GreenClouds

Cees de Laat

University of Amsterdam

Towards Hybrid Networking!

- Costs of photonic equipment 10% of switching 10 % of full routing
  - for same throughput!
  - Photonic vs Optical (optical used for SONET, etc, 10-50 k$/port)
  - DWDM lasers for long reach expensive, 10-50 k$
- Bottom line: look for a hybrid architecture which serves all classes in a cost effective way
  - map A -> L3 , B -> L2 , C -> L1 and L2
- Give each packet in the network the service it needs, but no more !
- Lower == greener

L1 ≈ 2-3 k$/port

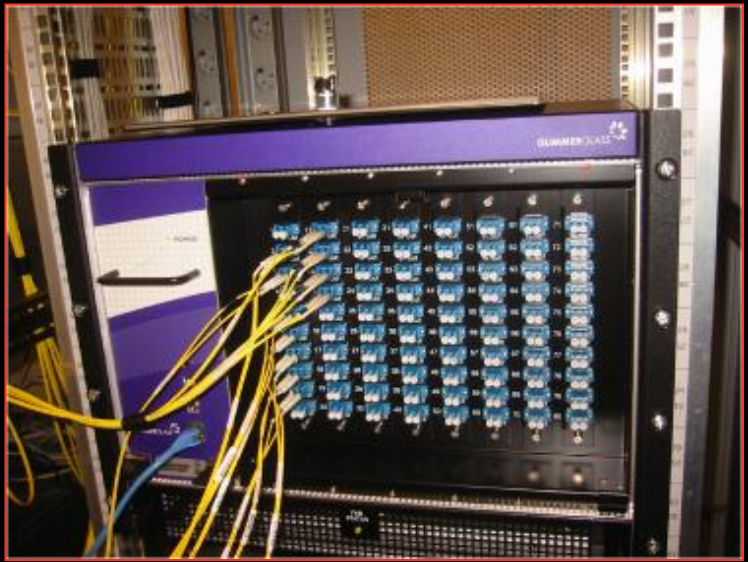

L2 ≈ 5-8 k$/port

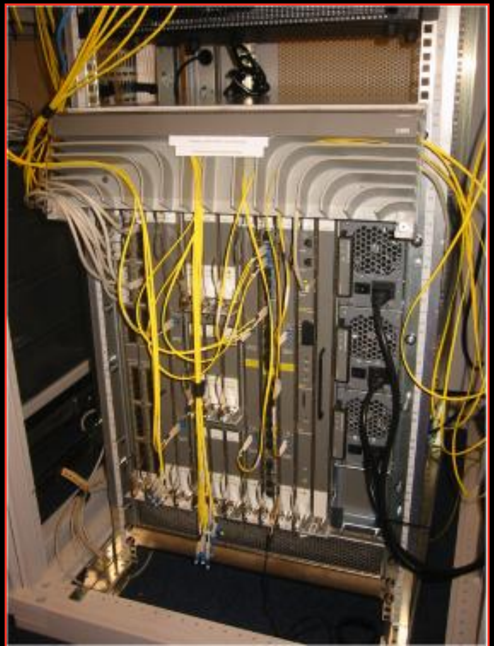

L3 ≈ 75+ k$/port

How low can you go?

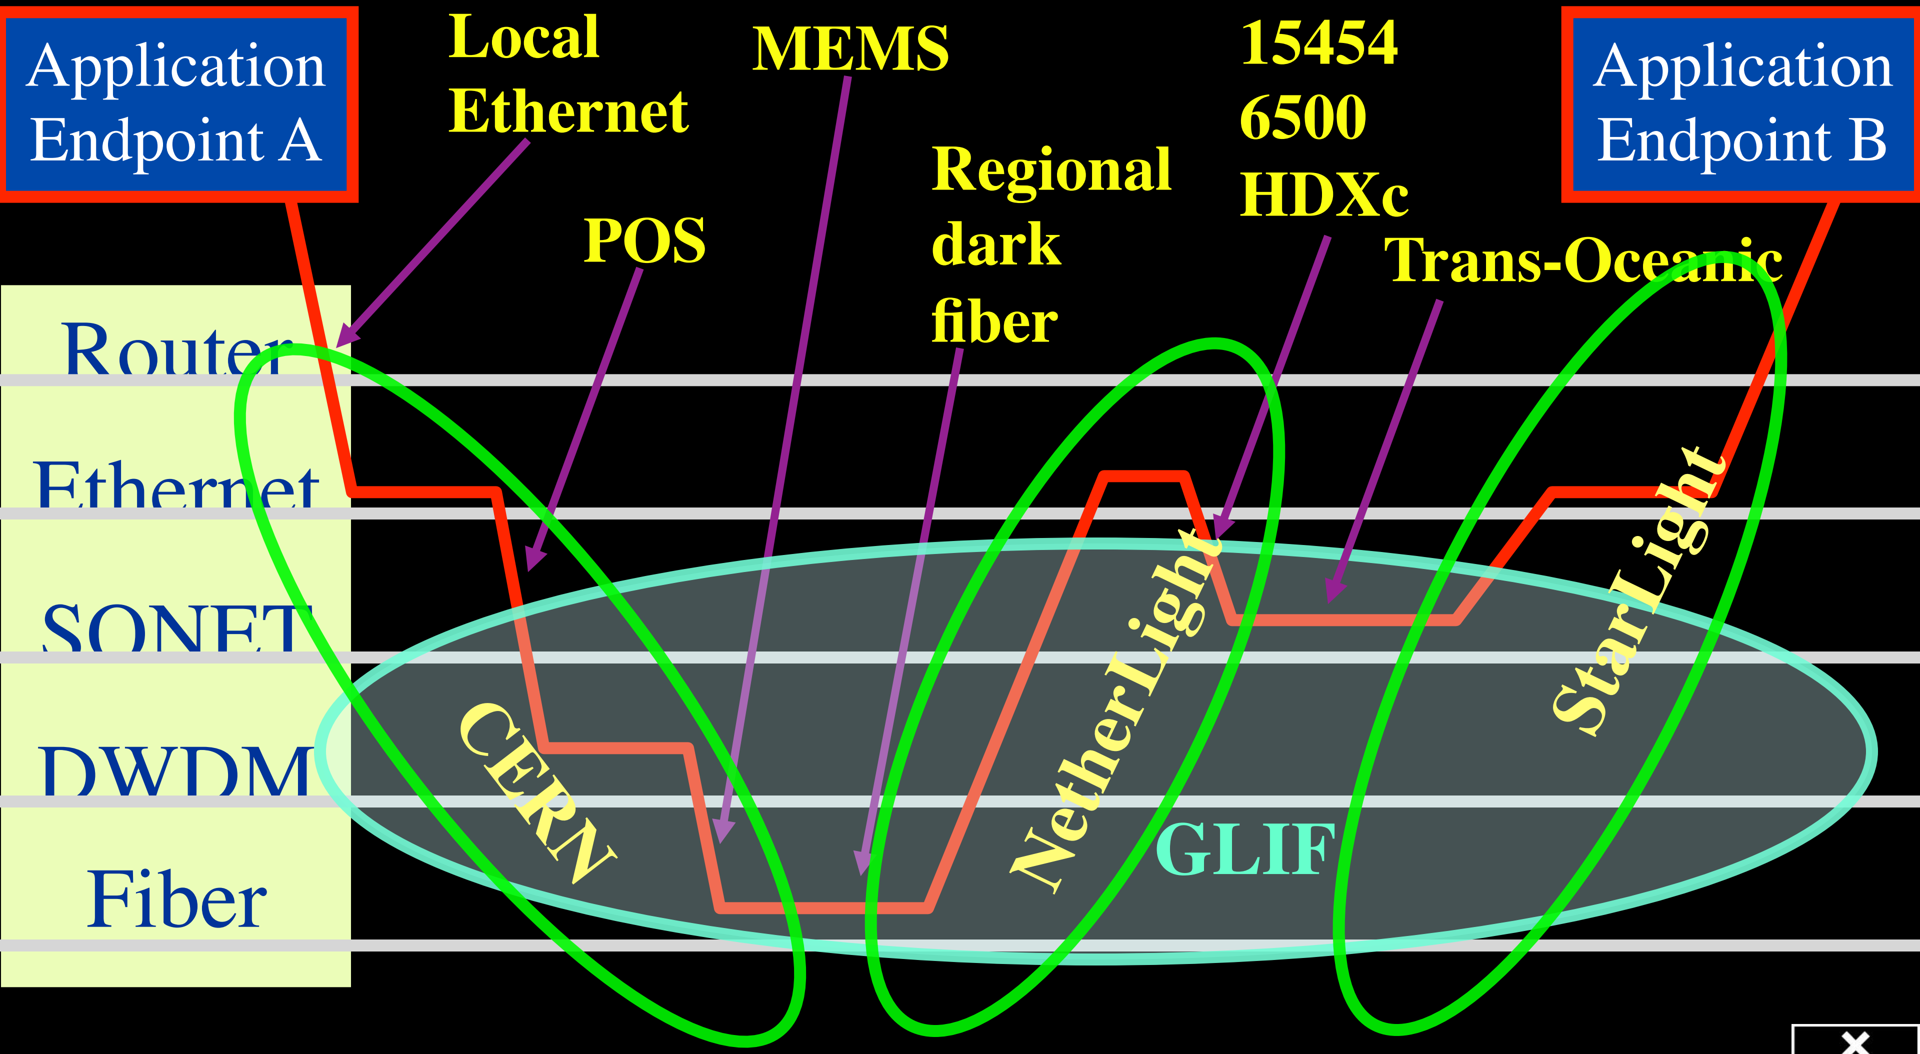

Hybrid computing

Routers	← →	Supercomputers
Ethernet switches	← →	Grid & Cloud
Photonic transport	← →	GPU's

What matters:

Energy consumption/multiplication

Energy consumption/bit transported

DAS-3 Cluster Architecture

2006!!!

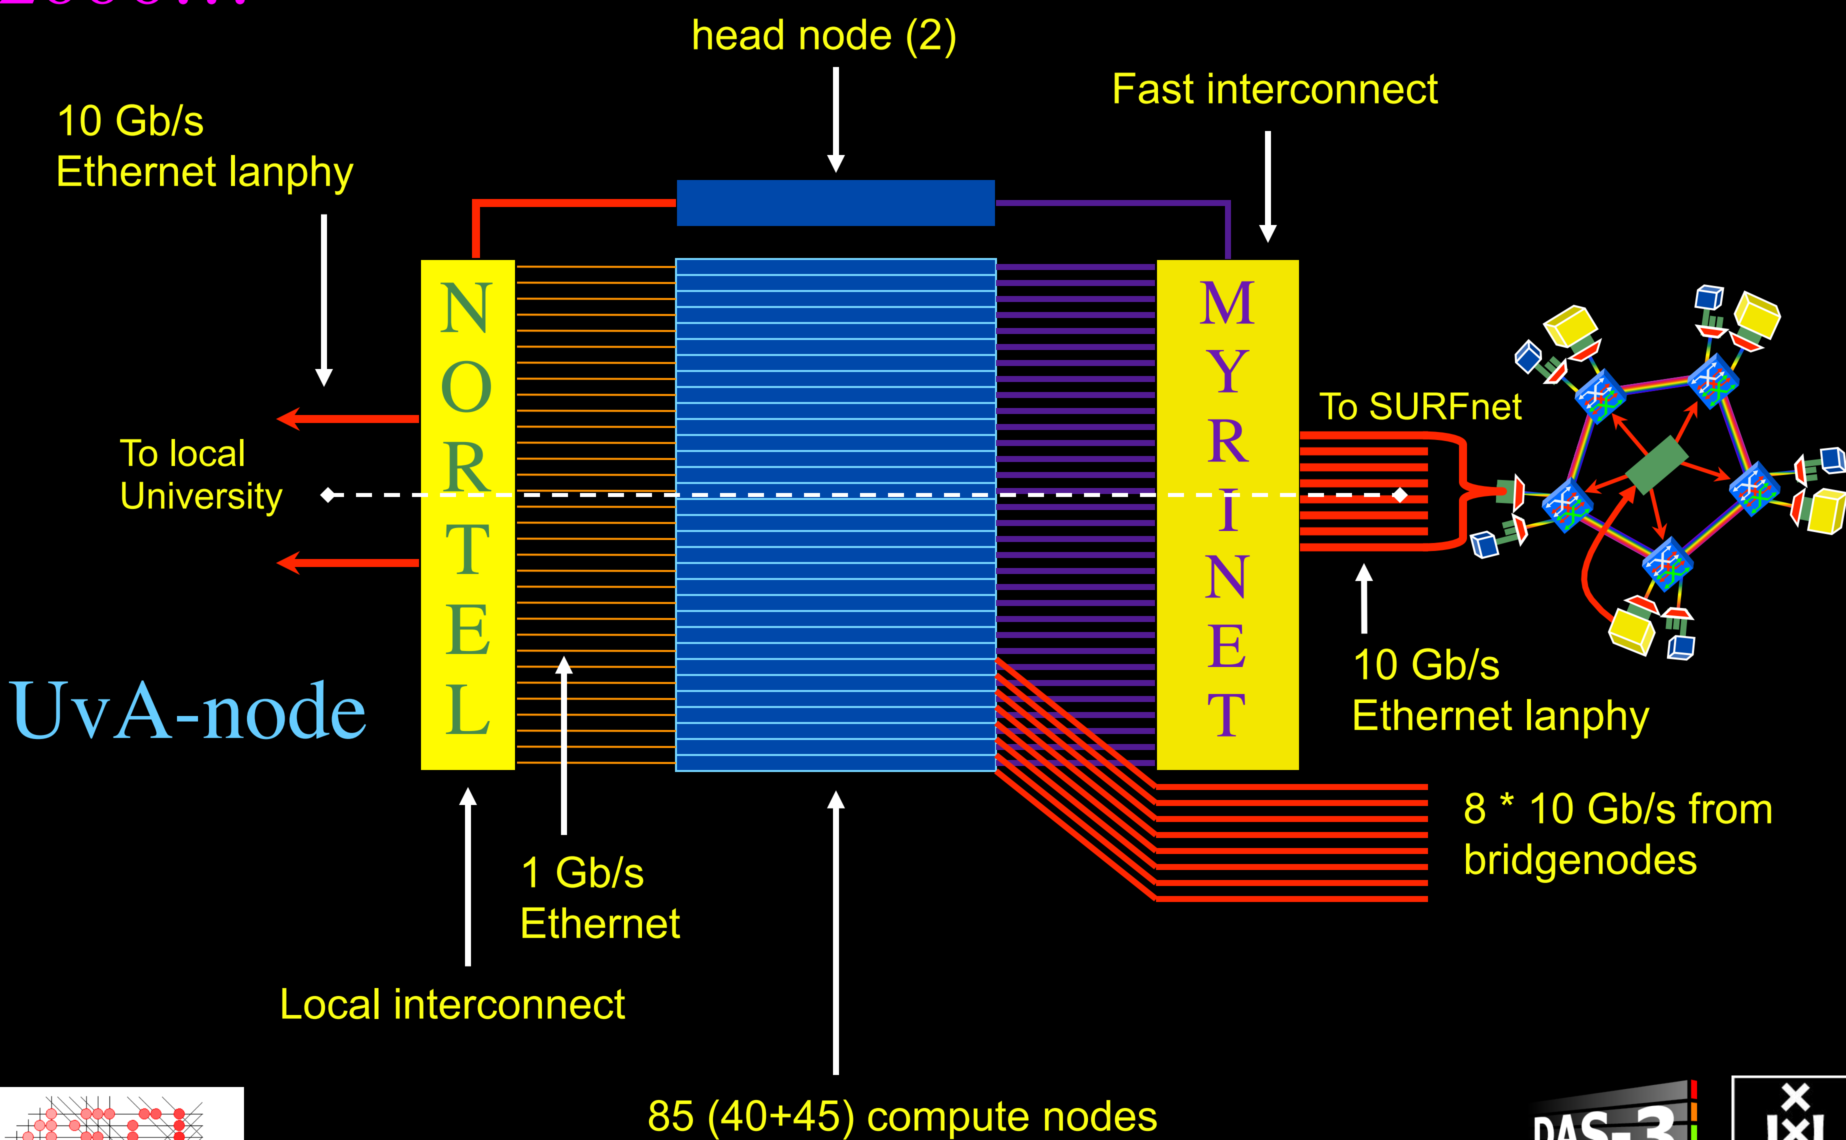

Advanced School for Computing and Imaging

2006!!!

