COMMIT/

Managing e-Infrastructures

Dr. Paola Grosso
Email: p.grosso@uva.nl
URL: http://staff.science.uva.nl/~grosso

Universiteit van Amsterdam

System and Network Engineering

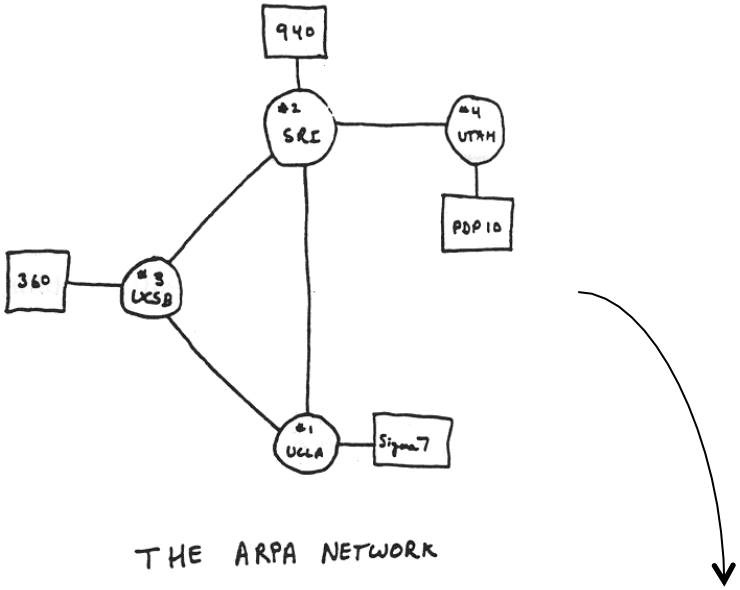

Complex (network) infrastructures

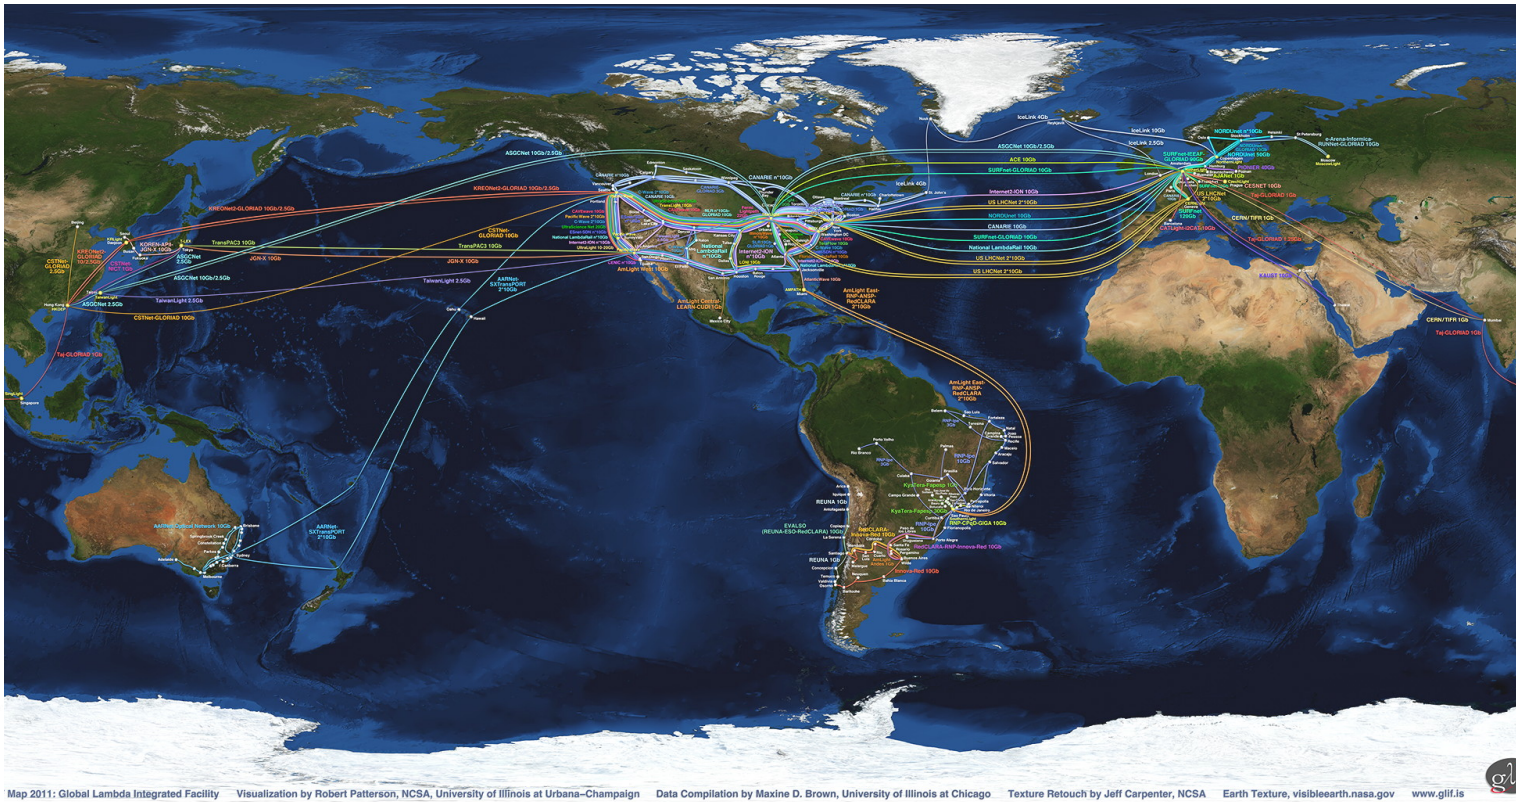

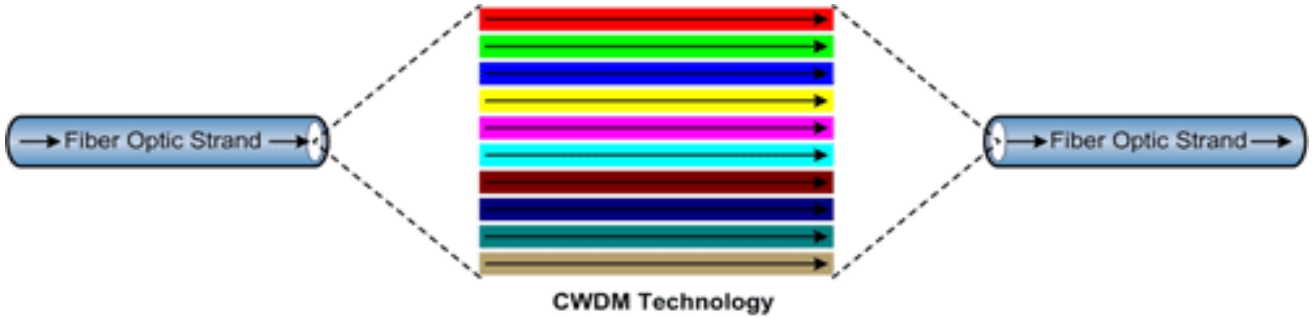

Optical transmission

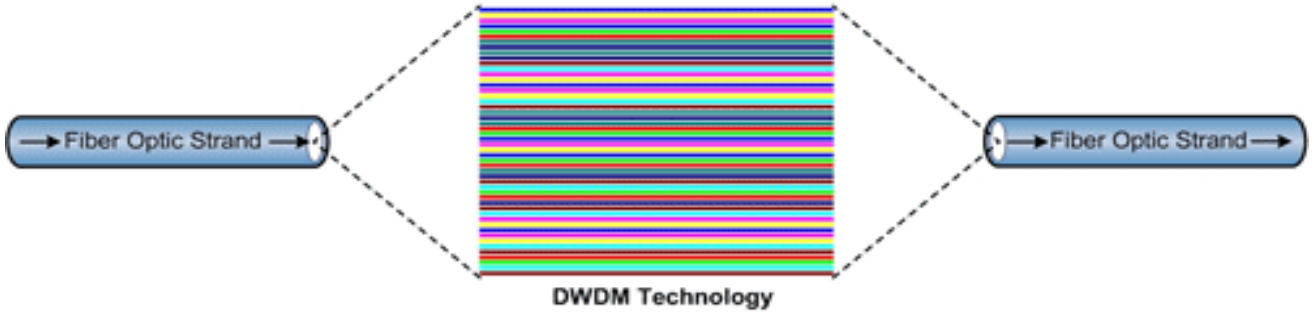

…with more possibilities

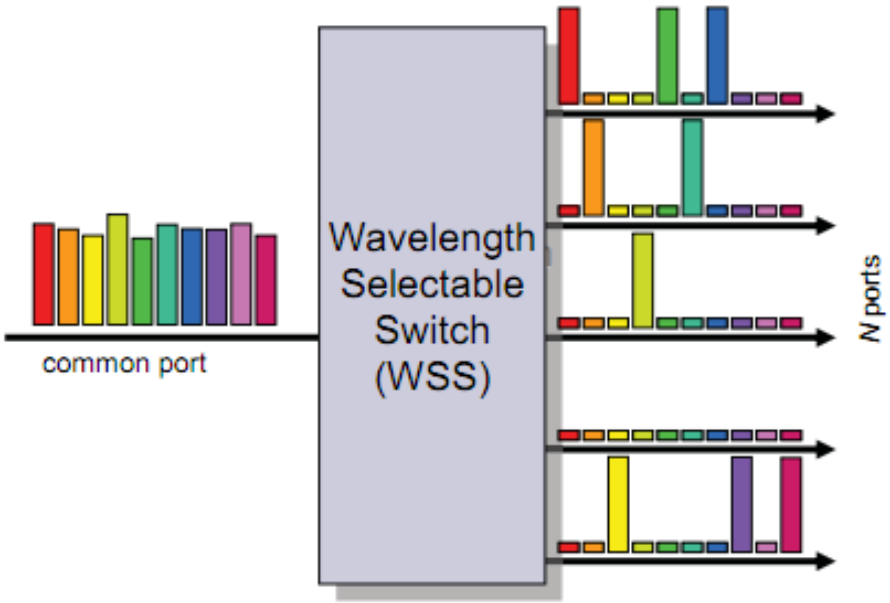

Virtualization

The Big Data Challenge

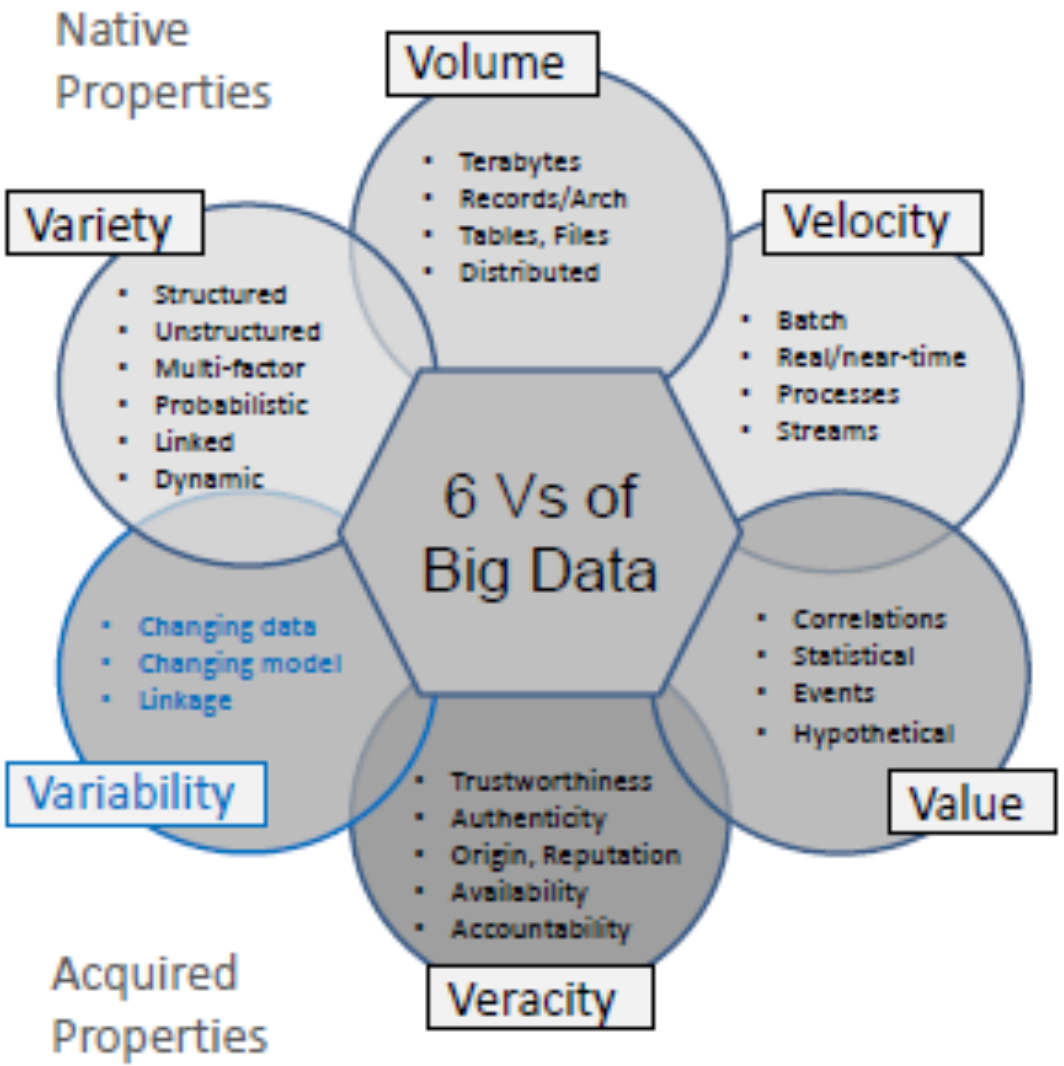

The SE ExoGENI Rack

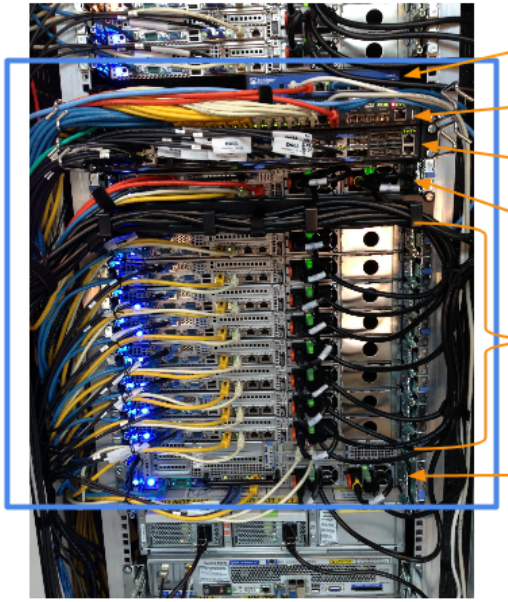

VPN Gateway - Juniper SRX100

Management Switch - Dell Force10 S55

OpenFlow Switch - Dell Force10 S4810P

Head Node - Dell R620

Compute Nodes - 8 x Dell R620

Storage Node - Dell R720

The OpenLab in Amsterdam

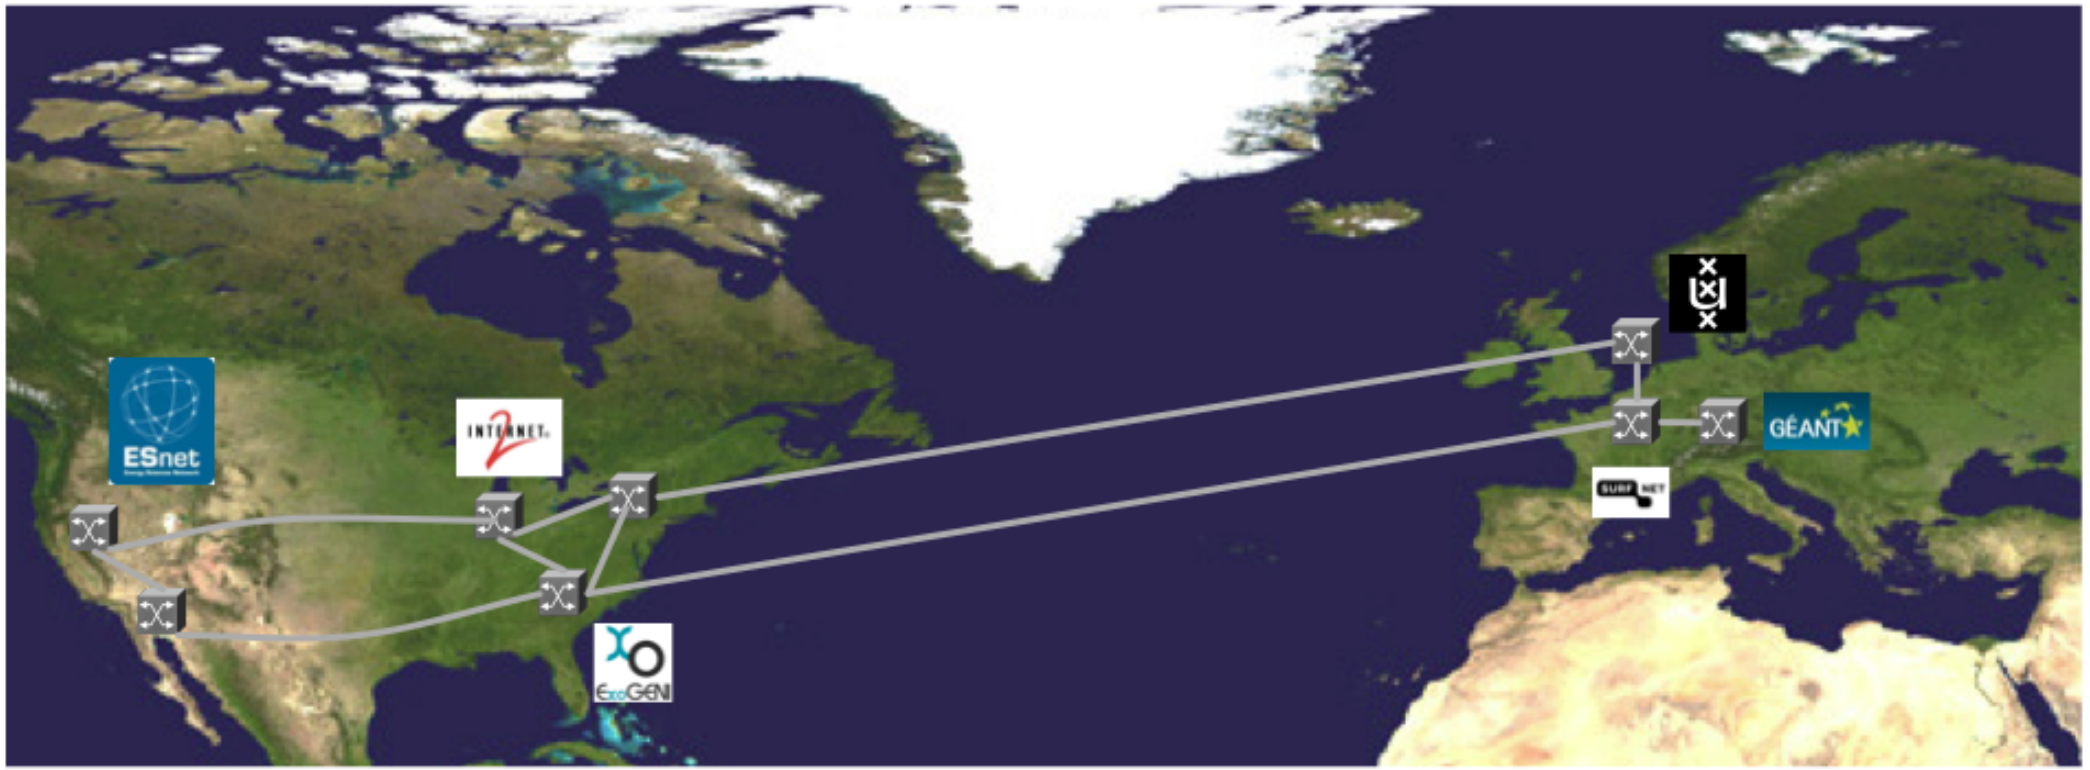

COMMIT/

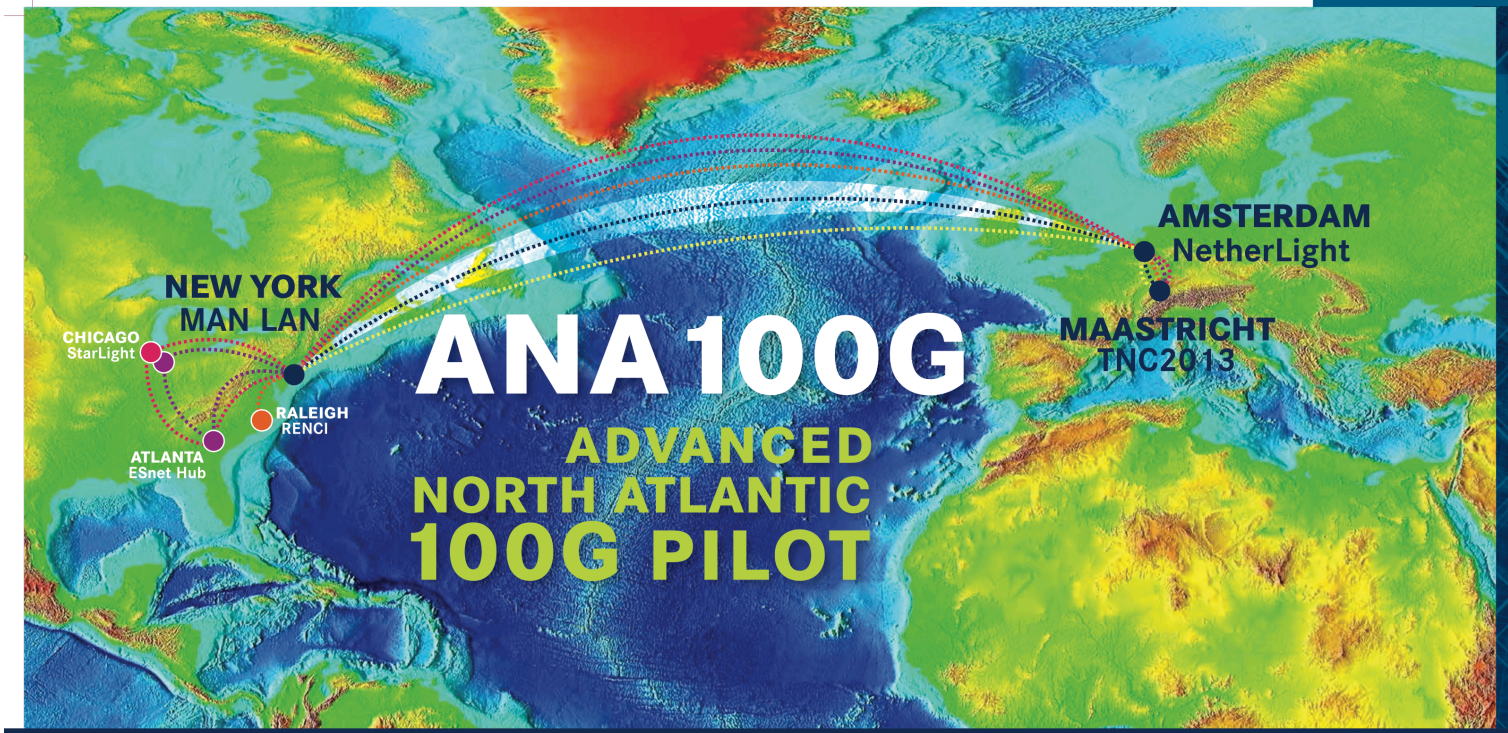

INTERNET2 · NORDUnet · ESnet · SURFnet · canarie 1993-2013 · ciena · JUNIPER NETWORKS · GÉANT · TATA · TATA COMMUNICATIONS · UNIVERSITY OF AMSTERDAM

TNC2013 DEMOS JUNE, 2013

DEMO	TITLE	OWNER	AFFILIATION	E-MAIL	A-SIDE	Z-SIDE	PORTS(S) MAN LAN	PORTS(S) TNC2013	DETAILS
