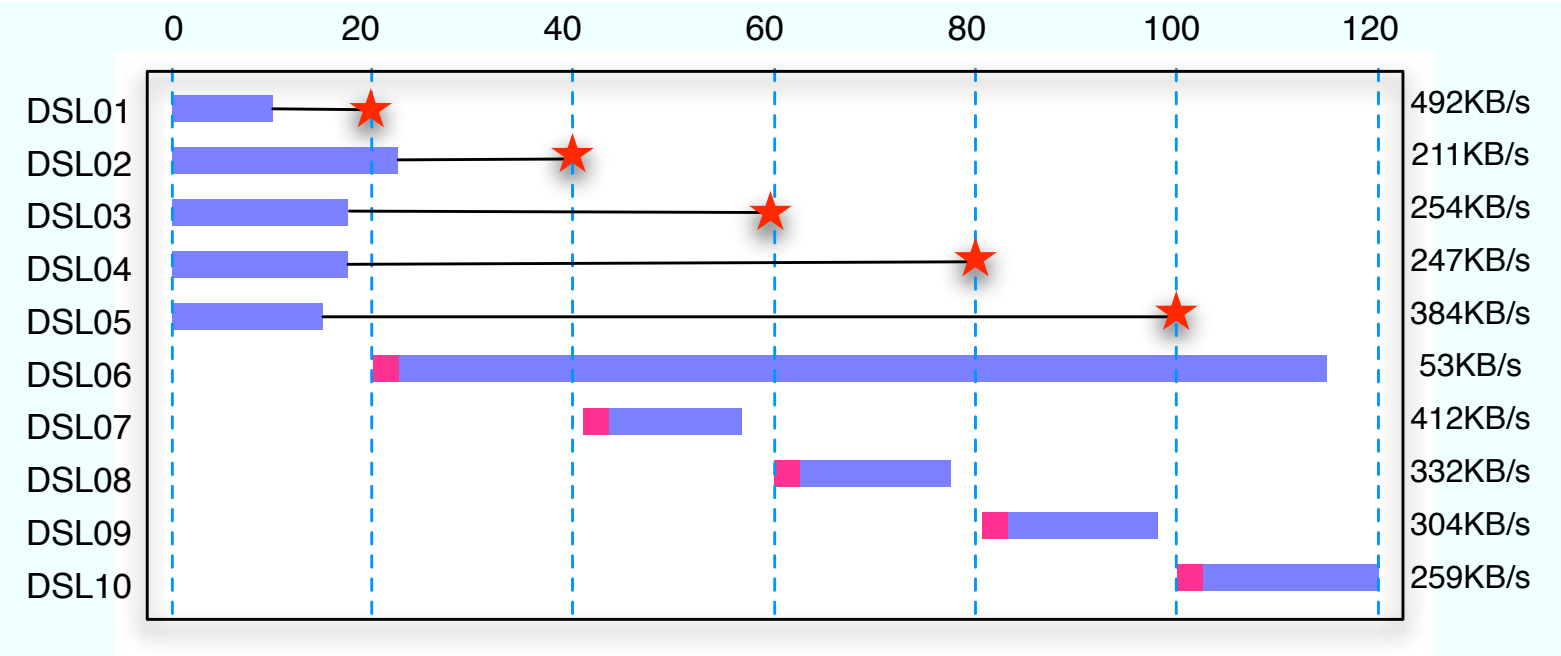

Bridging experimental platforms together

- IDHAL experiment : connecting several experimental platforms
  DSL-Lab, Grid'5000, PlanetLab, desktop computing nodes, clusters in Luxembourg and Porto Alegre, virtual machines
- Challenges :
  - Connect nodes from different platforms together
    Despite firewalls, NATs, private networks, ...
  - Make them compute together efficiently
    N-Queens problem
- Two support tools used :
  - TakTuk : adaptative deployment tool
  - KAAPI : Adaptative, Asynchronous, Parallel programming
    Work-stealing capabilities to use both slower and faster nodes

Bridging experimental platforms together

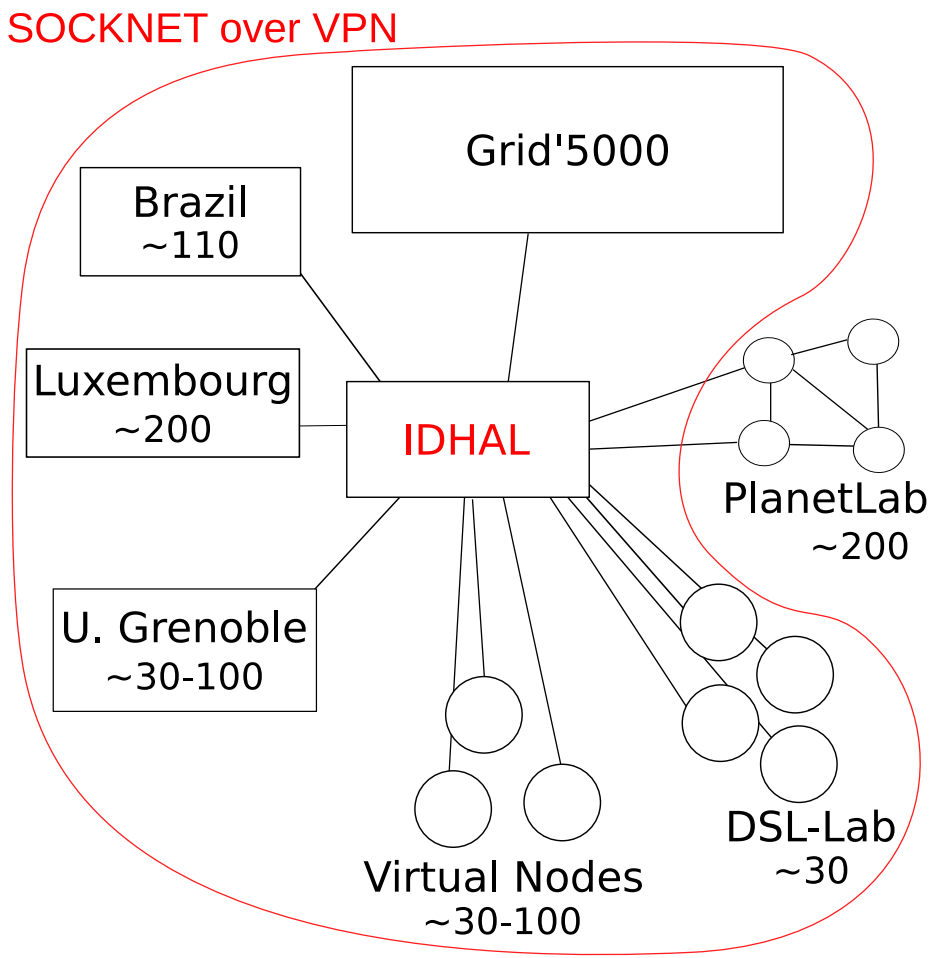

Computed several large N-Queens instances

Conclusions

DSL-Lab : platform to experiment on the broadband DSL Internet

40 nodes located on the edge of the Internet, available to researchers

Design :

- Low power, low noise nodes
- Remote OS deployment
- Cluster-like connectivity
- Nodes reservation using batch scheduler

Experiments :

- Latency / Bandwidth
- Direct communication between nodes
- Evaluation of desktop grid platforms
- Connecting experimental platforms together