Power is a big issue

- UvA cluster uses (max) 30 kWh
- 1 kWh ~ 0.1 €
- per year -> 26 k€/y
- add cooling 50% -> 39 k€/y
- Emergency power system -> 60 k€/y
- per rack 10 kWh is now normal
- YOU BURN HALF THE CLUSTER OVER ITS LIFETIME!

VM opportunity

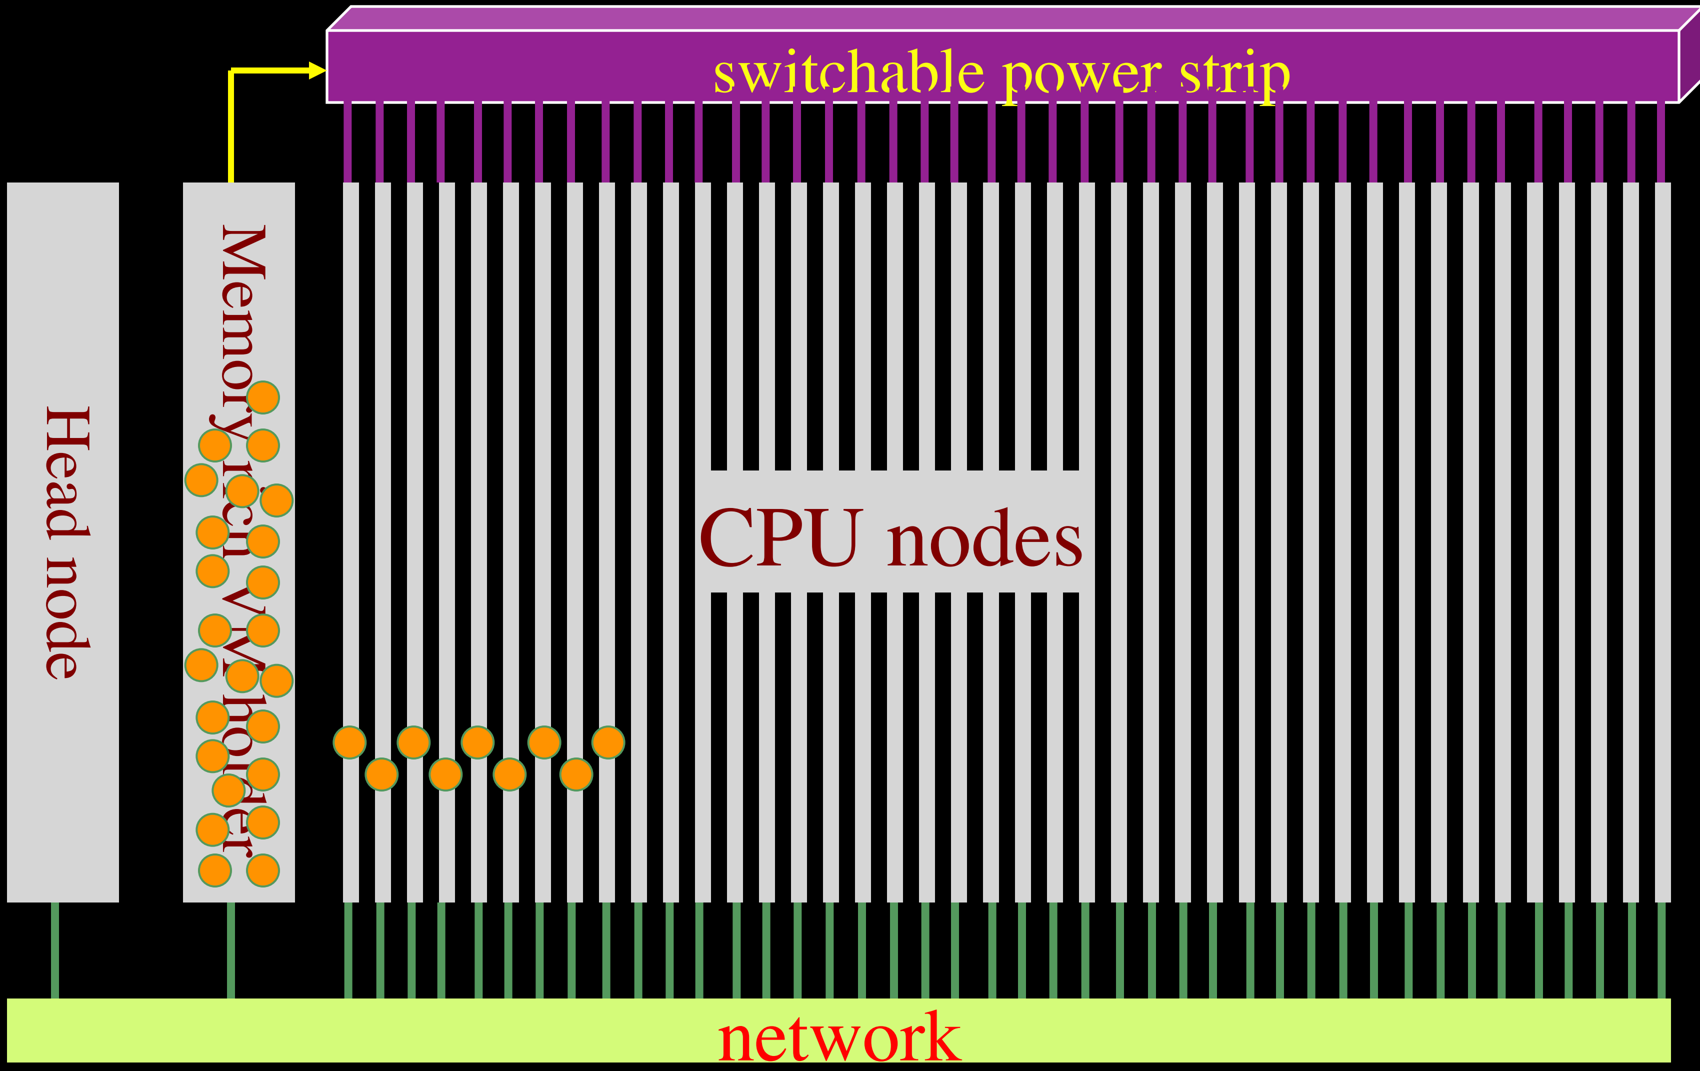

VM opportunity - B

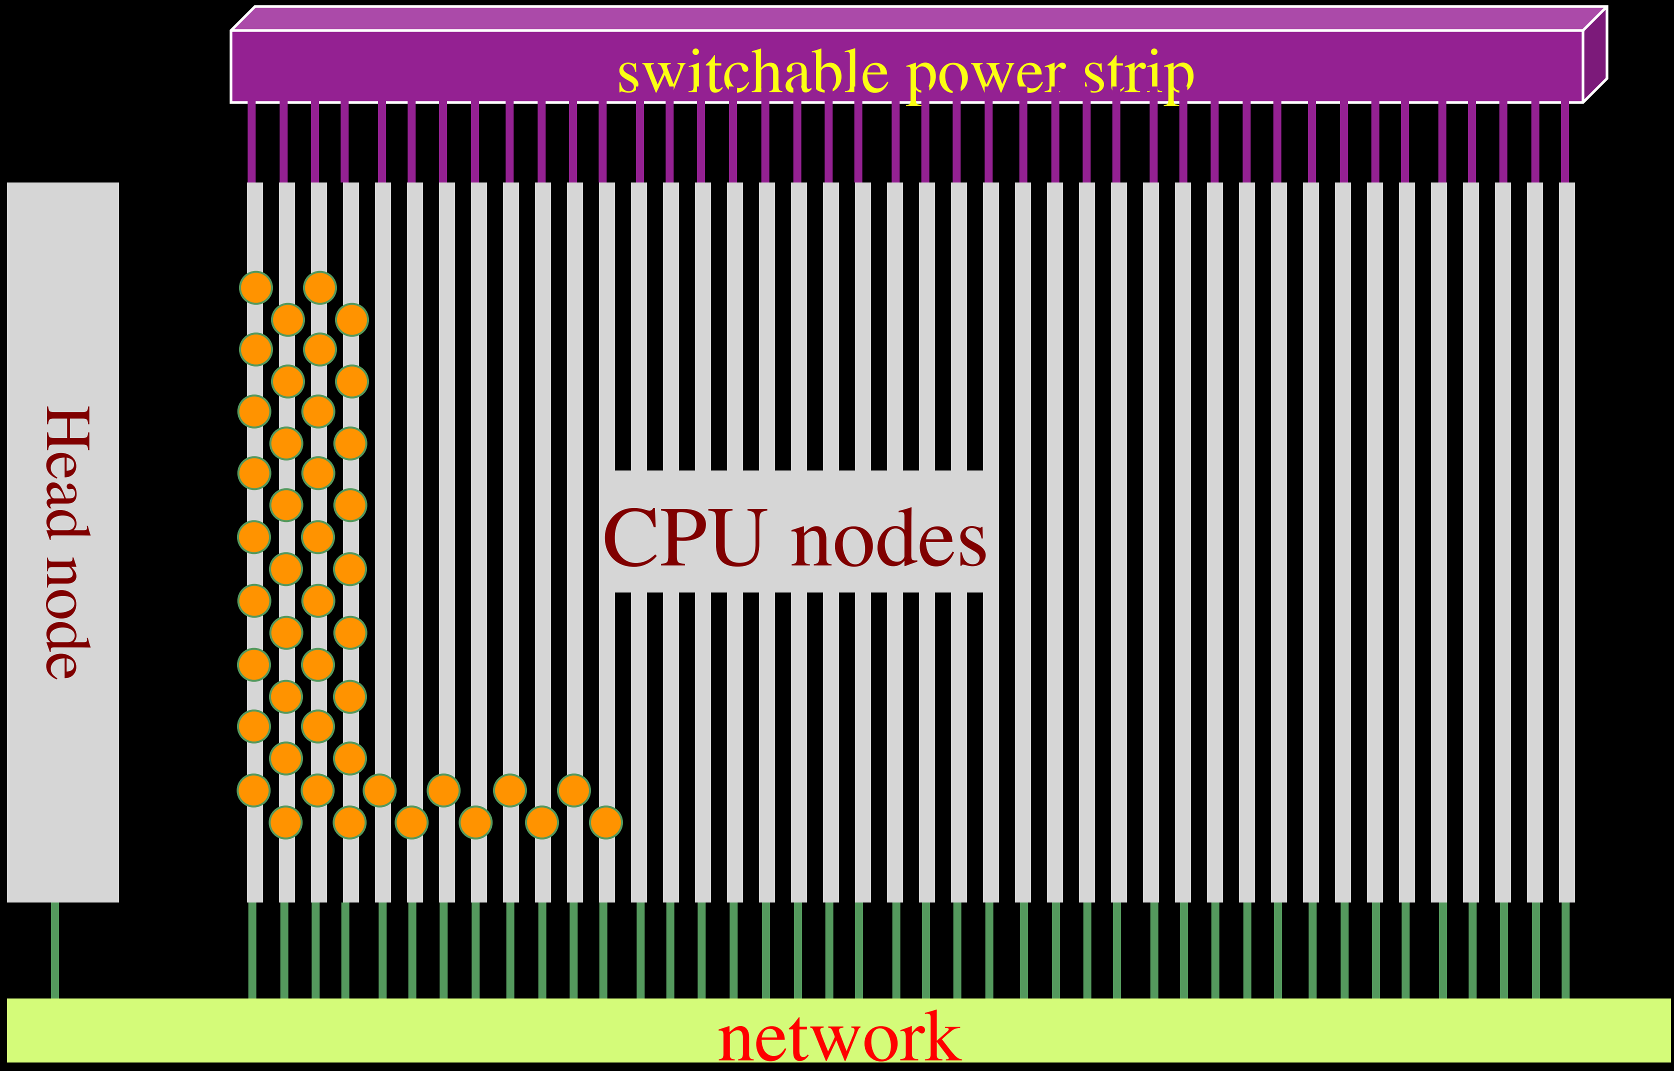

Dutch Science Foundation (NWO)

Smart Energy Systems call

- Spring 2010
- Awards in september 2010
- Start in 2011
- UvA & VU teamed up to submit GreenClouds
- http://www.nwo.nl/SES

Four focus areas

1. Smart ICT methods for energy saving, storage and generation in building environments
2. Smart control systems for flexible electricity networks (smart grids)
3. Energy reduction in processing and storing of information
4. Energy reduction in communication

Focus area 1: Energy management in buildings

- Ultimate goal is to create
  - a zero energy building (ZEB) (no energy import or export)
  - or energy neutral buildings (ENB) (net energy import and export over a year is zero)

  … without reduction in Quality of Living (QoL) and with acceptable costs
- Why interesting
  - 40 % of all energy is consumed in houses
- Consensus
  - ZEB and ENB with equal QoL is only possible <u>with ICT</u> technology

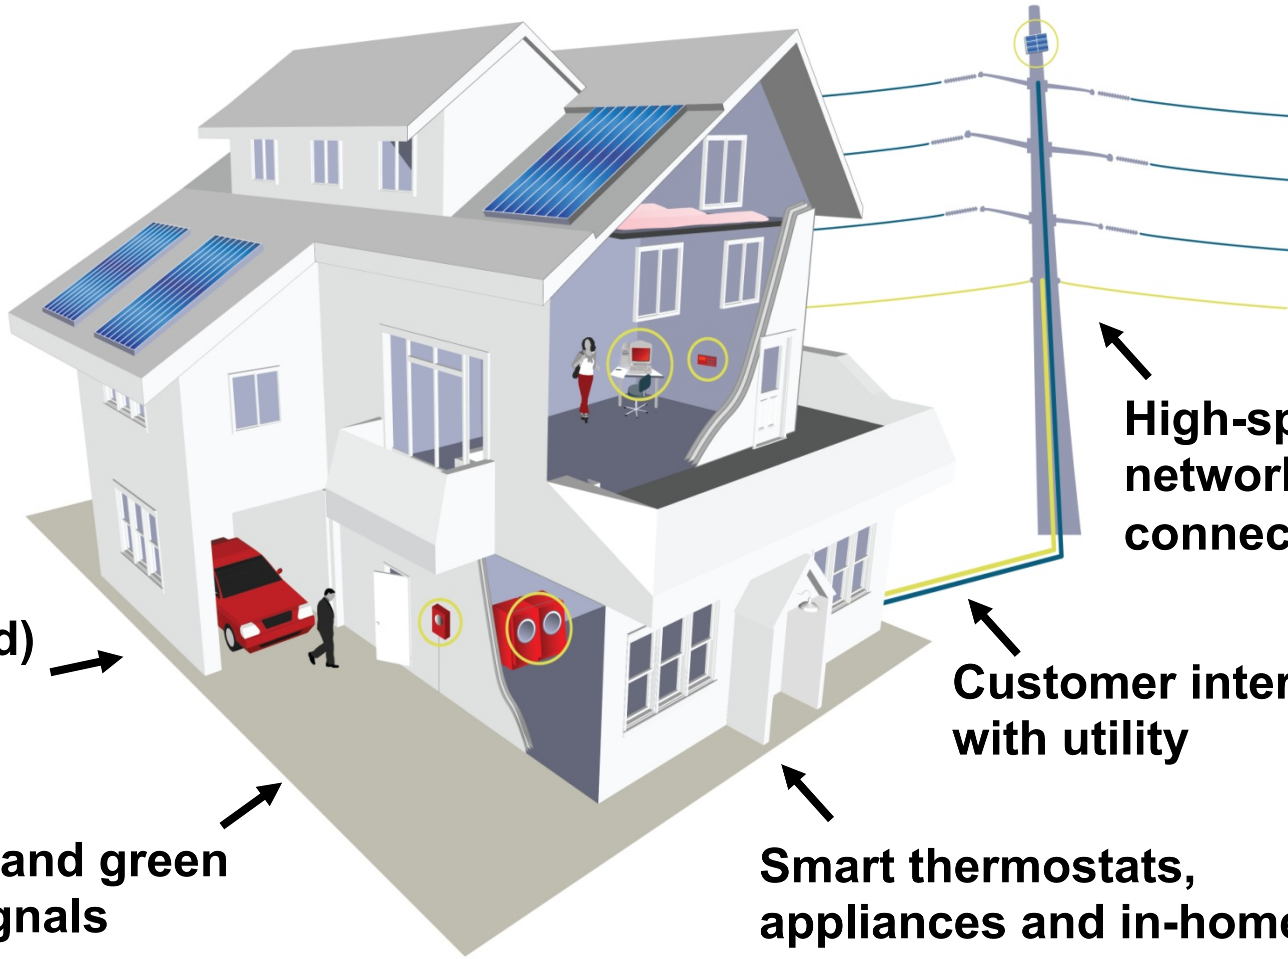
Added green power sources
Plug-in (hybrid) electric cars
Real-time and green pricing signals
High-speed, networked connections
Customer interaction with utility
Smart thermostats, appliances and in-home control devices

Smart House

ICT challenges

- Reduce stand-by power (10% of all power)
- Prediction of energy profile
- When to store electricity locally, when to export when to import
- Sensor networks to sense / predict / control energy consumption and production
- Efficient in-building communication infrastructures to control white-good & brown-good appliances & lighting and micro-generators
- Create energy-awareness of house owners
- Trade-off between energy-efficiency and Quality of Living
- ….