1	Big data transfers with multipathing, OpenFlow and MPTCP	Ronald van der Pol	SURFnet	ronald.vanderpol@surfnet.nl	TNC/MECC, Maastricht NL	Chicago, IL	Existing 100G link between internet2 and ESnet	2x40GE (Juniper)+ 2x10GE (OME6500)	In this demonstration we show how multipathing, OpenFlow and Multipath TCP (MPTCP) can help in large file transfers between data centres (Maastricht and Chicago). An OpenFlow application provisions multiple paths between the servers and MPTCP will be used on the servers to simultaneously send traffic across all those paths. This demo uses 2x40G on the transatlantic 100G link. ESnet provides 2x40G between MAN LAN and StarLight, ACE and USLHCnet provide additional 10GEs.
2	Visualize 100G traffic	Inder Monga	ESnet	imonga@es.net					Using an SNMP feed from the Juniper switch at TNC2013,and/or Brocade AL2S node in MANLAN, this demo would visualize the total traffic on the link, of all demos aggregated. The network diagram will show the transatlantic topology and some of the demo topologies.
3	How many modern servers can fill a 100Gbps Transatlantic Circuit?	Inder Monga	ESnet	imonga@es.net	Chicago, III	TNC showfloor	1x 100GE	8x 10GE	In this demonstration, we show that with the proper tuning and tool, only 2 hosts on each continent can generate almost 80Gbps of traffic. Each server has 4 10G NICS connected to a 40G virtual circuit, and