Focus area 2: Flexible electricity networks

- Goal
  - Create ICT technology for smart grids

Today's electricity grids and efficiency

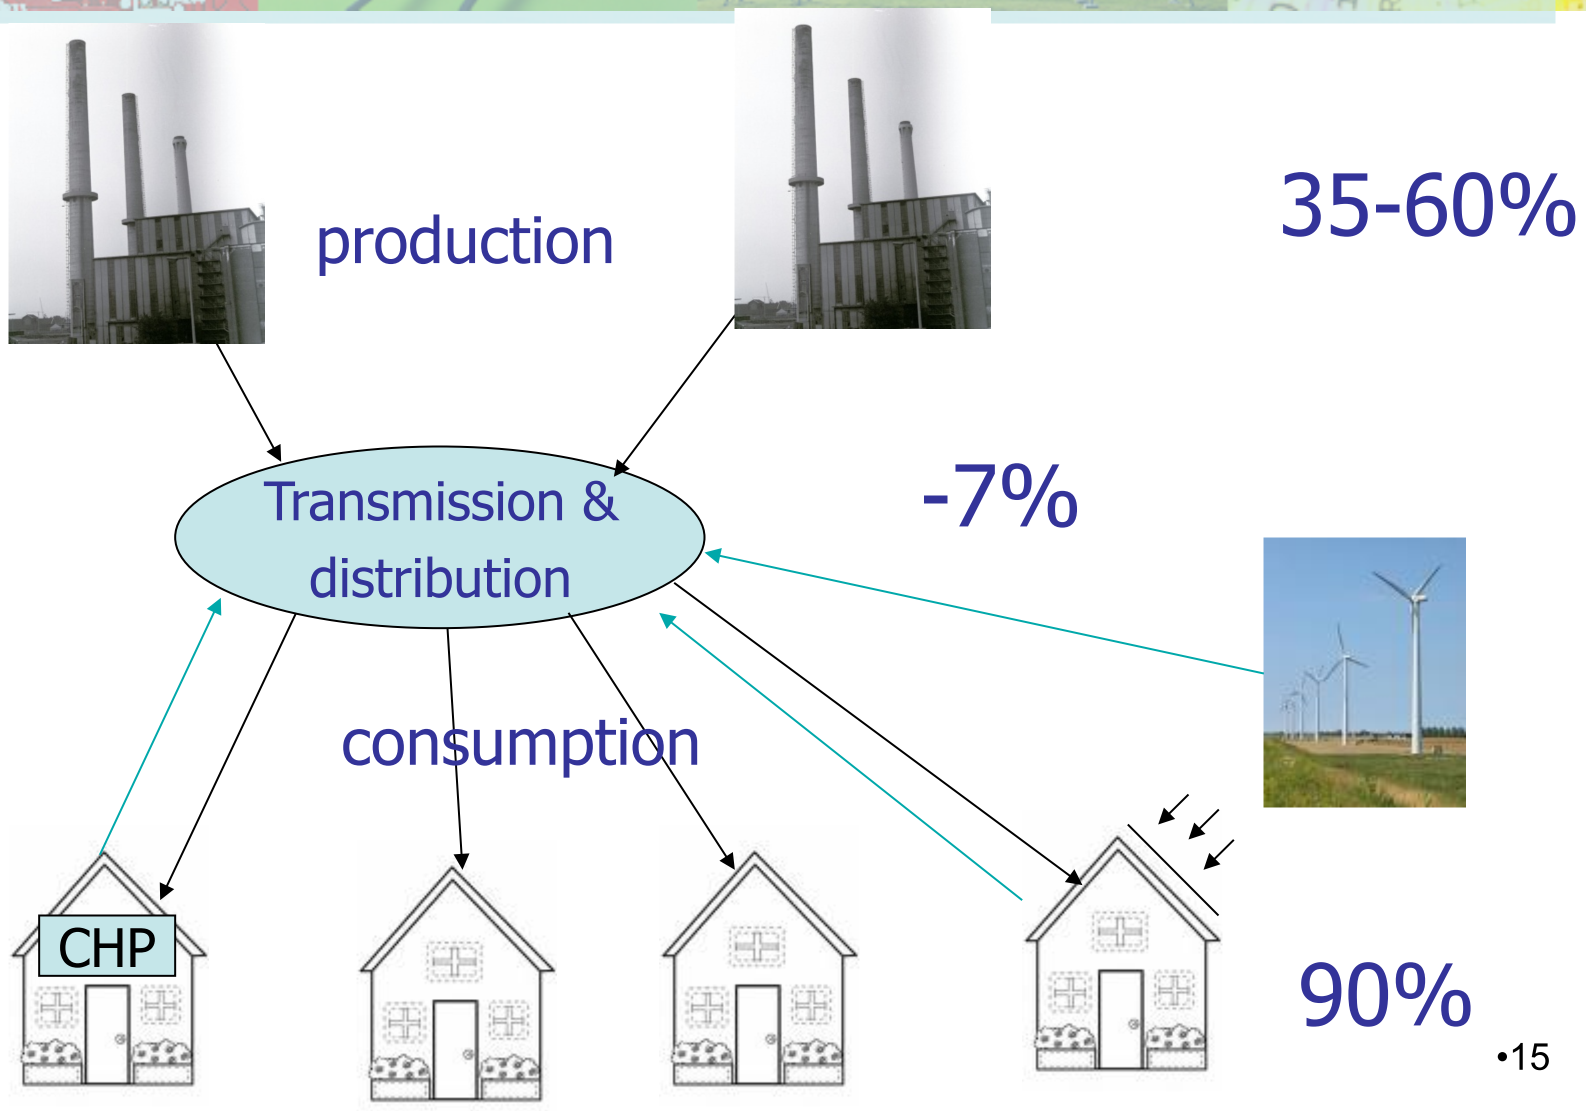

The future: smart grids

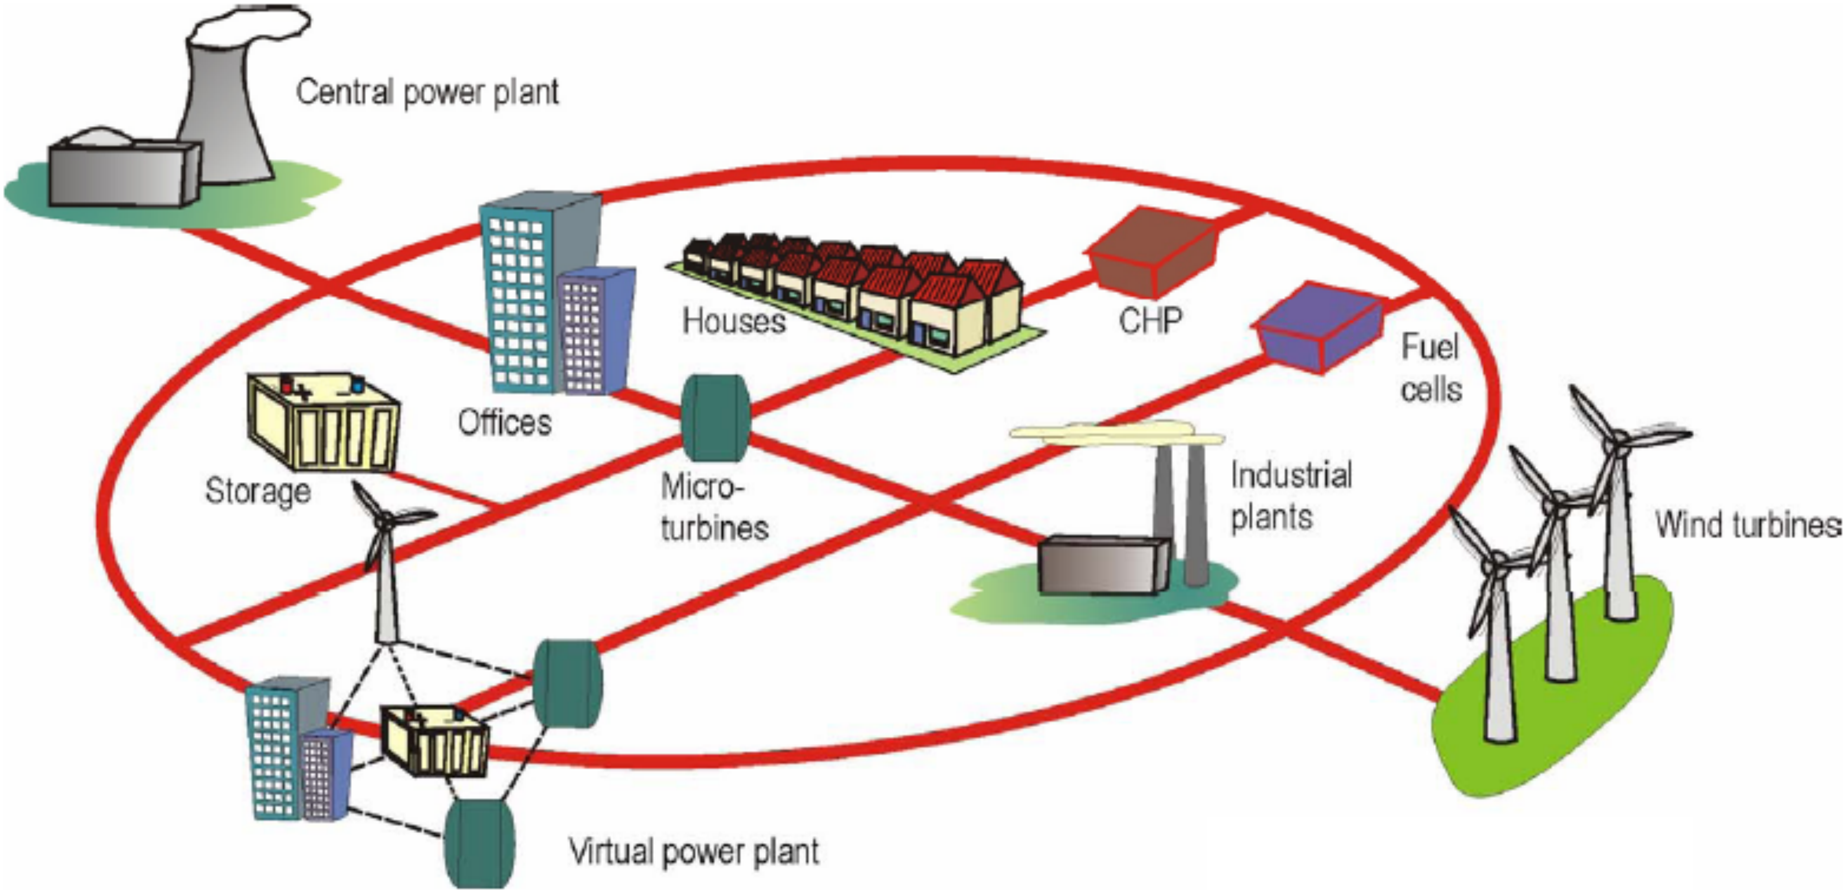

Some ICT challenges ahead

- Using ICT for efficiency implies efficient ICT
- Dependability of ICT
  - Smart grids are the life lines of our society
  - Should continue even when some parts fail
- Load balancing in the home / neighborhood
- Compensate for dynamics of generation (e.g. windmills)
- Scalability
  - Grid with thousands / millions of generators/consumers
  - Real-time control of thousands / millions of appliances
- Online optimization problems
  - Do I store energy locally or give it back to the grid?
  - Do I get energy from the battery or from the grid?

Focus area 3:
Energy reduction in processing

- Goal
  - Reduce energy consumption of ICT

Microprocessor Trends

- Single Thread performance power limited
- Multi-core throughput performance extended
- Heterogeneous extends performance and efficiency