has iperf3 running to generate traffic. ESnet's new "iperf3" throughput measurement tool, still in 'beta', combines the best features from other tools such as iperf, nuttcp, and netperf. See: https://my.es.net/demos/tnc2013/
4	First European ExoGENI at Work	Jeroen van der Ham	UvA	vdham@uva.nl	RENCI, NC	UvA, Amsterdam, NL	1x 10GE	1x 10GE	The ExoGENI racks at RENCI and UvA will be interconnected over a 10G pipe and be on continuously, showing GENI connectivity between Amsterdam and the rest of the GENI nodes in the USA.
5	Up and down North Atlantic @ 100G	Michael Enrico	DANTE	michael.enrico@dante.net	TNC showfloor	TNC showfloor	1x 100GE	1x 100GE	The DANTE 100GE test set will be placed at the TNC2013 showfloor and connected to the Juniper at 100G. When this demo is running a loop @ MAN LAN's Brocade switch will ensure that the traffic sent to MAN LAN returns to the showfloor. On display is the throughput and RTT (to show the traffic traveled the Atlantic twice)

This brings new fundamental questions with regard to scalability, robustness and sustainability.

A new model for the Internet

Can we build an Internet that is smart and sustainable?

With:

- Determinist and federative behaviors
- Flexible and dynamic communication
- More intelligence in the network