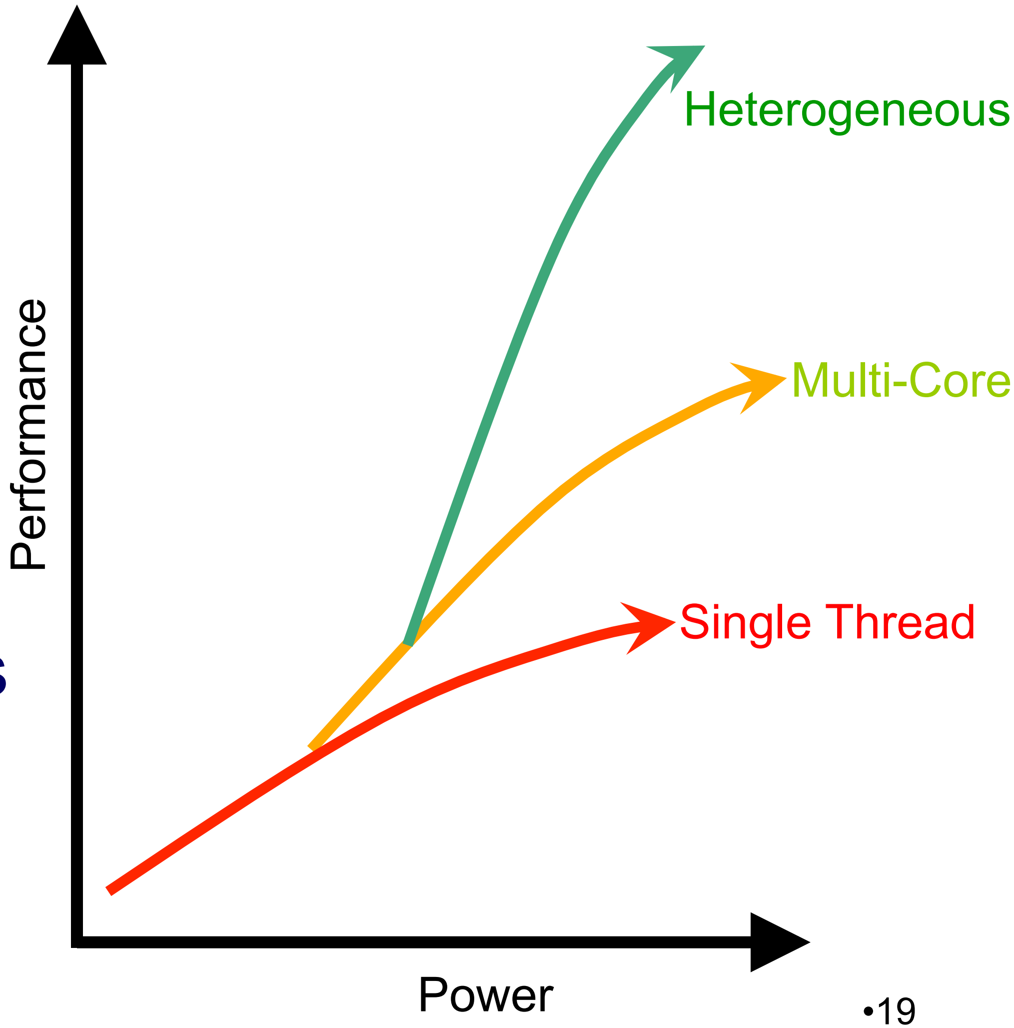

Slide from Peter Hofstee IBM

Future is in heterogeneous MPSoC Platforms

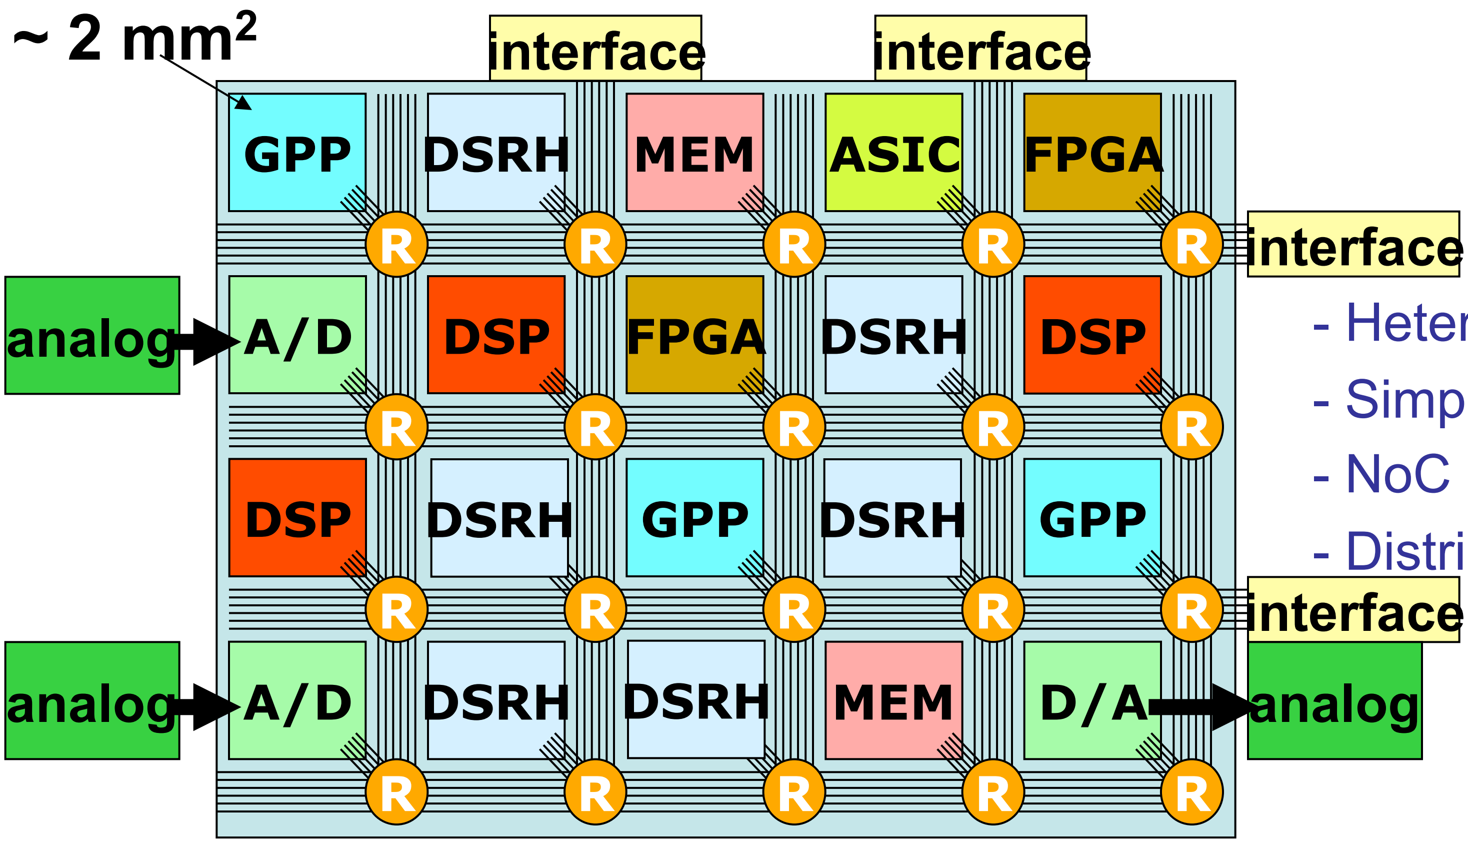

- Heterogeneous
- Simple tiles
- NoC
- Distributed Memory

ICT challenges

- Efficient processing platforms
- Efficient memory hierarchy
- Efficient software / compilers
- How to program multi-core systems
- Make applications, compilers and operating systems energy aware
- System can adapt (at run-time) to the environment
- Promote energy awareness of PC users

Focus area 4
Energy reduction in communication

- Goal
  - Energy reduction in communication by using
    - Optical communication techniques
    - Wireless communication techniques
    - Intelligent networking techniques

ICT challenges

- Optical fiber access networks
  - optical access by GPON consumes about 18x less energy per user than VDSL2
  - all-optical packet switching by avoiding power-hungry EO conversions
- Optimum combination of radio technologies with optical fiber technologies
- Low power cognitive radio transceivers
- Wideband transceivers and wake-up radios for small and adaptive cell sizes
- Low-power transceivers with strong spatial selectivity, MIMO and adaptive beamforming

Partners in GreenClouds

- Free University of Amsterdam
  - Henri Bal
- (really free) University of Amsterdam
  - Paola Grosso, Cees de Laat
- SARA
  - Axel Berg
- In context of:
  - ASCI
  - DAS4

GreenClouds @ VU & UvA

- The GreenClouds project studies how to reduce the energy footprint of modern High Performance Computing systems (like Clouds) that are distributed, elastically scalable, and contain a variety of hardware (accelerators and hybrid networks). The project takes a system-level approach and studies the problem of how to map high-performance applications onto such distributed systems, taking both performance and energy consumption into account.

- We will explore three ideas to reduce energy:
  1. Exploit the diversity of computing architectures (e.g. GPUs, multicores) to run computations on those architectures that perform them in the most energy-efficient way;
  2. Dynamically adapt the number of resources to the application needs accounting for computational and energy efficiency;
  3. Use optical and photonic networks to transport data and computations in a more energy-efficient way.

GreenClouds @ VU & UvA

- The project will create the GreenClouds Knowledge Base System (GKBS) based on semantic web technology, which will provide detailed information on the energy characteristics of various applications (e.g., obtained from previous execution runs) and the different parts of the distributed system, including the network. Also, the project will study a broad range of applications and determine which classes of applications can reduce their energy consumption using accelerators. Finally, it will study energy reductions through dynamic adaptation of computing and networking resources. The project will make extensive use of the DAS-4 infrastructure, which is a wide-area testbed for computer scientists, to be equipped with many types of accelerators, a photonic network, and energy sensors.

- The results of the project will be utilized by the SARA national HPC center that operates a supercomputer, clusters, accelerator systems, and an HPC cloud. Today, the costs of energy over the lifetime of these systems are already larger than their acquisition costs, so reducing energy is vitally important for centers like SARA. Moreover, the results will be utilized in DAS-4 itself.

DAS-4 @ UvA

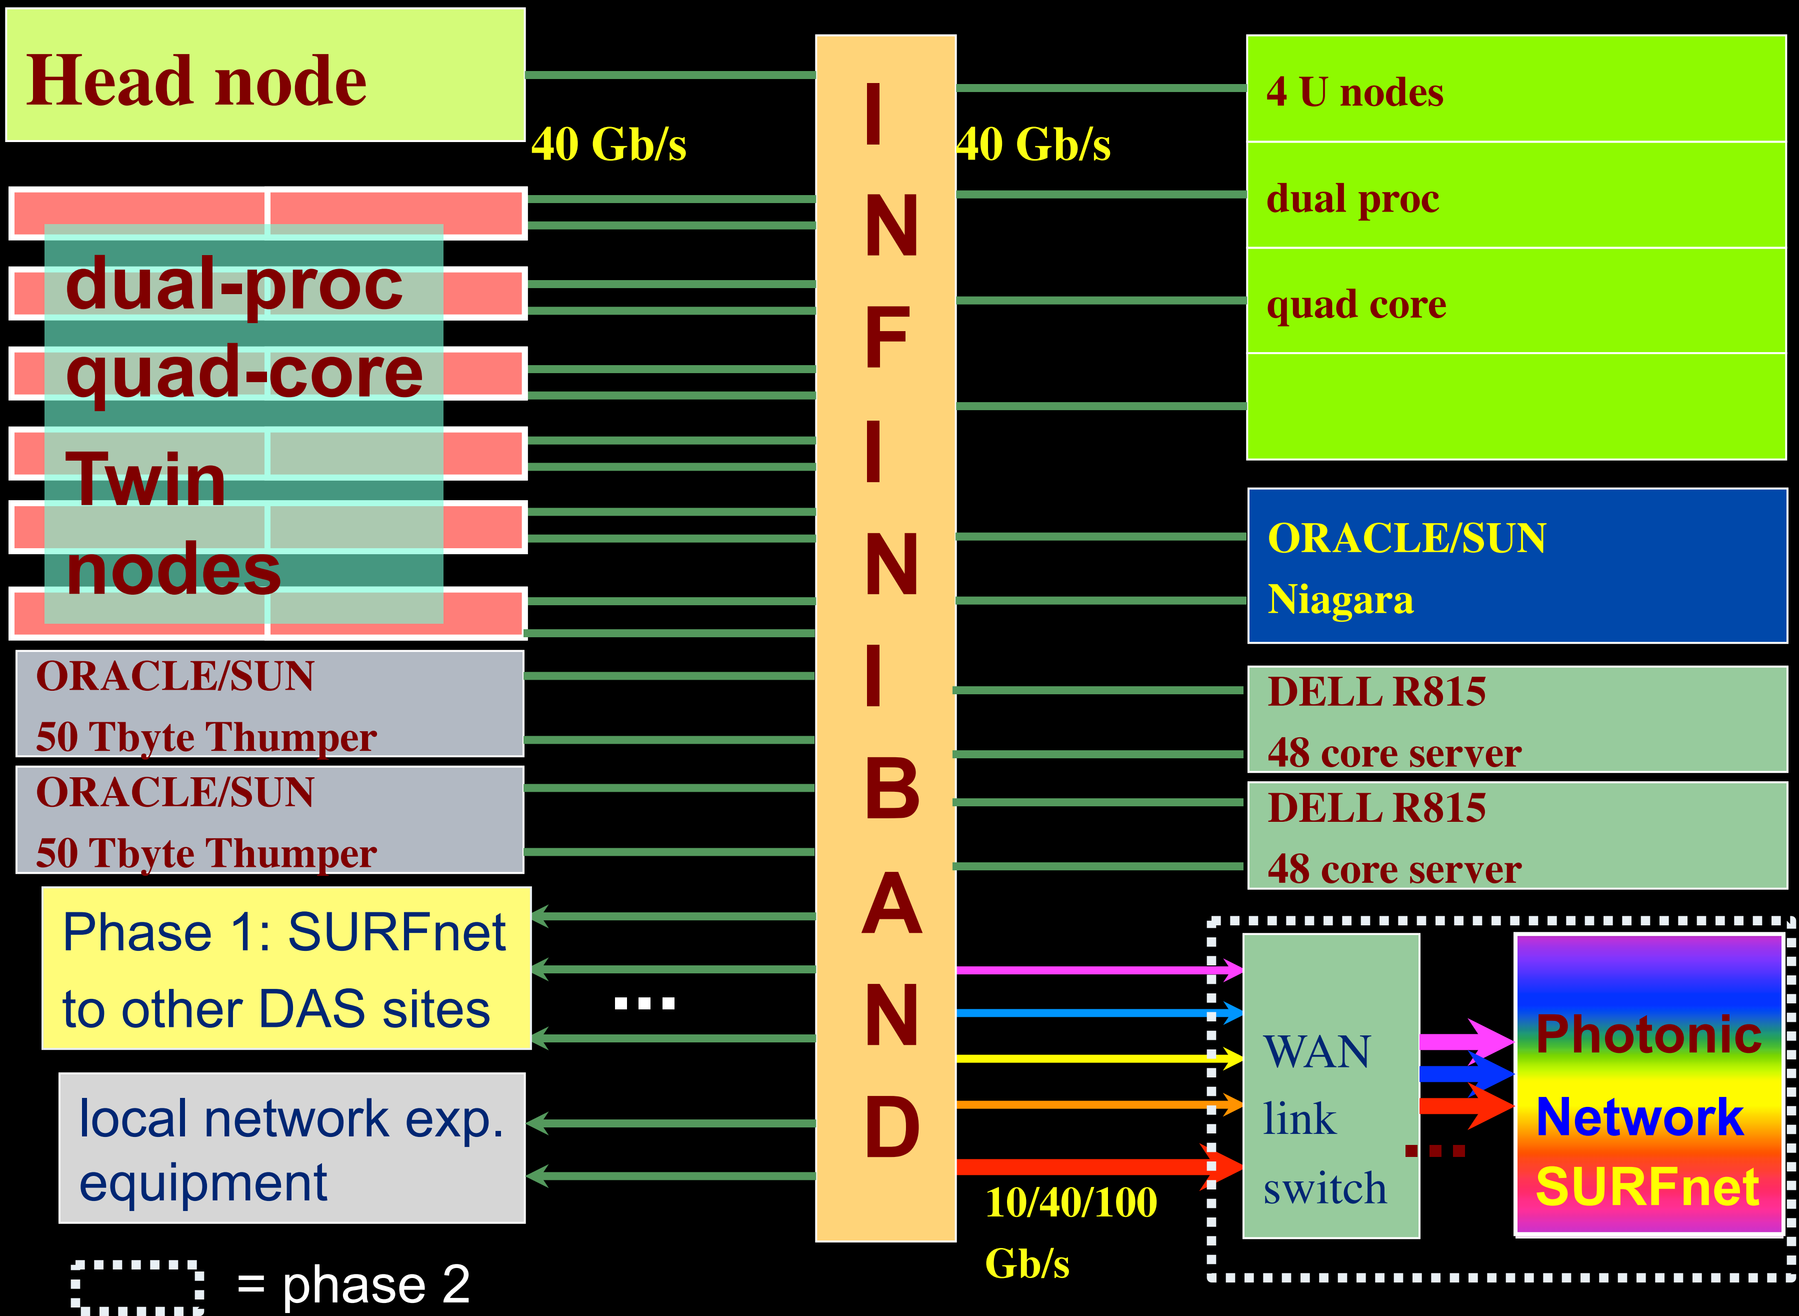

Head node
40 Gb/s
40 Gb/s
dual-proc
quad-core
Twin
nodes
ORACLE/SUN
50 Tbyte Thumper
ORACLE/SUN
50 Tbyte Thumper
Phase 1: SURFnet
to other DAS sites
...
local network exp.
equipment
= phase 2
INFINIBAND
4 U nodes
dual proc
quad core
ORACLE/SUN
Niagara
DELL R815
48 core server
DELL R815
48 core server
WAN
link
switch
...
Photonic
Network
SURFnet
10/40/100
Gb/s

CPU freq

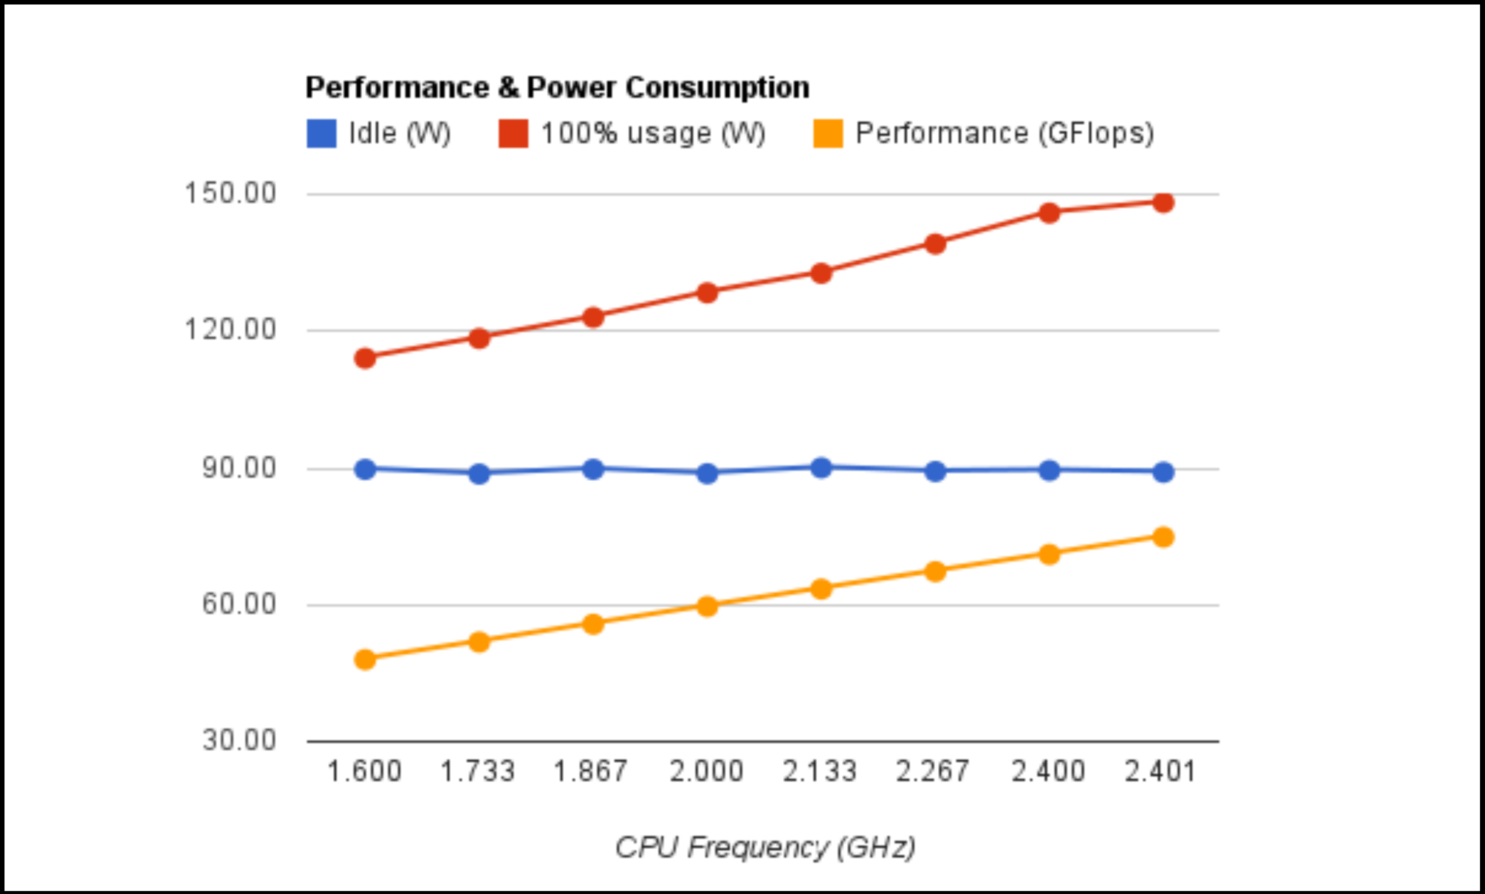

Within this benchmark, measured the power consumption & performance of the twin-node with various fixed frequencies by disabling the CPU-frequency-scaling feature and fixing the CPU frequency to one of available frequencies.

Efficiency (Multiplications/Watt)

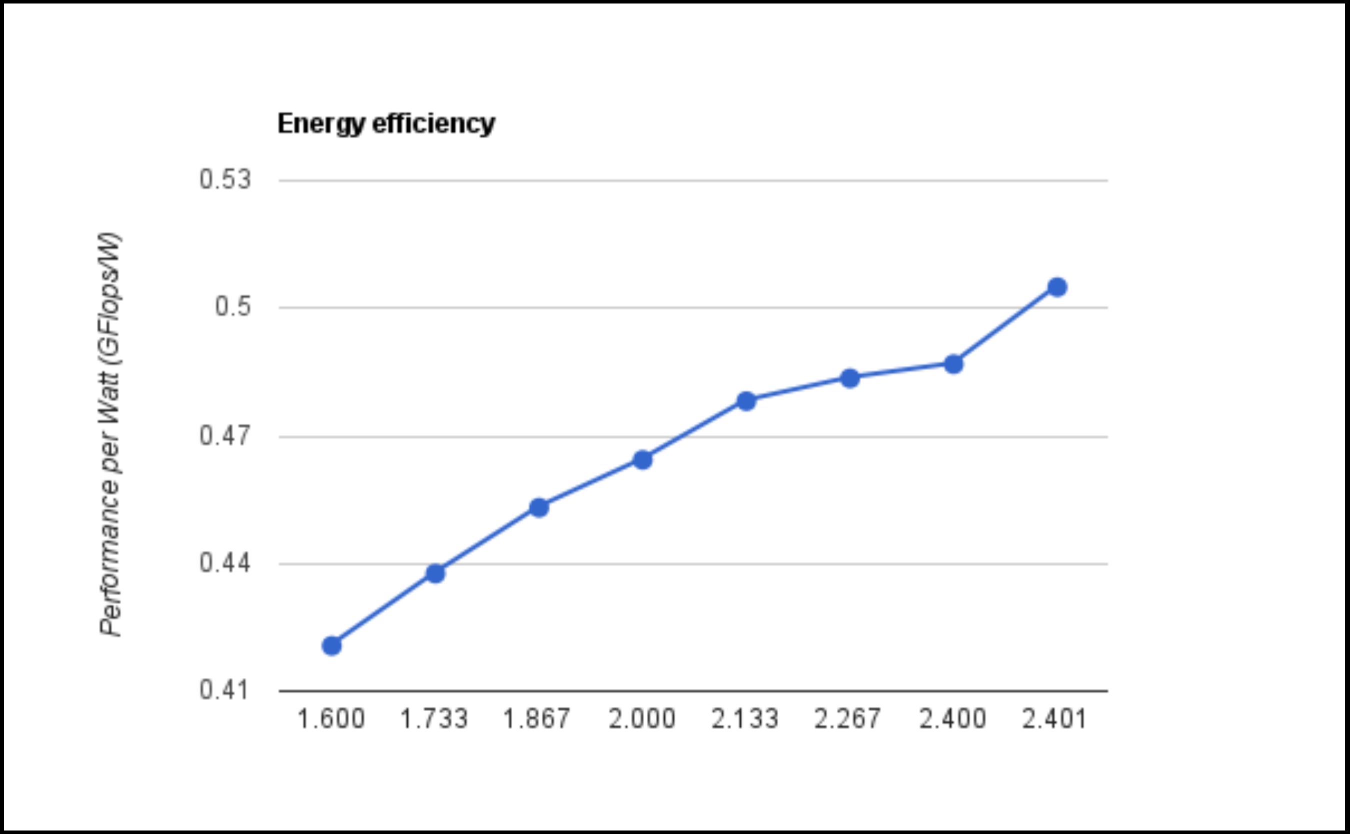

Figure 1(a) and Figure 1(b) are the performance and energy efficiency graphs for the CPU-Frequency-Scaling experiments, from where we may conclude that:

- Idle power consumption remains a constant in regardless of its CPU frequency;
- Higher CPU frequency results in better energy efficiency
- CPU consumes the most significant power within a system.

of threads 1-16

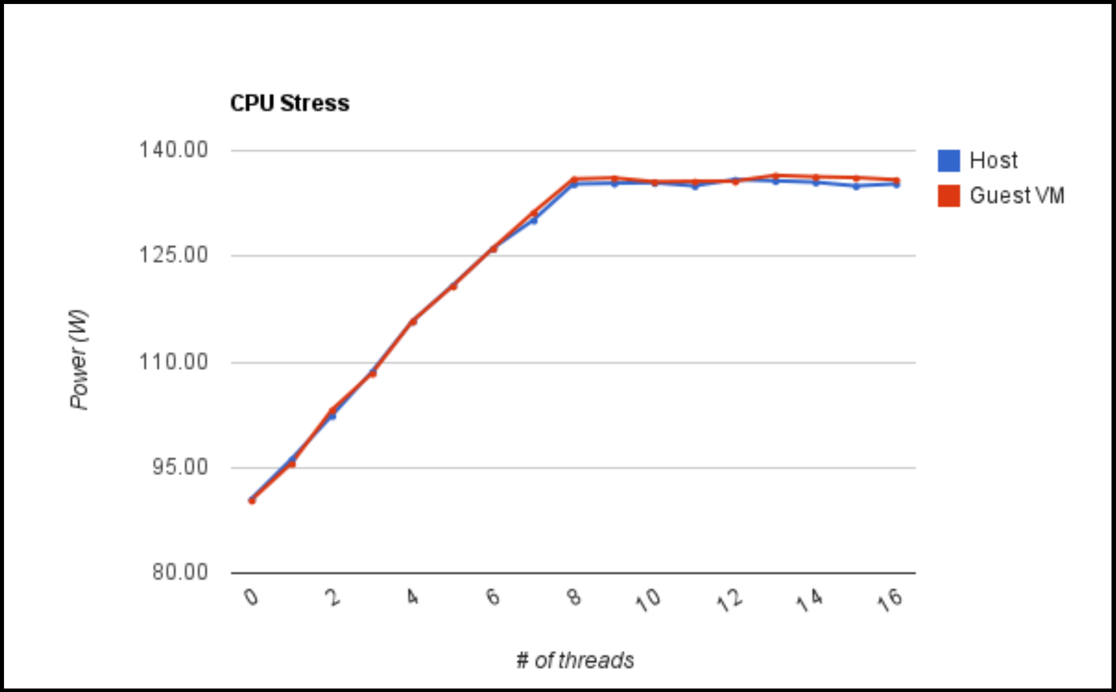

VM or host

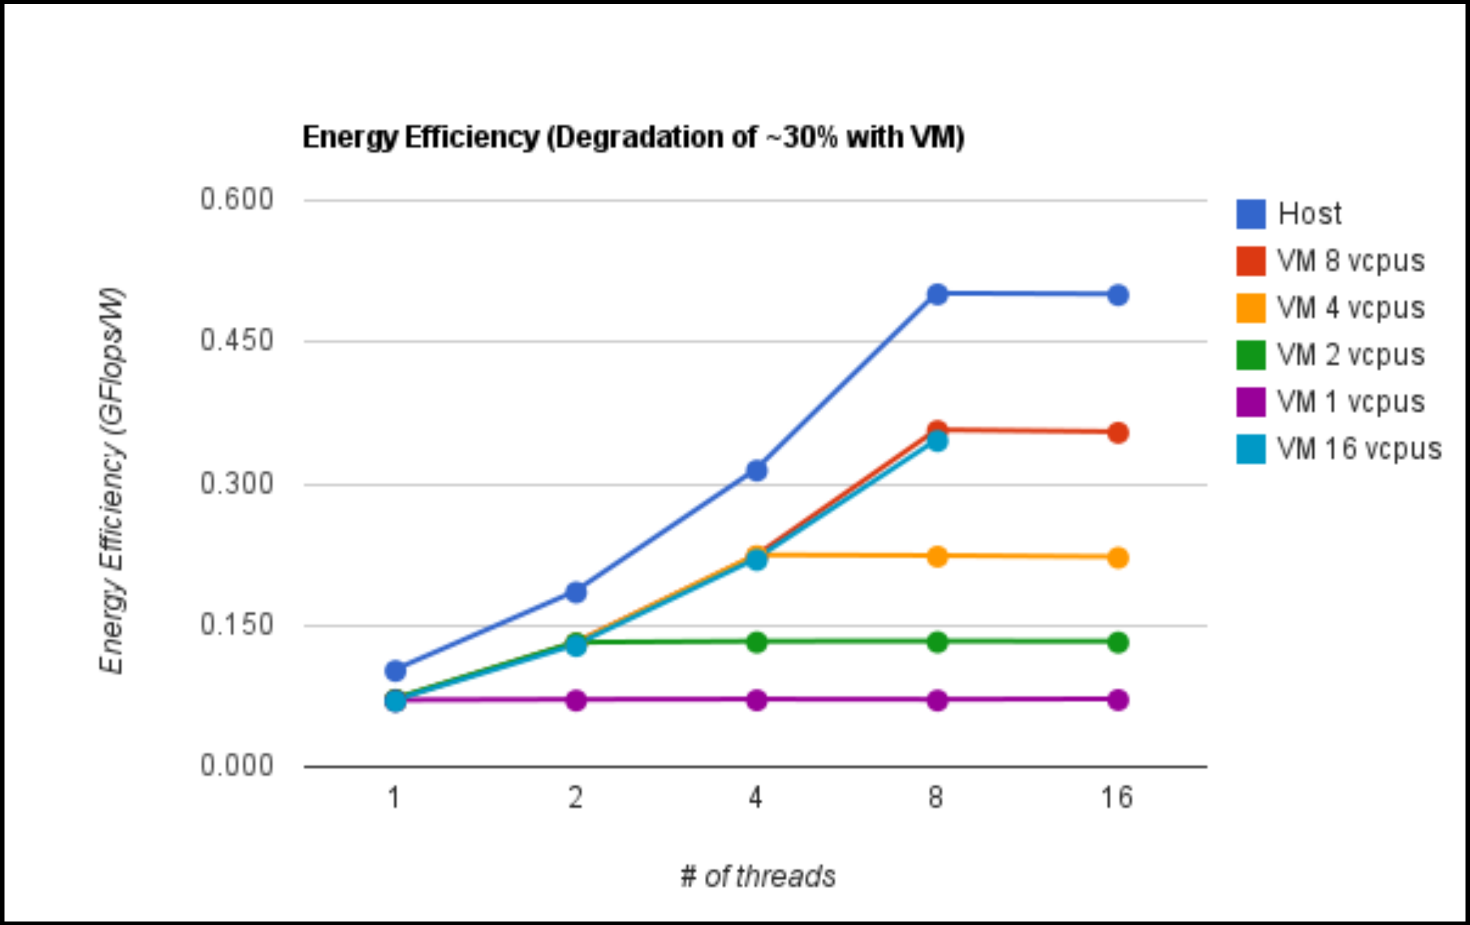

Each benchmark is run with the same amount of memory.

The degradation in energy efficiency of VMs is around 30% compared with the host.