Towards smart and sustainable e-infrastructures

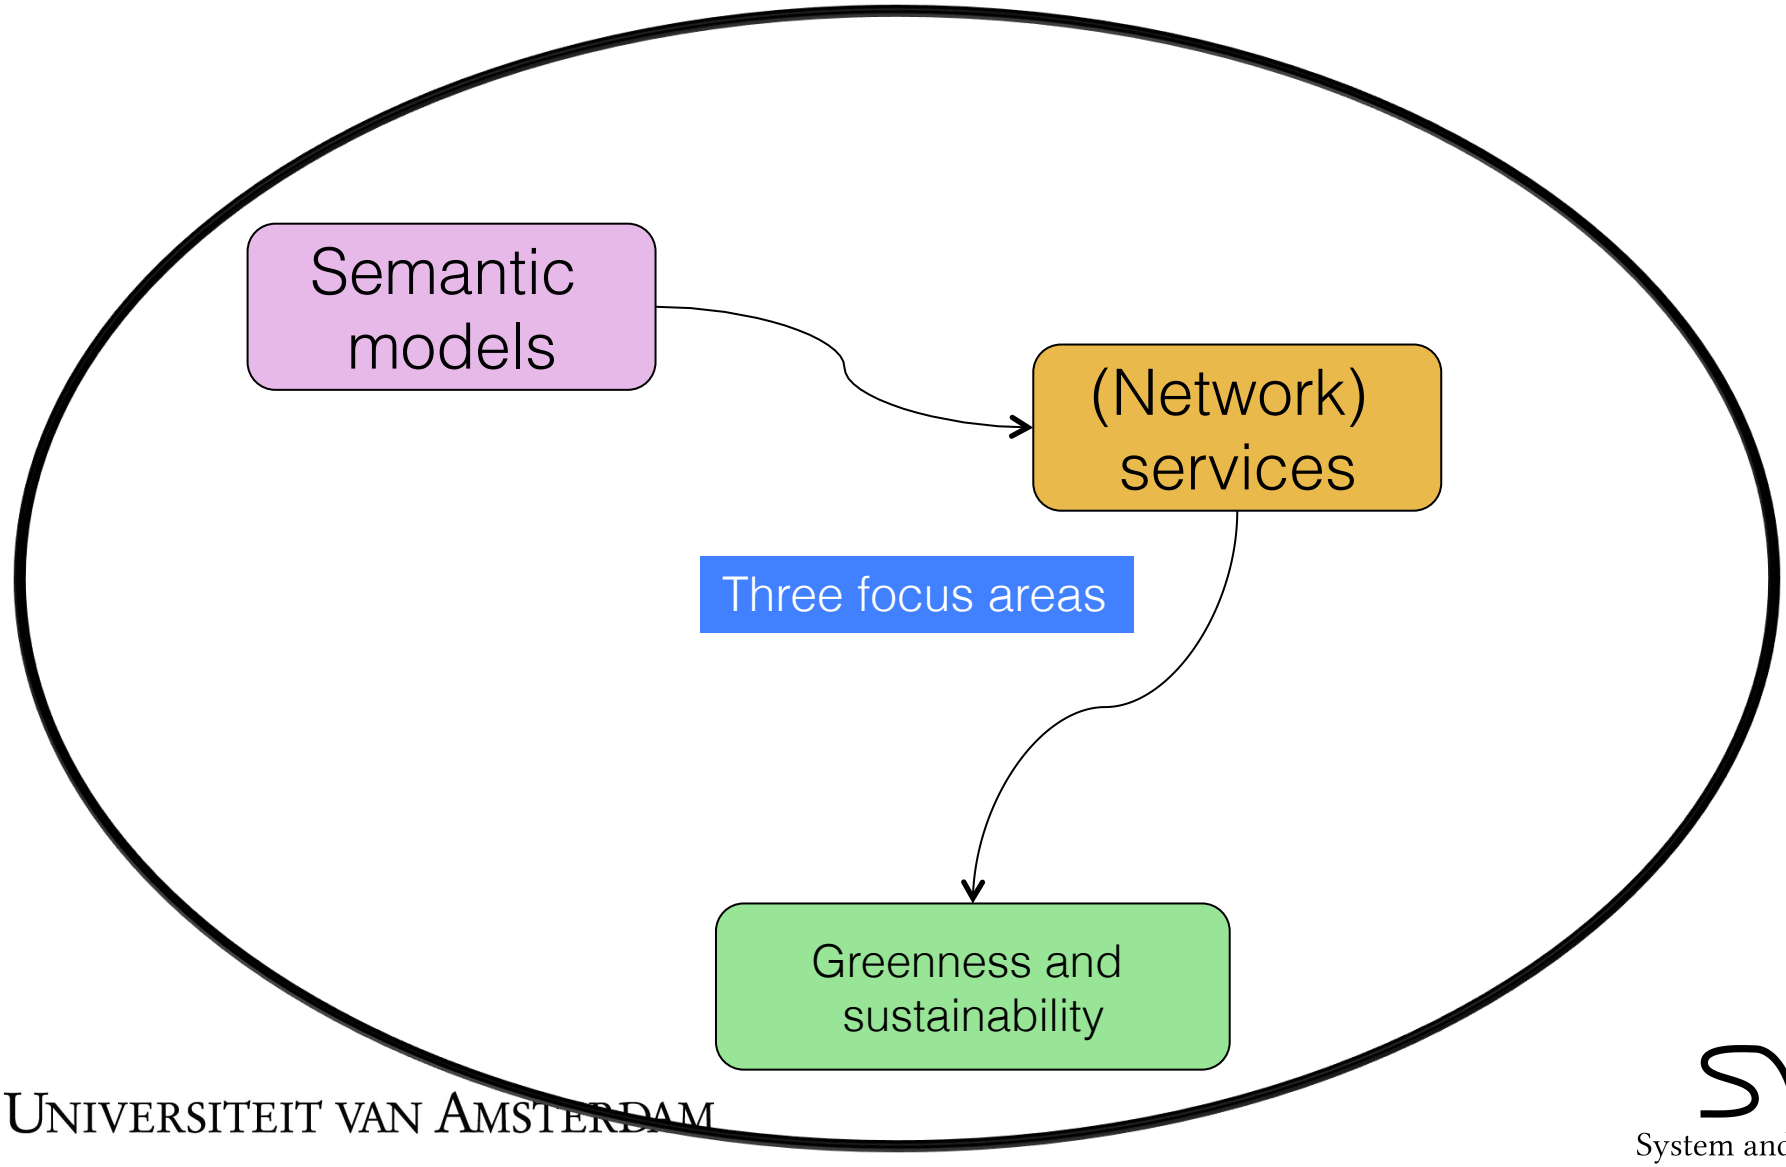

- *How can we represent complex e-Infrastructure?*

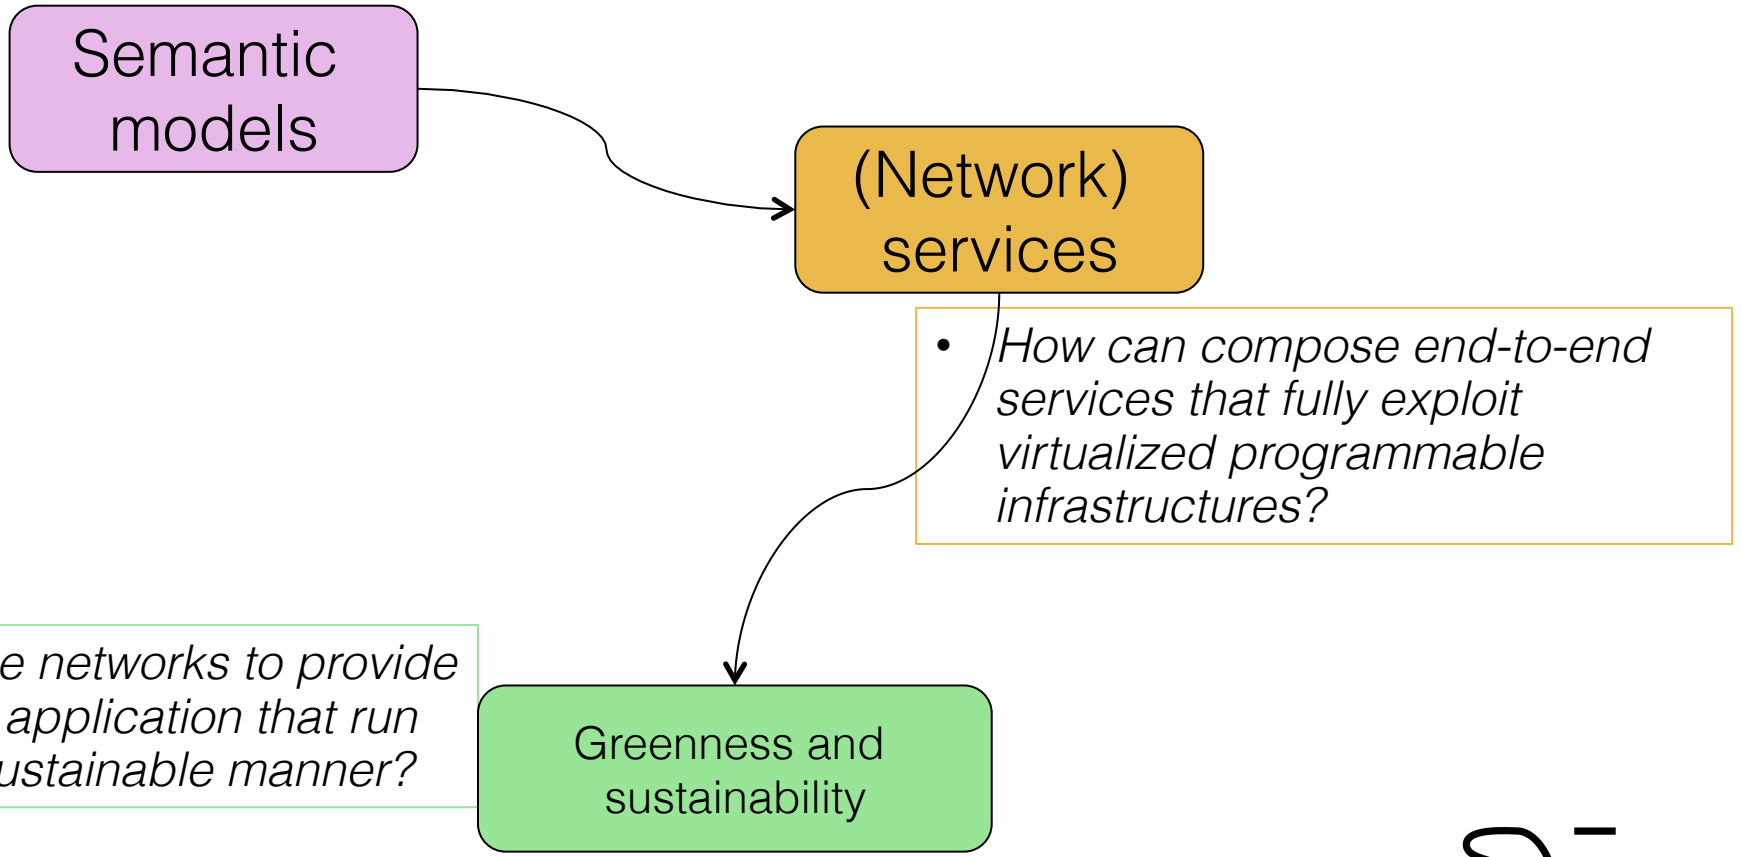

- *How can compose end-to-end services that fully exploit virtualized programmable infrastructures?*

- *Can we use networks to provide support for application that run in a more sustainable manner?*

Semantic models

"One of the **main ingredients** *in the design, implementation and operation of computing infrastructures is the* **information model**. *This information model must describe both the physical infrastructure and its virtualization aspects"*

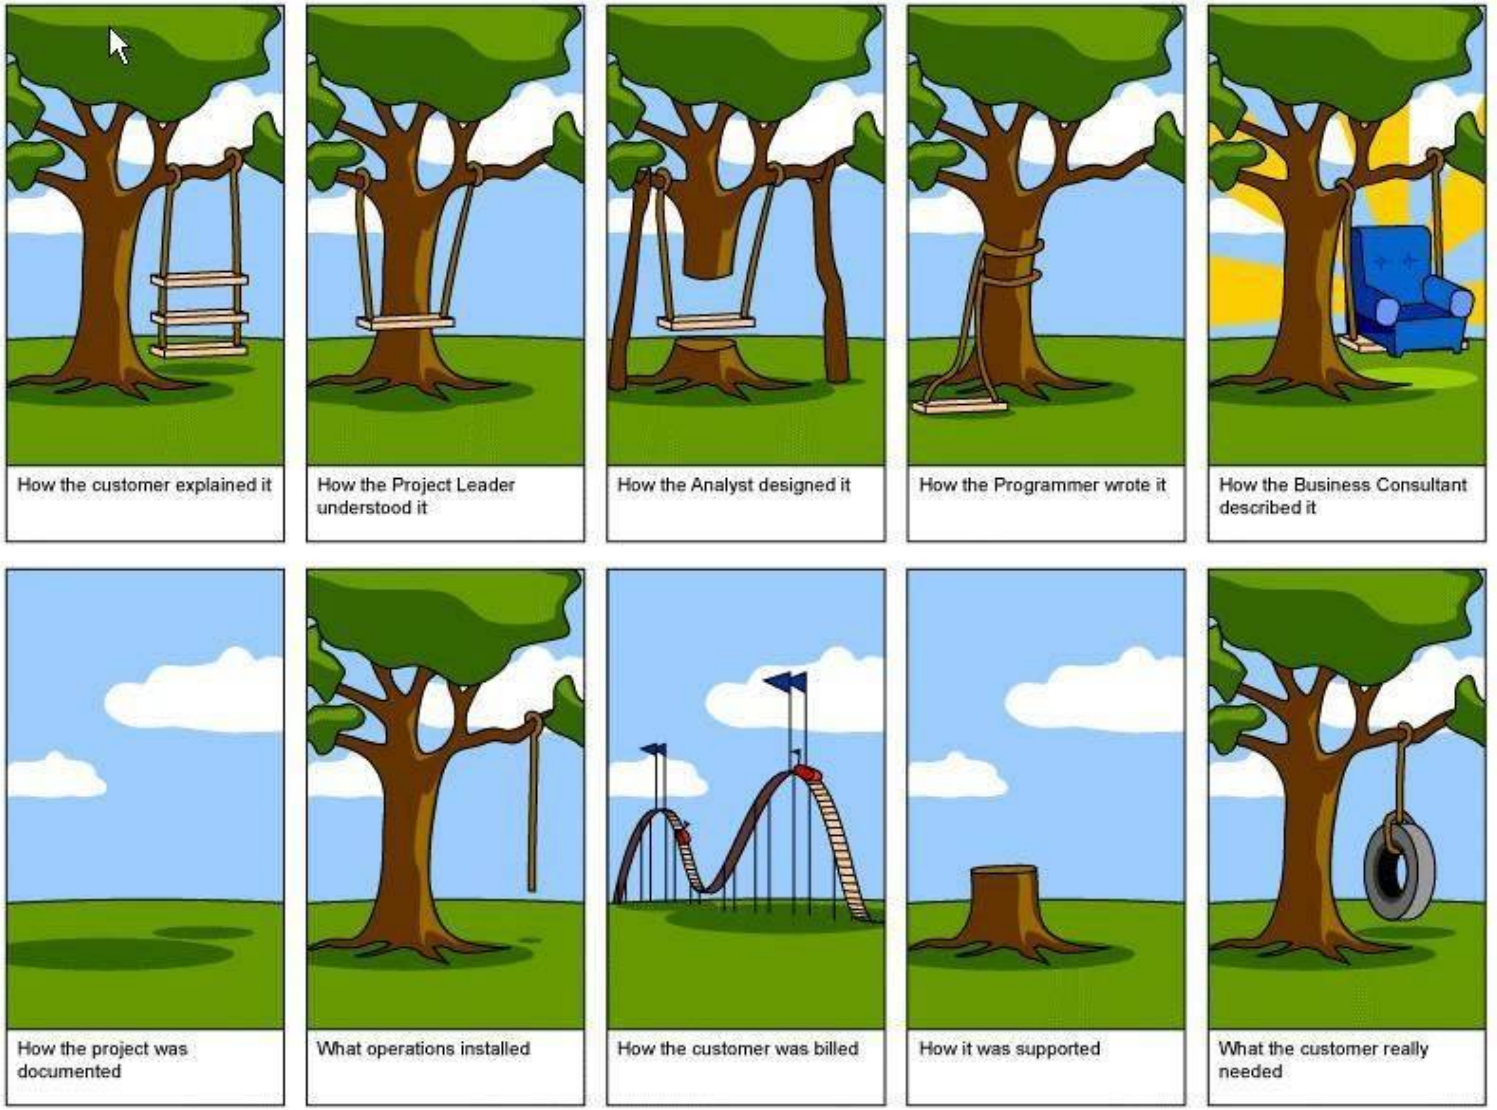
How the customer explained it
How the Project Leader understood it
How the Analyst designed it
How the Programmer wrote it
How the Business Consultant described it
How the project was documented
What operations installed
How the customer was billed
How it was supported
What the customer really needed

INDL

An effort started in 2010 (in parallel with our involvement in the FP7 projects Geysers and NOVI).

The goal was to capture the concept of virtualization in computing infrastructures and to describe the storage and computing capabilities of the resources.

A key feature is the decoupling of virtualization, connectivity and functionalities.

It is built upon the NML ontology.
It uses the nml:node concept as basic entity to describe resources in computing infrastructures.

It can be used as:

- a stand-alone model (i.e. without any network descriptions),
- in combination with NML by importing the NML ontology into the INDL definition.