Memory dependence

Semantic web approach in GreenClouds

- Distributed info system describing current and historical load on infrastructure including parameters of jobs running
- Describe contextual parameters (energy sources, etc.)
- Dynamically optimize and migrate if context changes

The VM Turntable Demonstrator

iGrid2005

SC2005

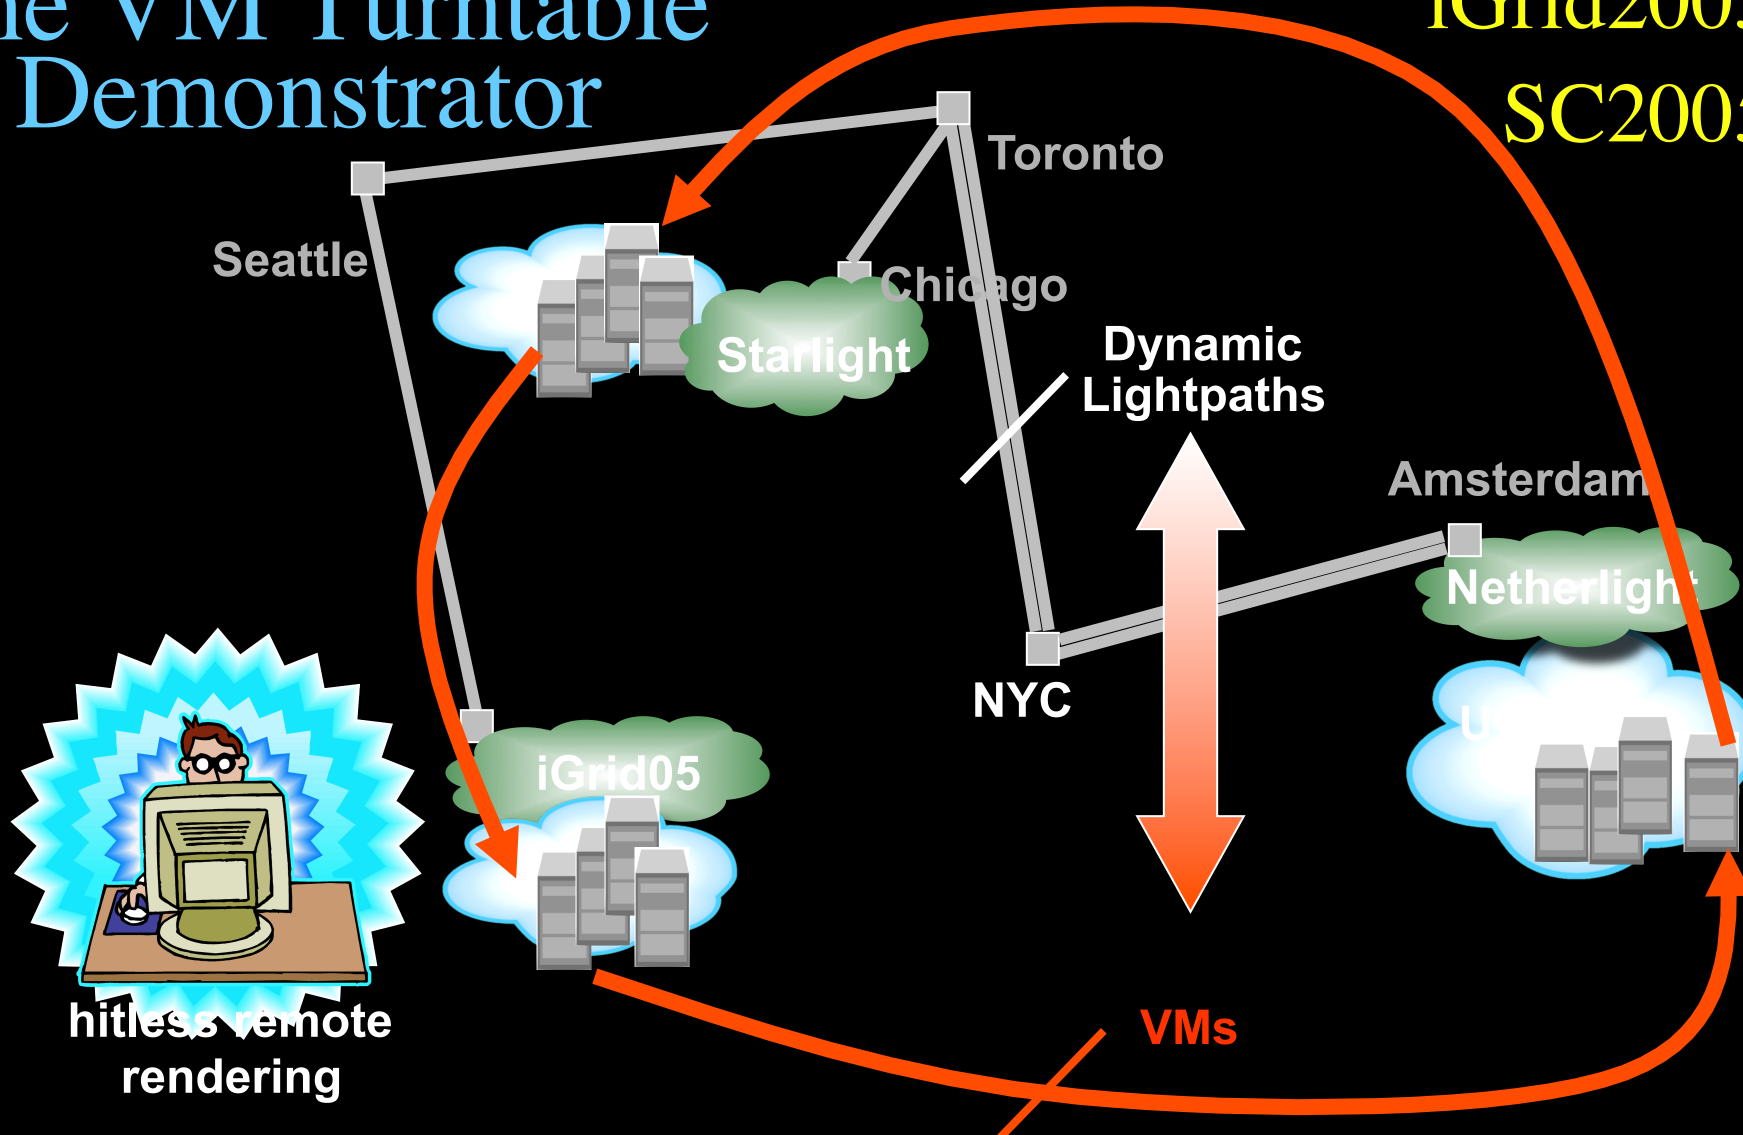

The VMs that are live-migrated run an iterative search-refine-search workflow against data stored in different databases at the various locations. A user in San Diego gets hitless rendering of search progress as VMs spin around

Network Description Language

- From semantic Web / Resource Description Framework.
- The RDF uses XML as an interchange syntax.
- Data is described by triplets:

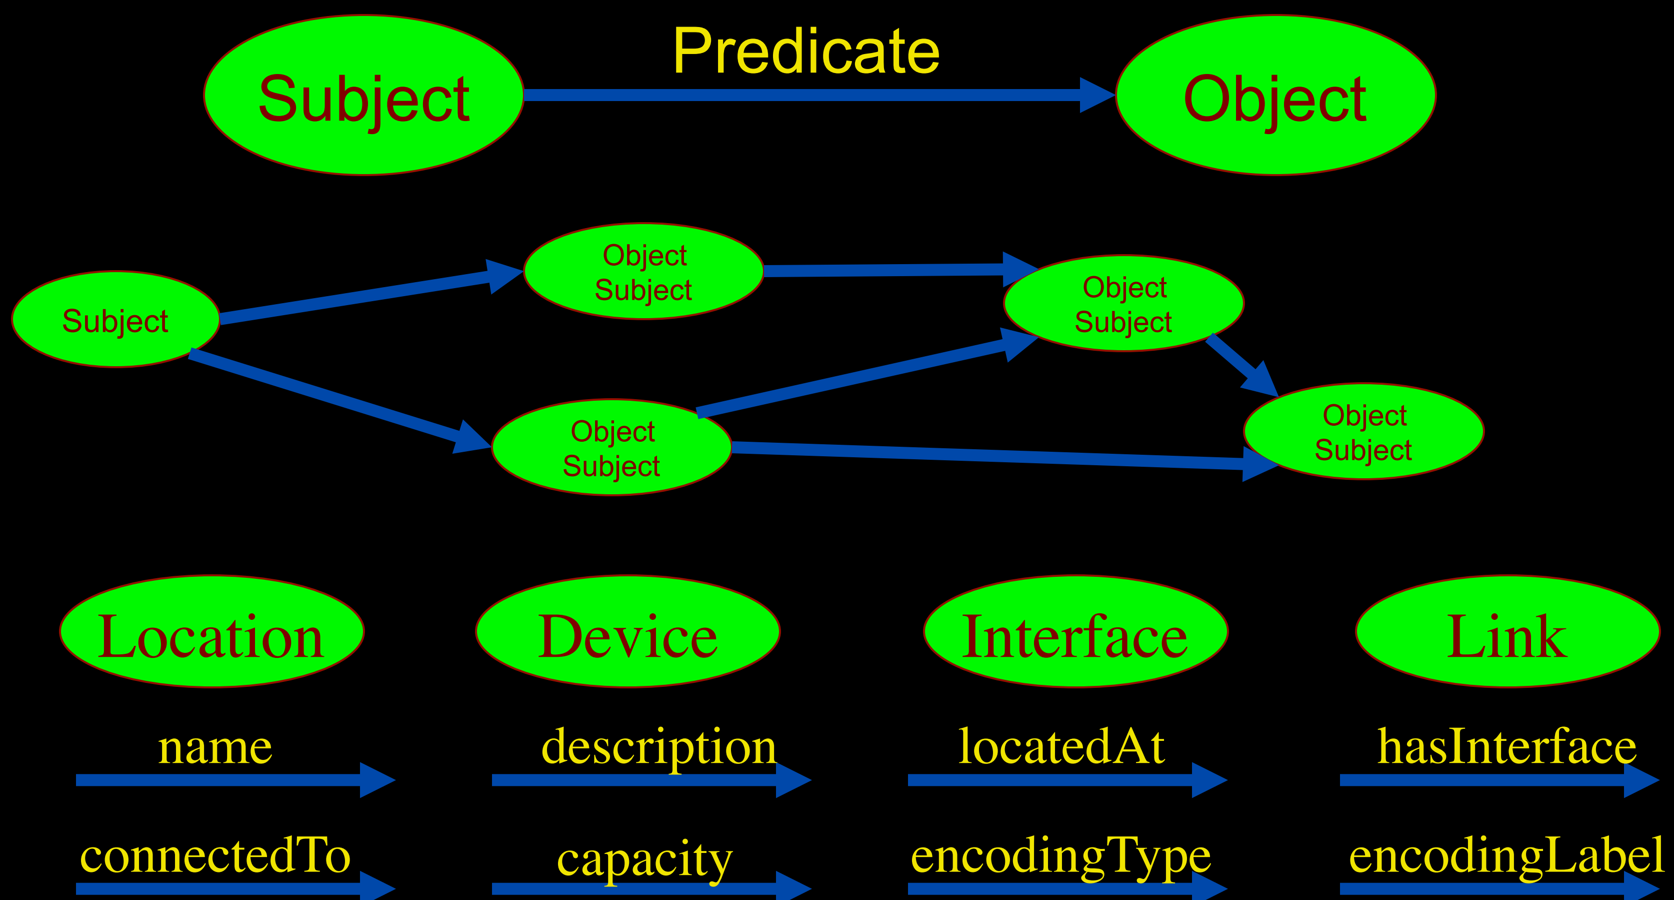

NDL + PROLOG

Research Questions:

- order of requests
- complex requests
- usable leftovers

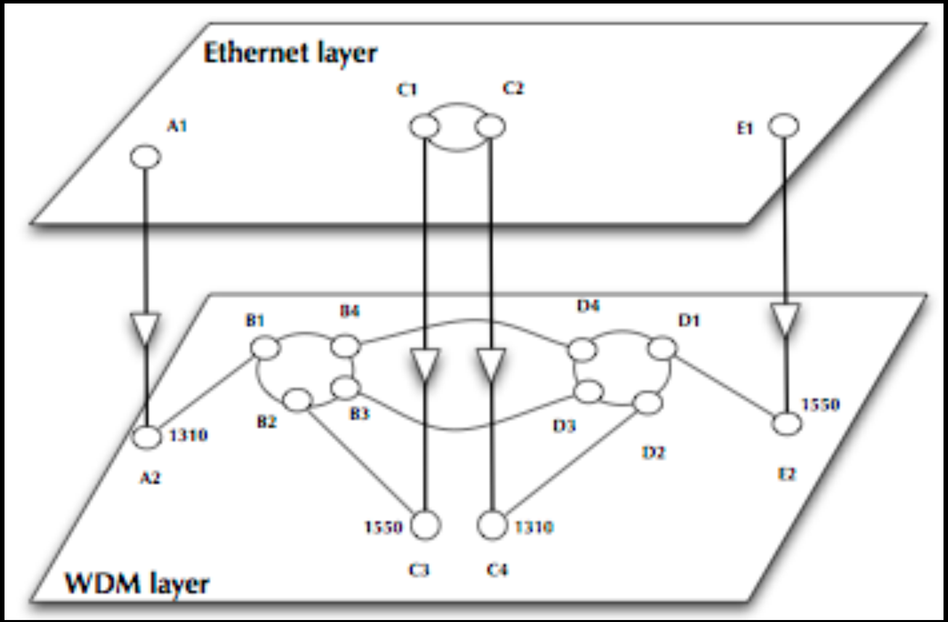

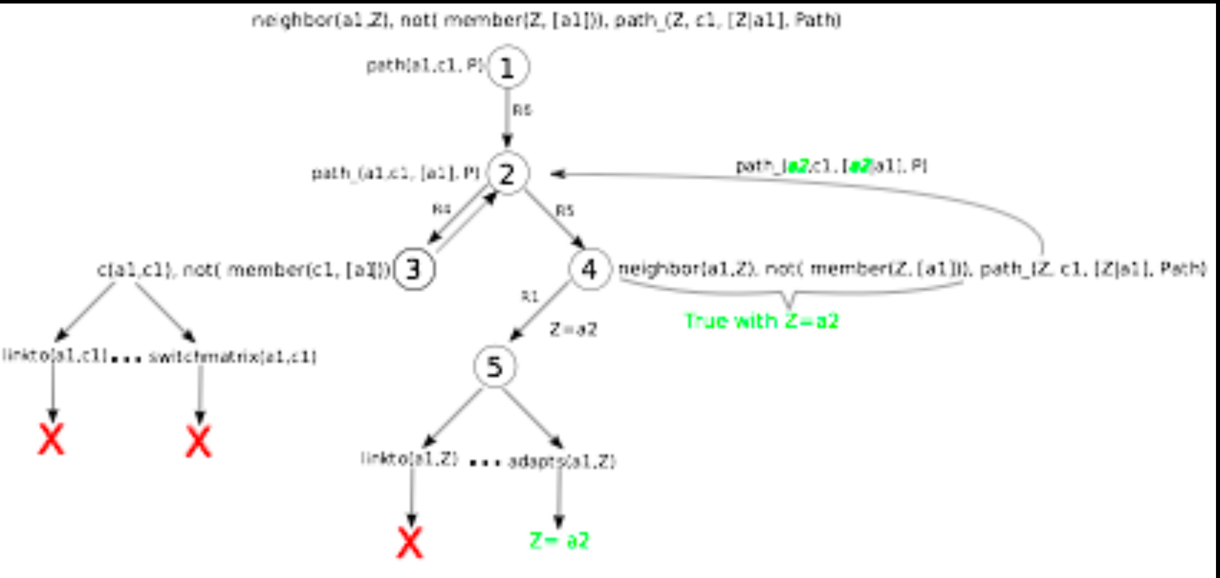

- Reason about graphs
- Find sub-graphs that comply with rules

Applications and Networks become aware of each other!

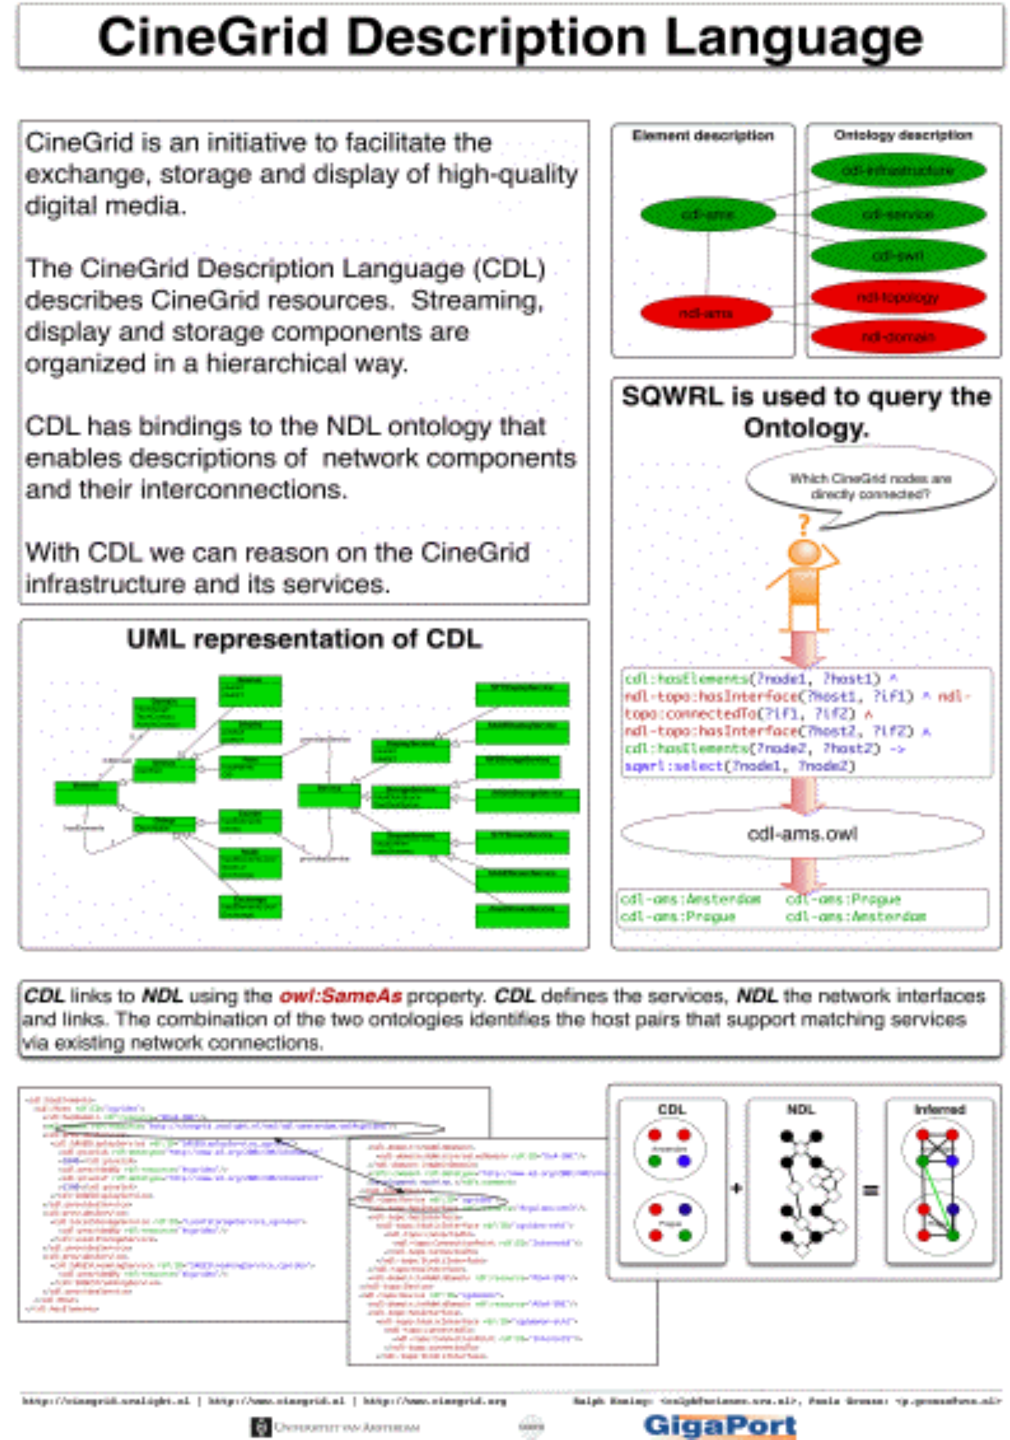

RDF describing Infrastructure

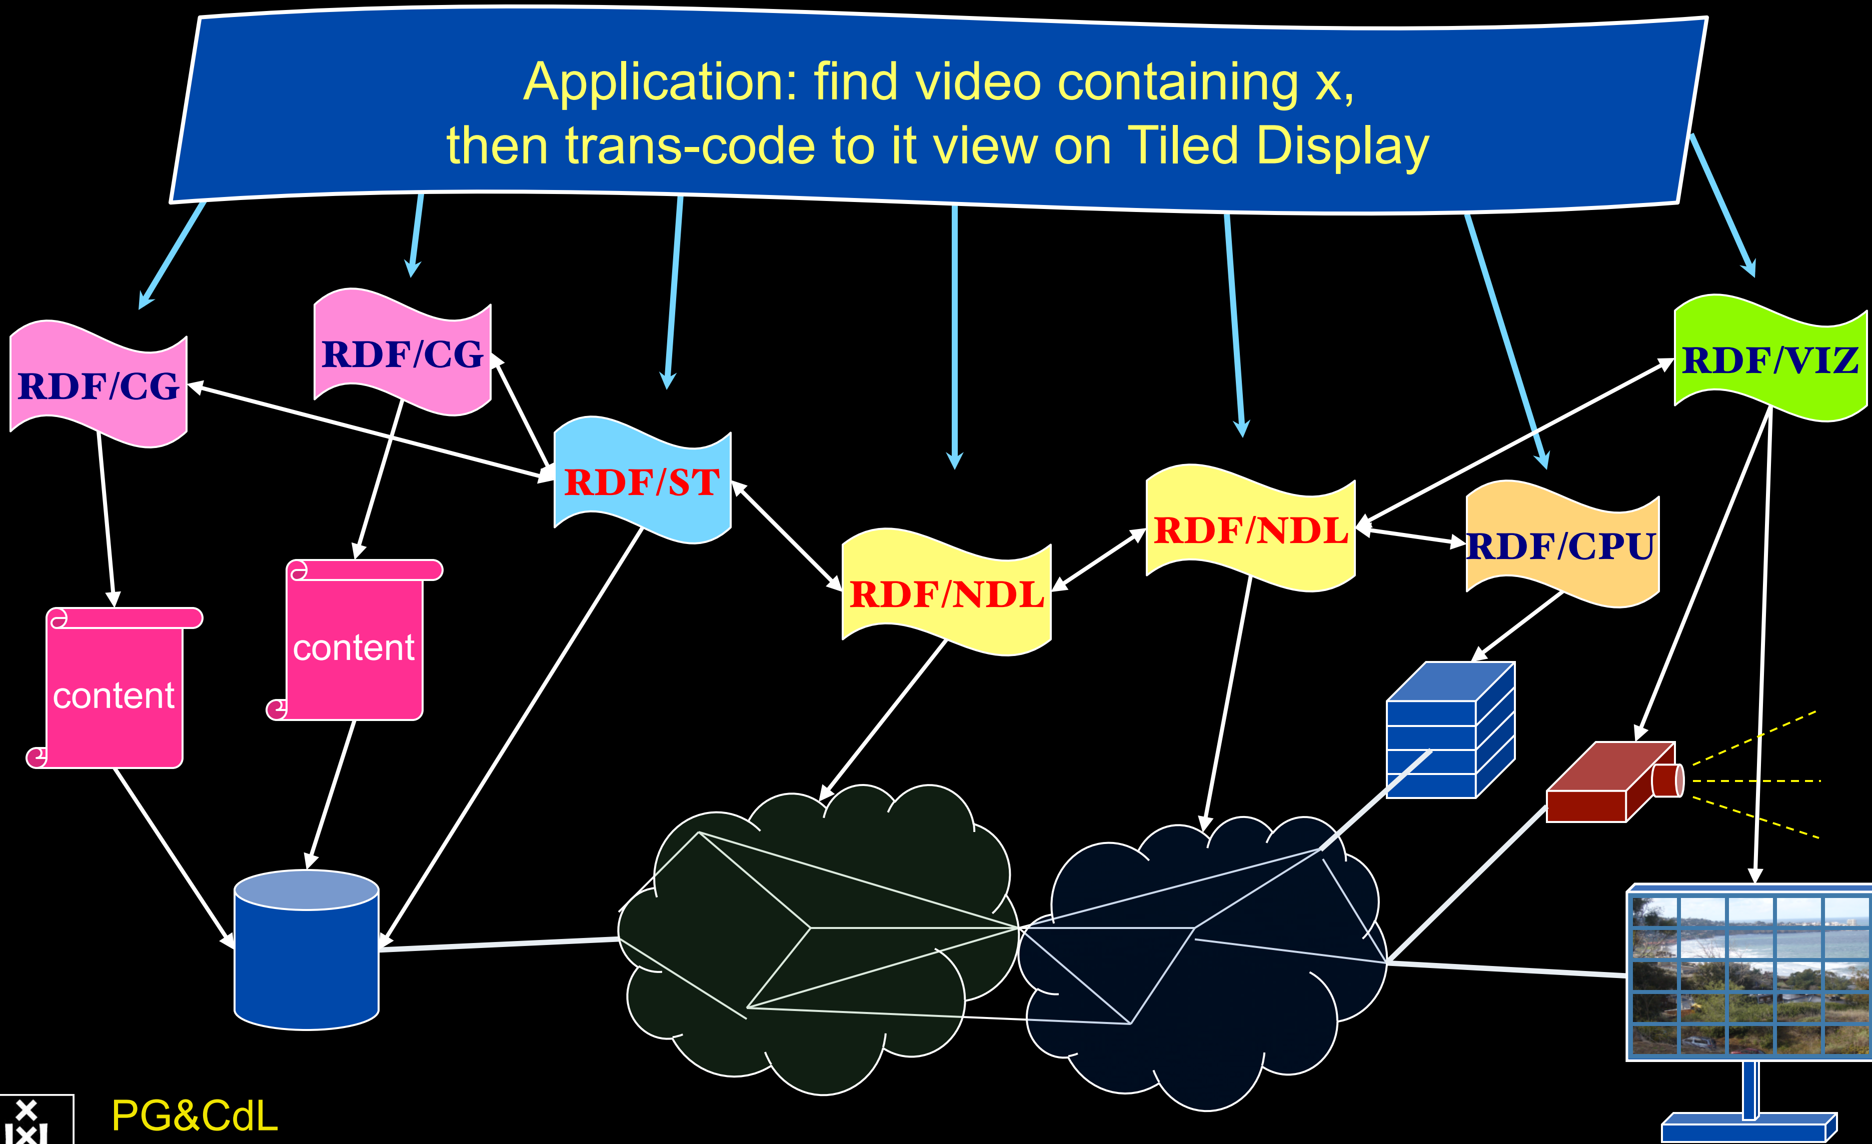

PG&CdL

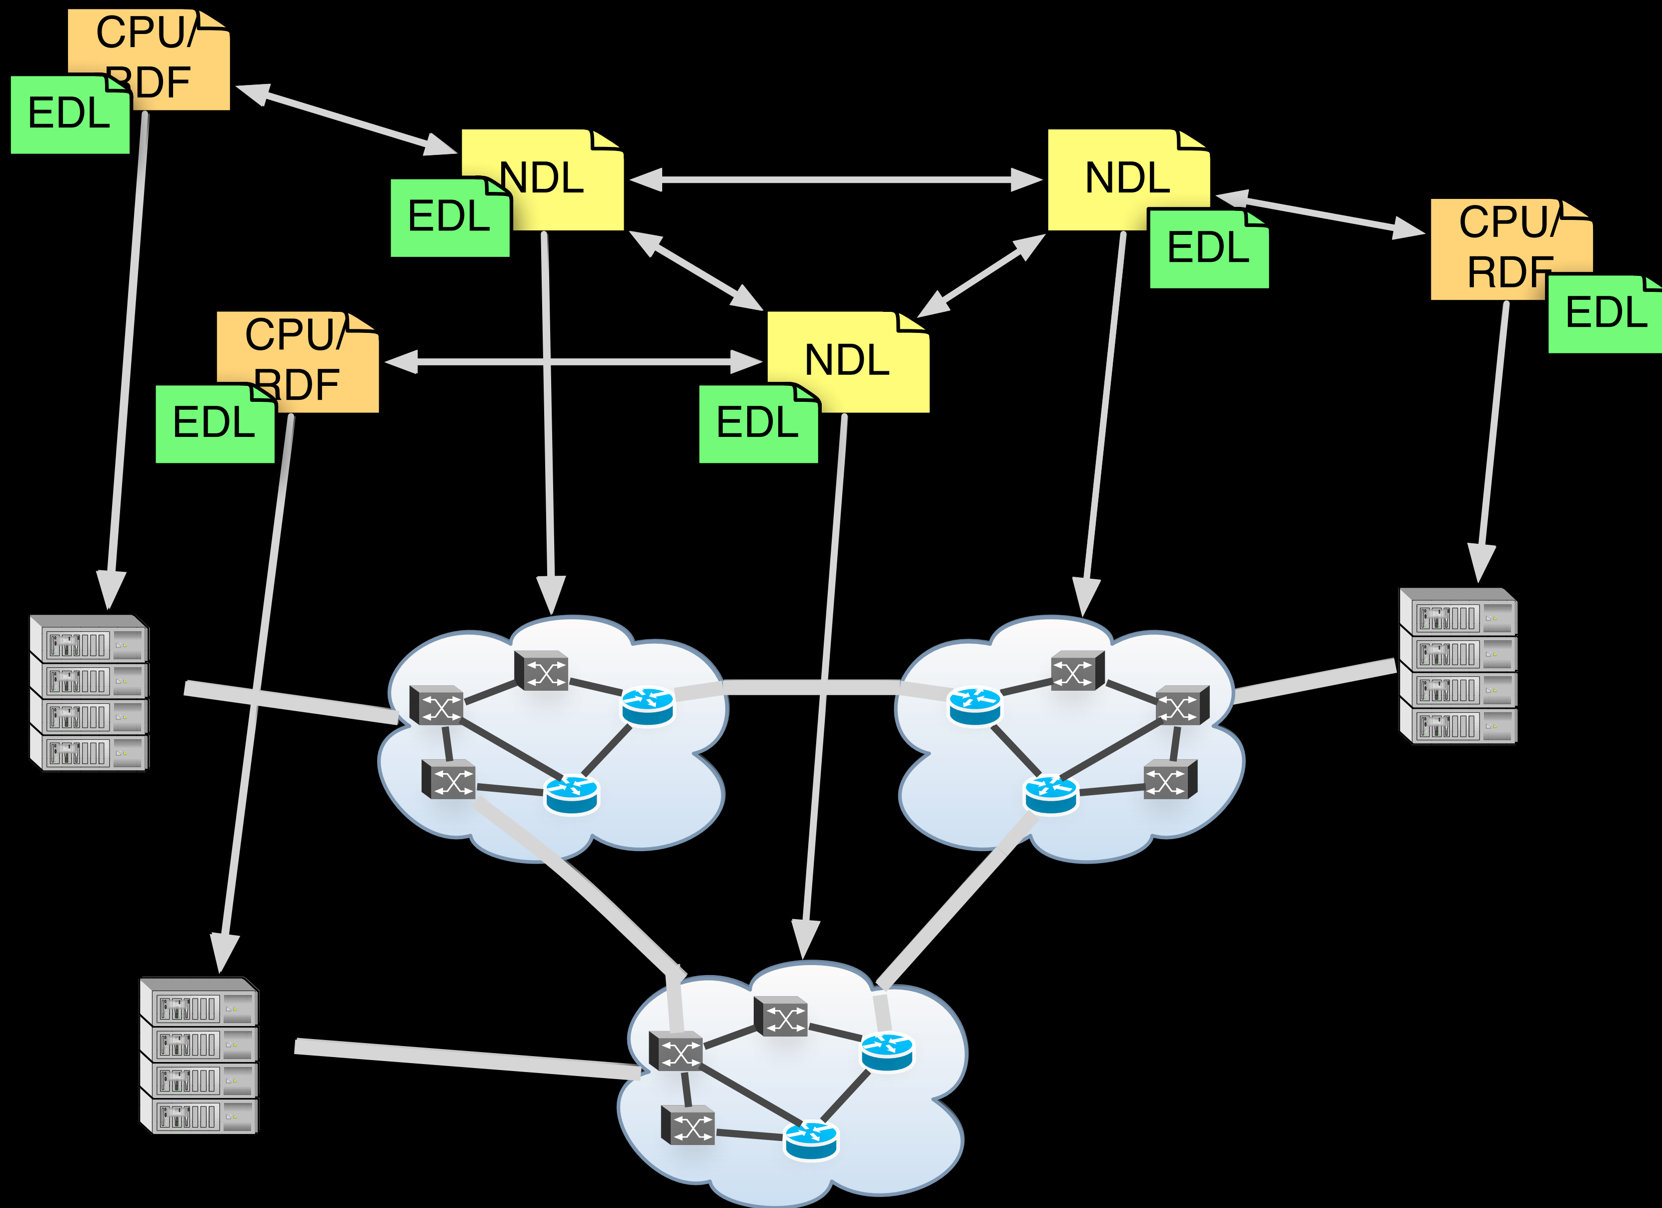
CPU/RDF
EDL
NDL
EDL
NDL
EDL
CPU/RDF
EDL
CPU/RDF
EDL
NDL
EDL

EU-NOVI

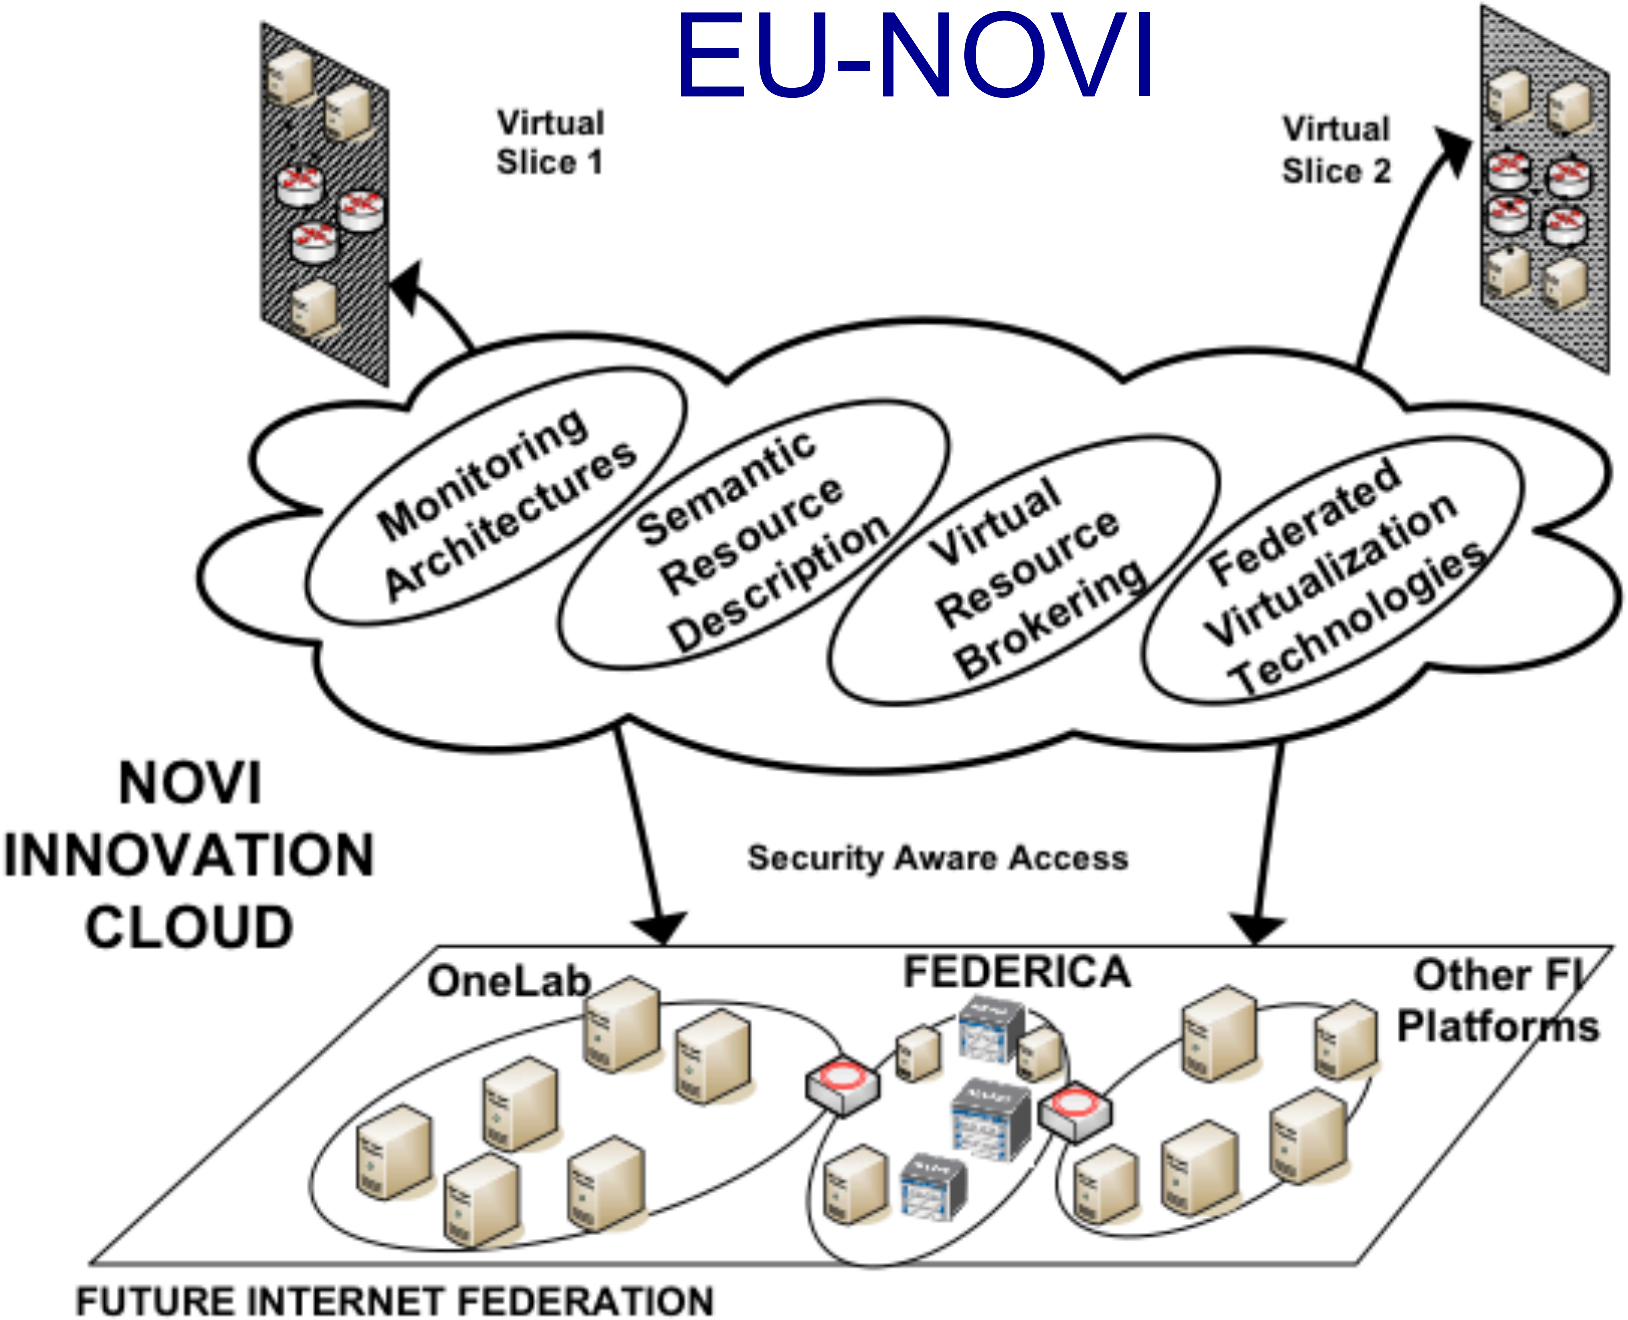
Virtual
Slice 1
Virtual
Slice 2
Monitoring
Architectures
Semantic
Resource
Description
Virtual
Resource
Brokering
Federated
Virtualization
Technologies
NOVI
INNOVATION
CLOUD
Security Aware Access
OneLab
FEDERICA
Other FI
Platforms
FUTURE INTERNET FEDERATION

GEYSERS EU

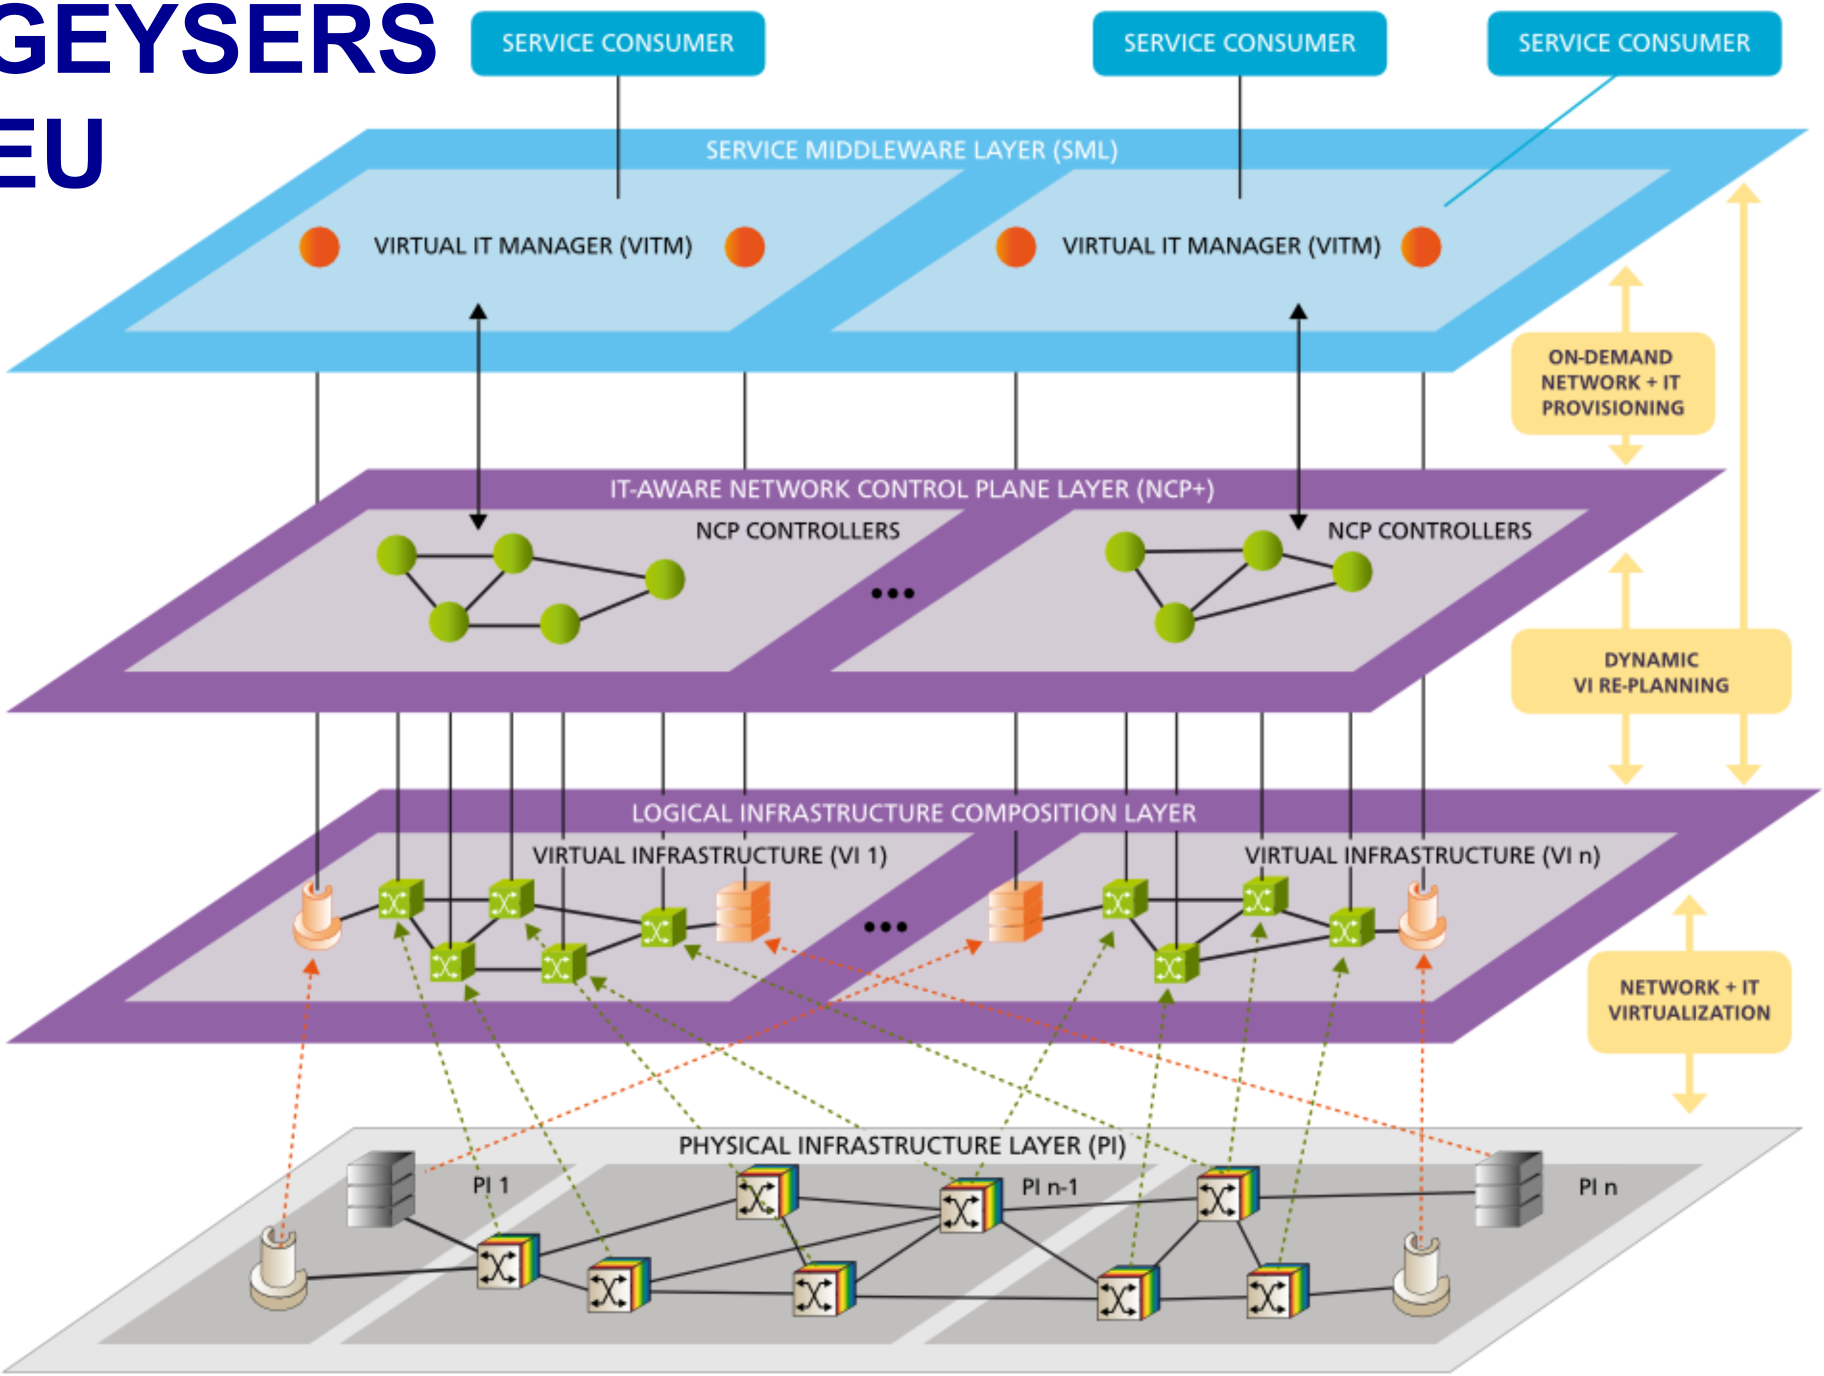
SERVICE CONSUMER
SERVICE CONSUMER
SERVICE CONSUMER
SERVICE MIDDLEWARE LAYER (SML)
VIRTUAL IT MANAGER (VITM)
VIRTUAL IT MANAGER (VITM)
ON-DEMAND NETWORK + IT PROVISIONING
IT-AWARE NETWORK CONTROL PLANE LAYER (NCP+)
NCP CONTROLLERS
NCP CONTROLLERS
DYNAMIC VI RE-PLANNING
LOGICAL INFRASTRUCTURE COMPOSITION LAYER
VIRTUAL INFRASTRUCTURE (VI 1)
VIRTUAL INFRASTRUCTURE (VI n)
NETWORK + IT VIRTUALIZATION
PHYSICAL INFRASTRUCTURE LAYER (PI)
PI 1
PI n-1
PI n

ECO-Scheduling

Future research

- Parametrize energy usage
- Describe energy properties of infrastructure
- (automatically) determine footprint of applications
- Dynamically optimize and migrate
- First order:

  $E = a*CPU + b*GPU + c*mem + d*trans + e*rw$

Q&A

http://ext.delaat.net/smartgreen/index.html