Node components

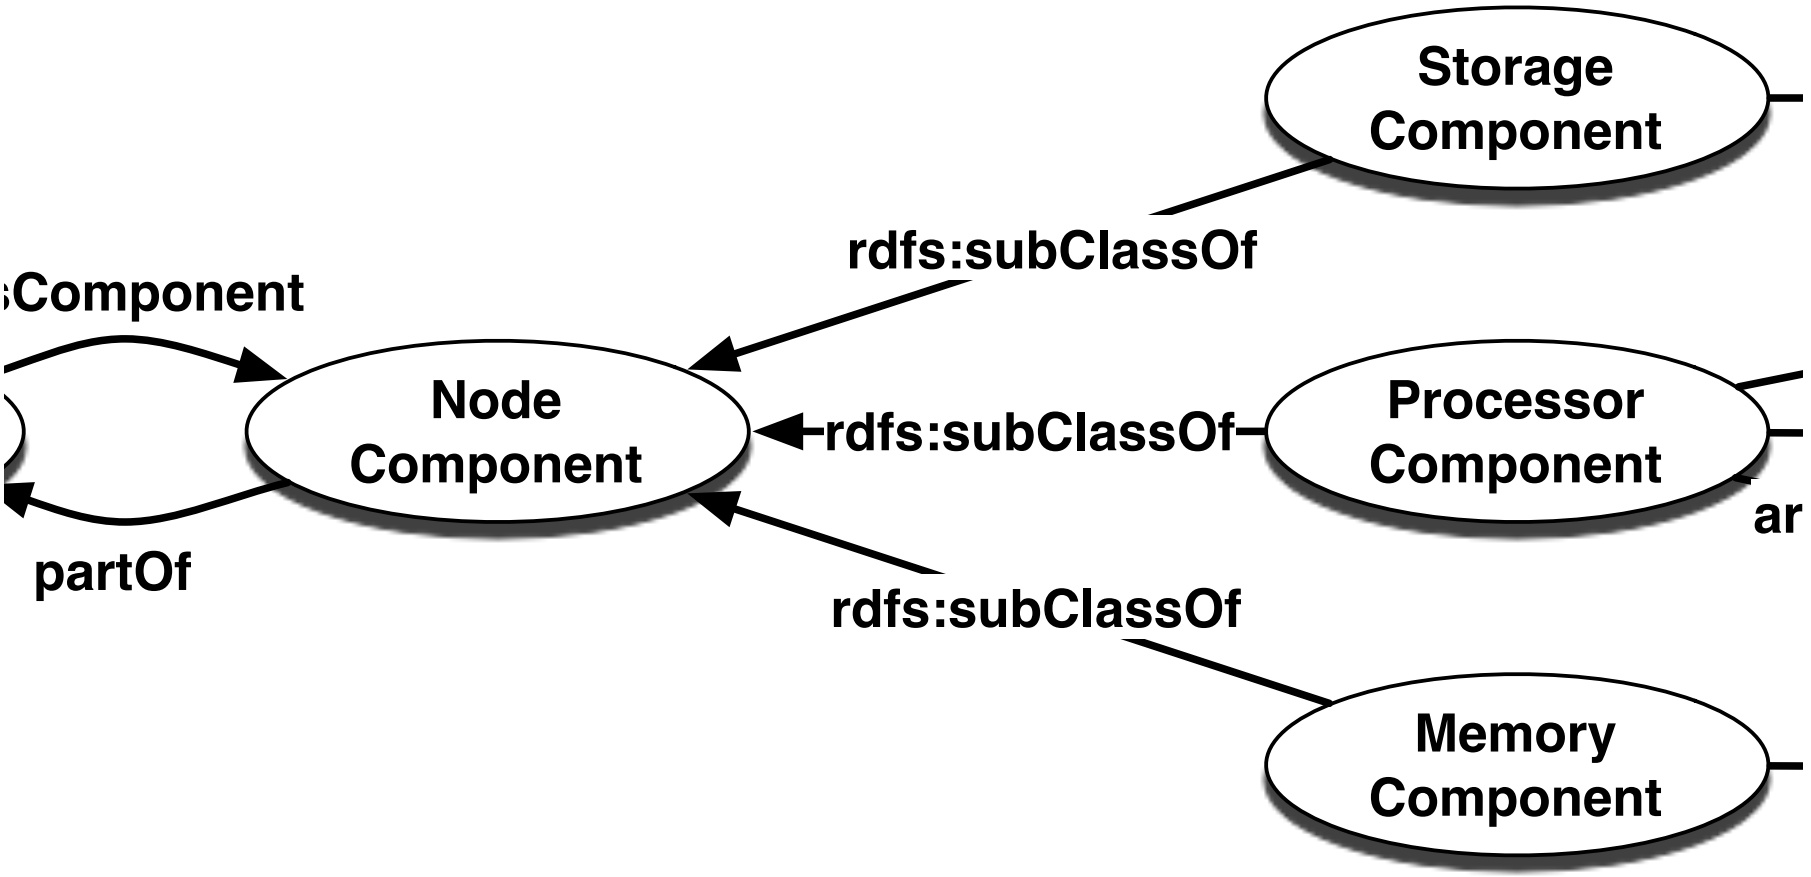

Our connecting models

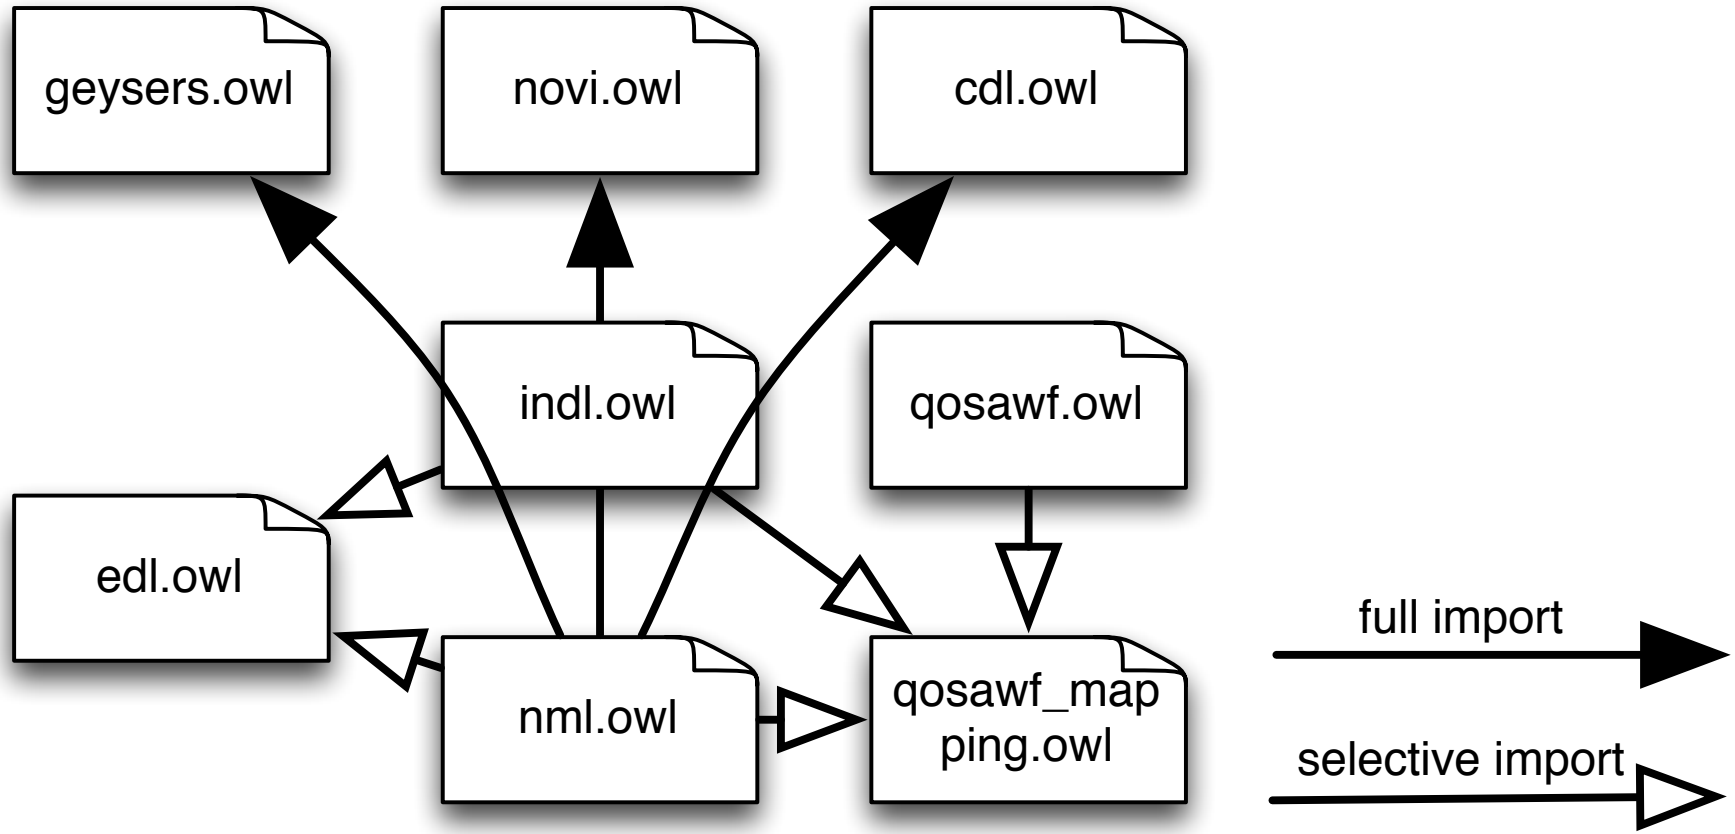

NOVI Federation

Our model is used for:

- Resource discovery
- Embedding algorithms

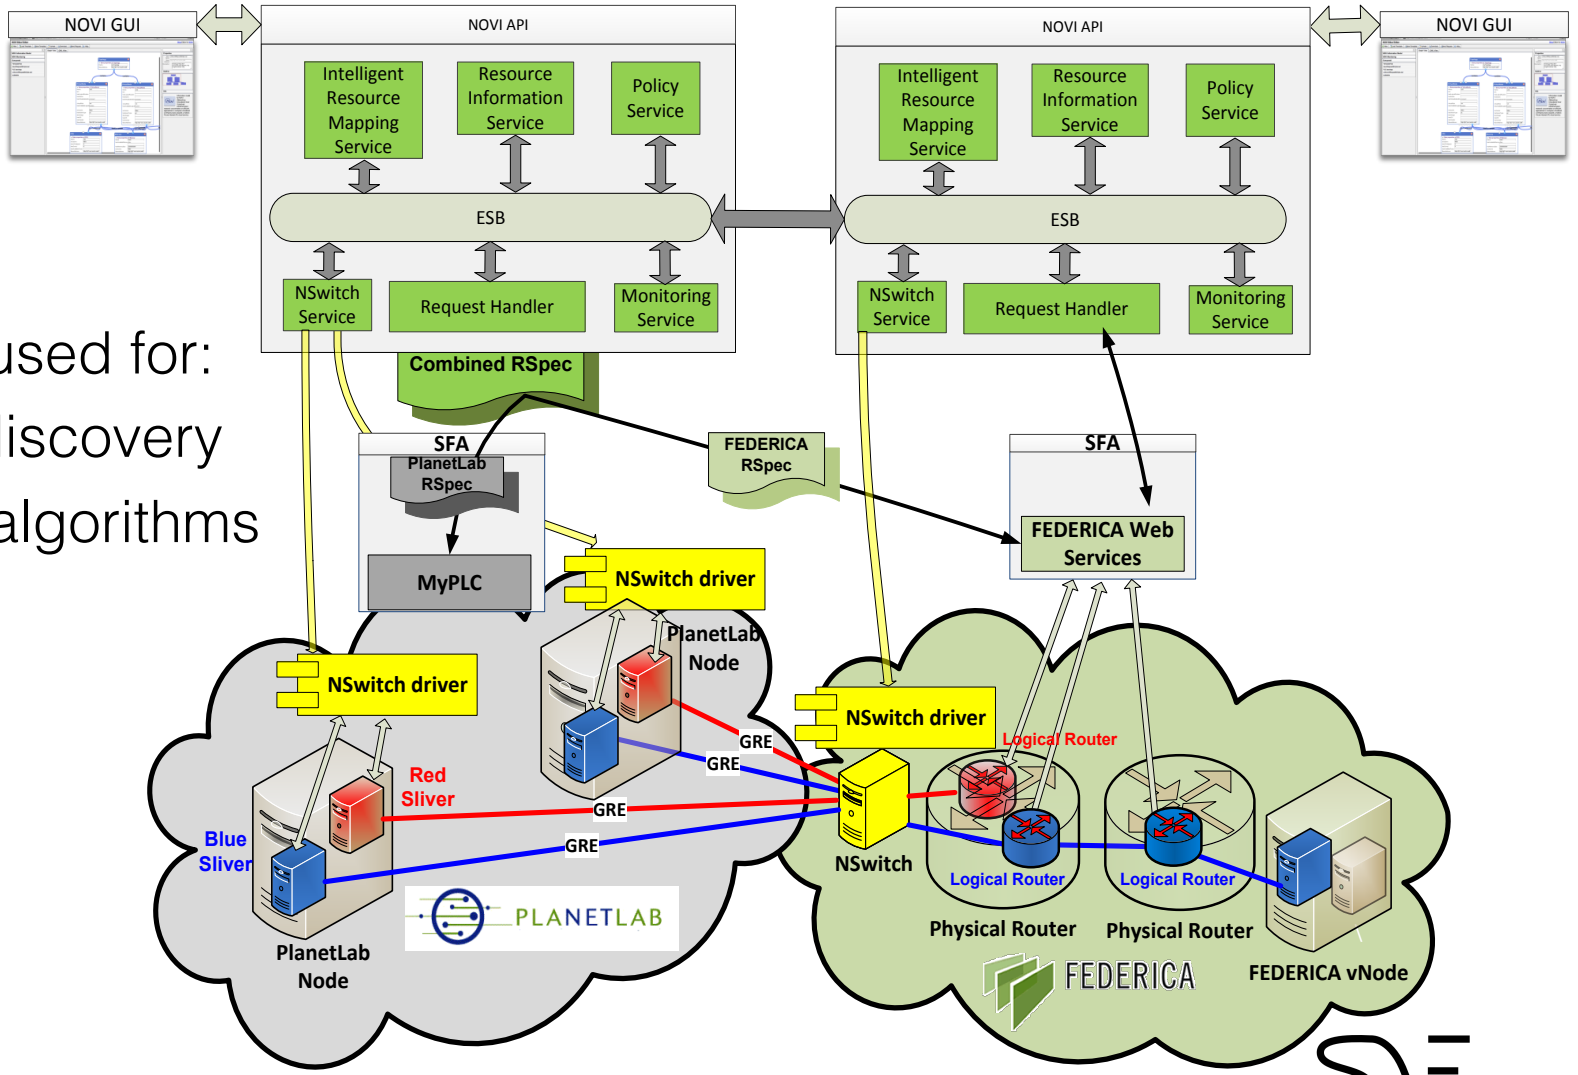

INDL use in NOVI

- Two nodes in the NOVI federation:

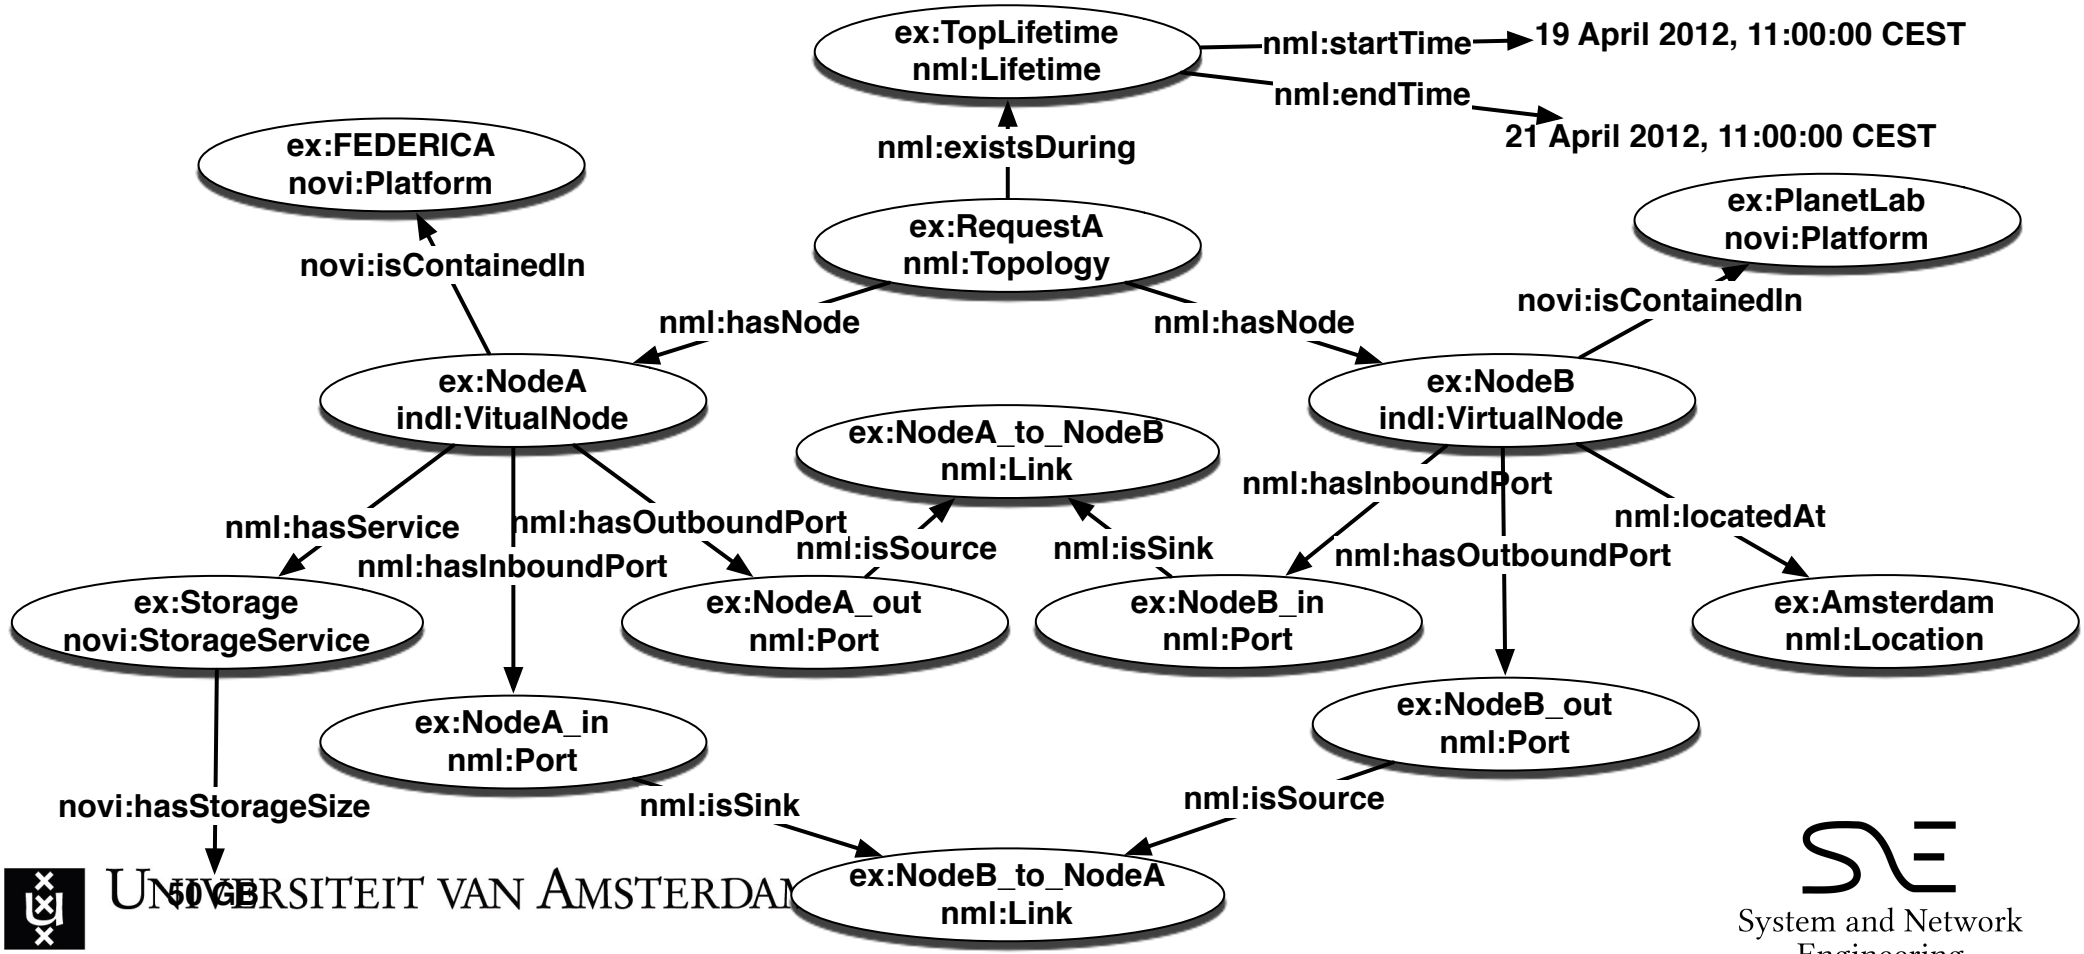

UNIVERSITEIT VAN AMSTERDAM

System and Network Engineering

INDL in Geysers

- The virtualization model:

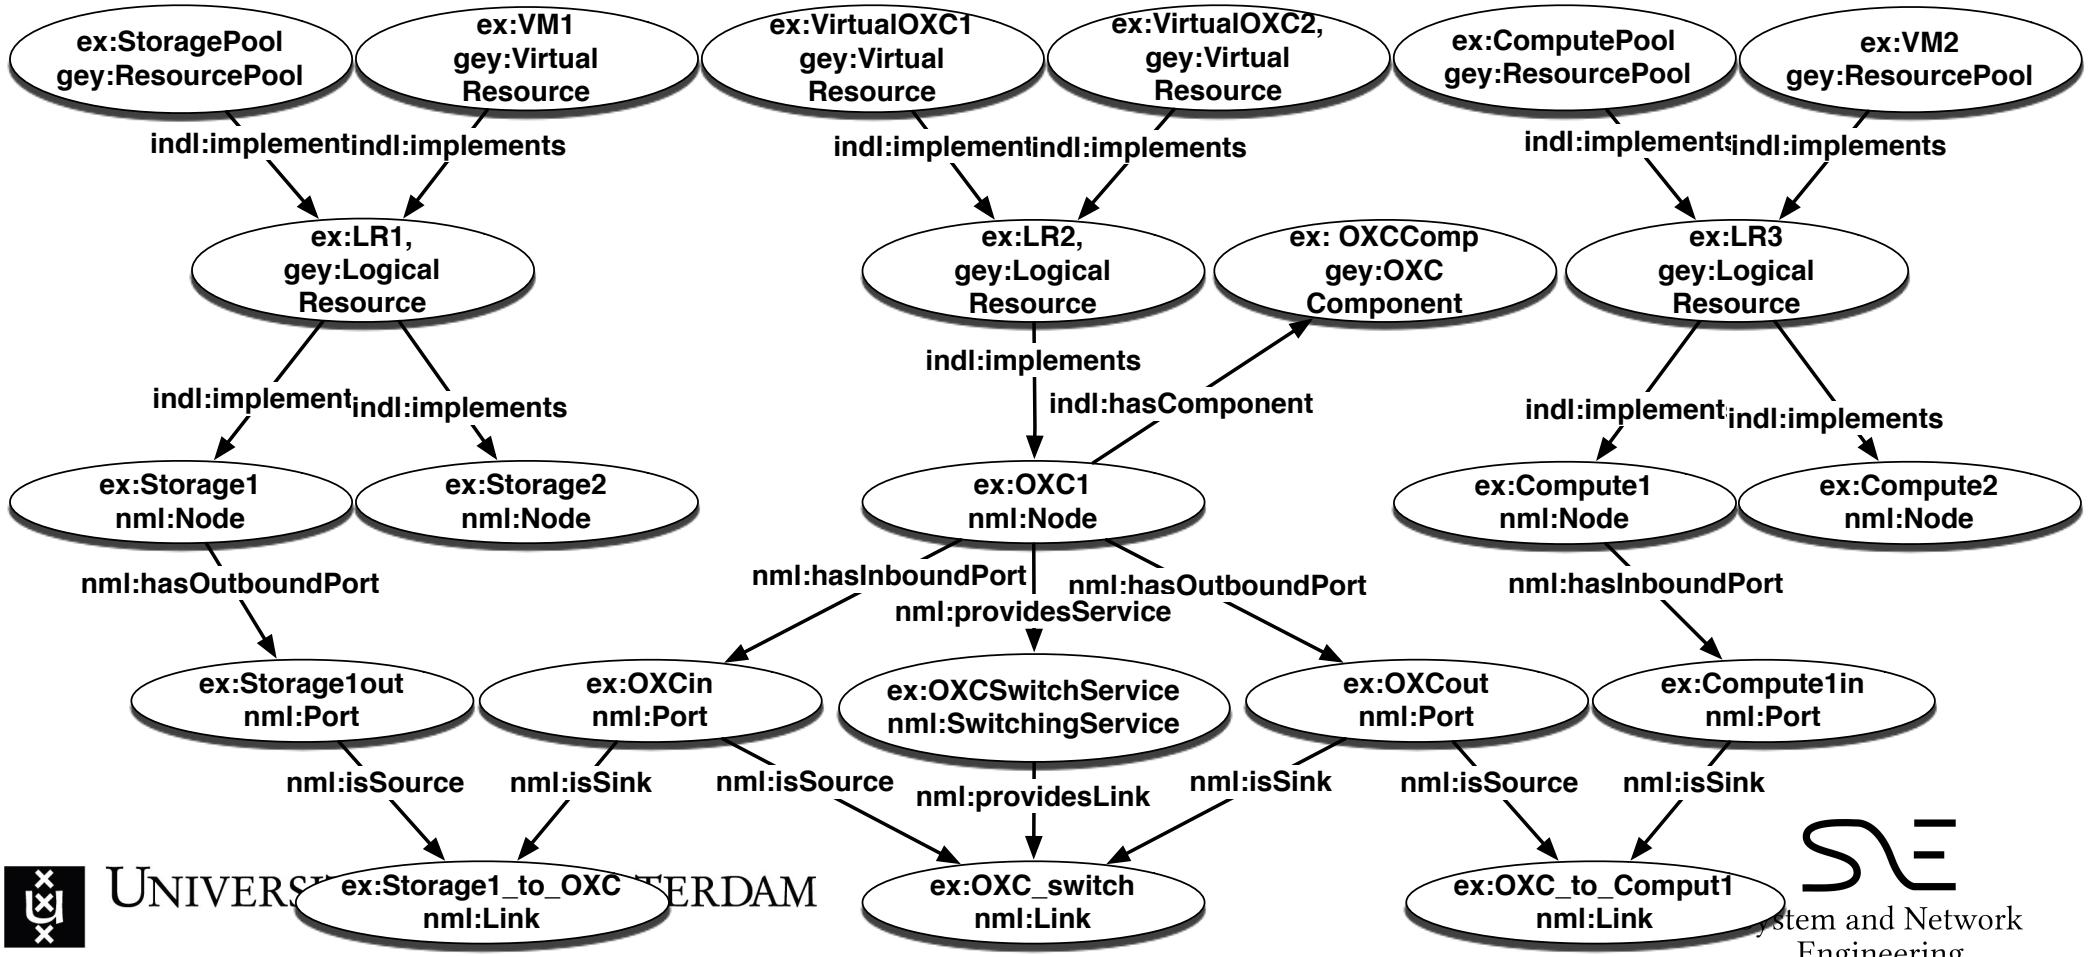

The OIntEd editor

http://indl-gui.appspot.com/

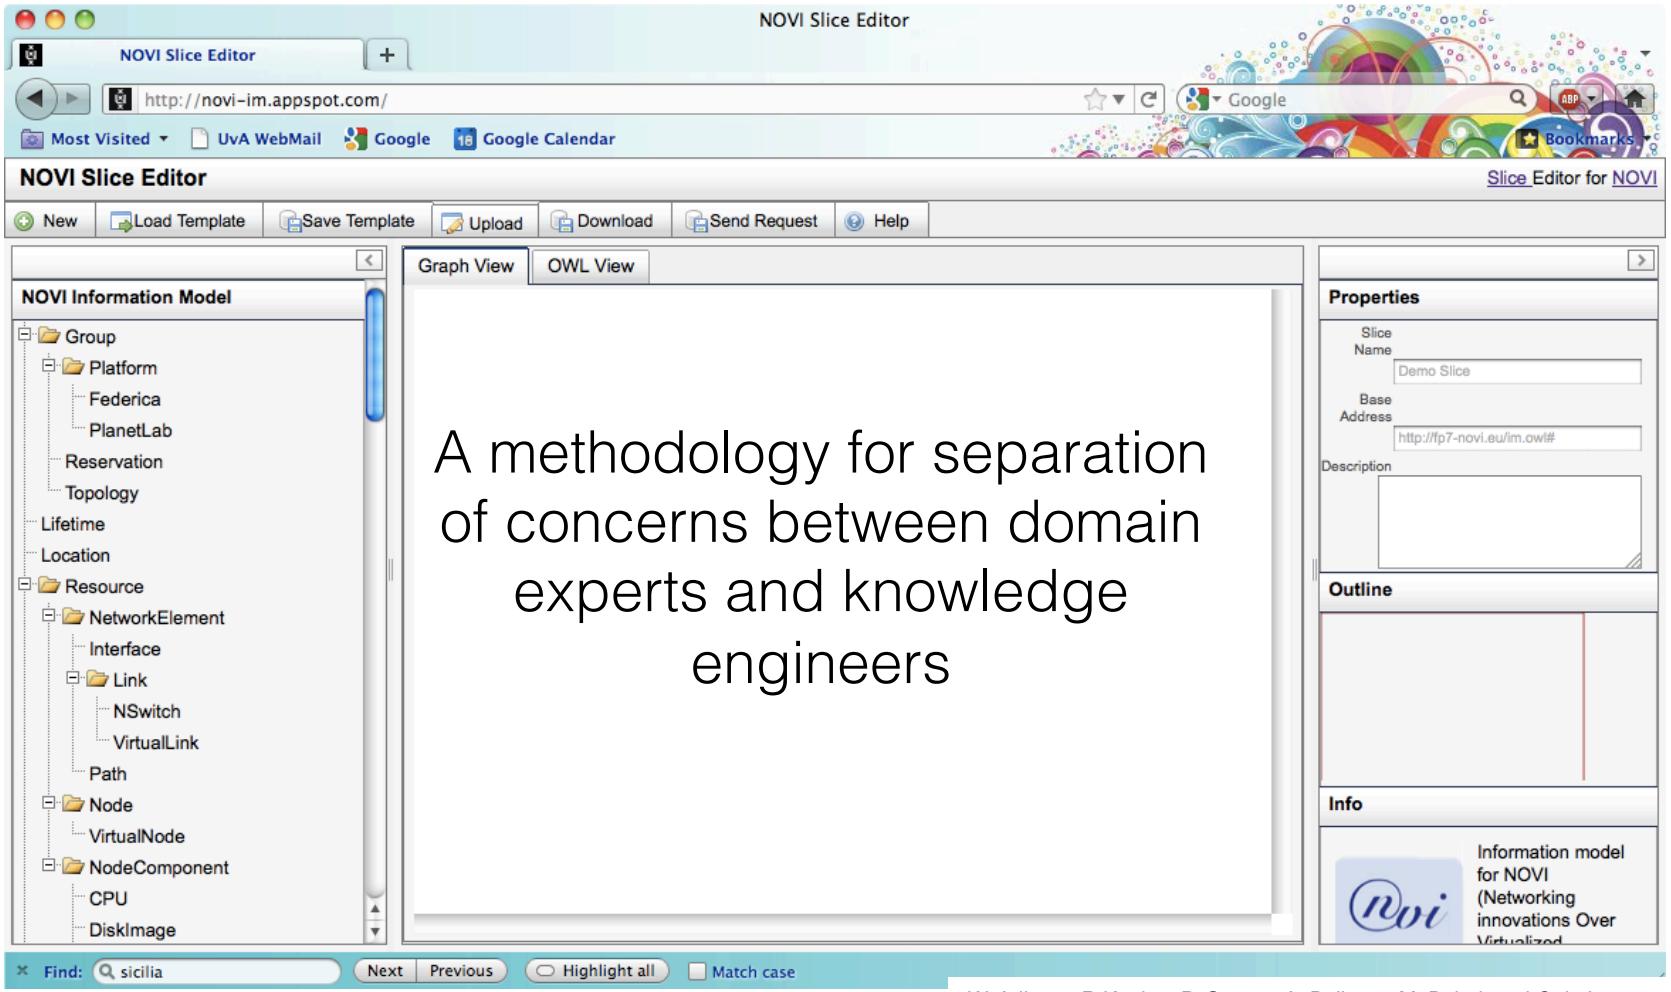

W. Adianto, R.Koning, P. Grosso, A. Belloum, M. Bubak and C.de Laat,
OIntEd: online ontology instance editor enabling a new approach to ontology development
In: Journal of "Software: Practice and Experience" 2012

NML and NSI

NML - Network Markup Language and NSI – Network Service Interface

within the OGF.

- See: "Network Markup Language Base Schema version 1"

The Network Markup Language purpose is to create a functional description of multi-layer and multi-domain networks.

It can be used for aggregated or abstracted topologies.

Under development: the Network Service Interface Topology Extensions

(Draft OGF Standard)

Publications (I)

- *Resource discovery and allocation for federated virtualized infrastructures*

C. Pittaras, C. Papagianni, A. Leivadeas, P. Grosso, J. van der Ham, S. Papavassiliou
To appear in: Future Generation Computer Systems 2014

- *A Semantic-Web Approach for Modeling Computing Infrastructures*

M.Ghijsen, J. van der Ham, P. Grosso, C.Dumitru, H. Zhu, Z. Zhao and C. de Laat
Elsevier Journal of Computers and Electrical Engineering – 2013

- *Addressing Big Data Issues in Scientific Data Infrastructure*

Y. Demchenko, P. Membrey, P. Grosso and C.de Laat
In: First International Symposium on Big Data and Data Analytics in Collaboration (BDDAC 2013). Part of The 2013 International Conference on Collaboration Technologies and Systems (CTS 2013), 2013

- *Semantic Distributed Resource Discovery for Multiple Resource Providers*

C. Pittaras. M. Ghijsen, W. Adianto, P. Grosso, J. van der Ham and C. de Laat,
In: Proceedings of the 8th International Conference on Semantics, Knowledge and Grids (SKG) 2012

- *OIntEd: online ontology instance editor enabling a new approach to ontology development*

W. Adianto, R. Koning, P. Grosso, A. Belloum, M. Bubak and C.de Laat,
In: Journal of "Software: Practice and Experience" 2012

(Network) services

“ Automated advanced capabilities to users […]: **intelligent resource mapping**, **policy-driven access and resource allocation**, **context aware resource discovery**, transparent data plane connectivity and monitoring of combined user slices and substrate resources across domains”

NewQOSPlanner

The NSI – Network Service Interface – creates on the fly connections between domains.

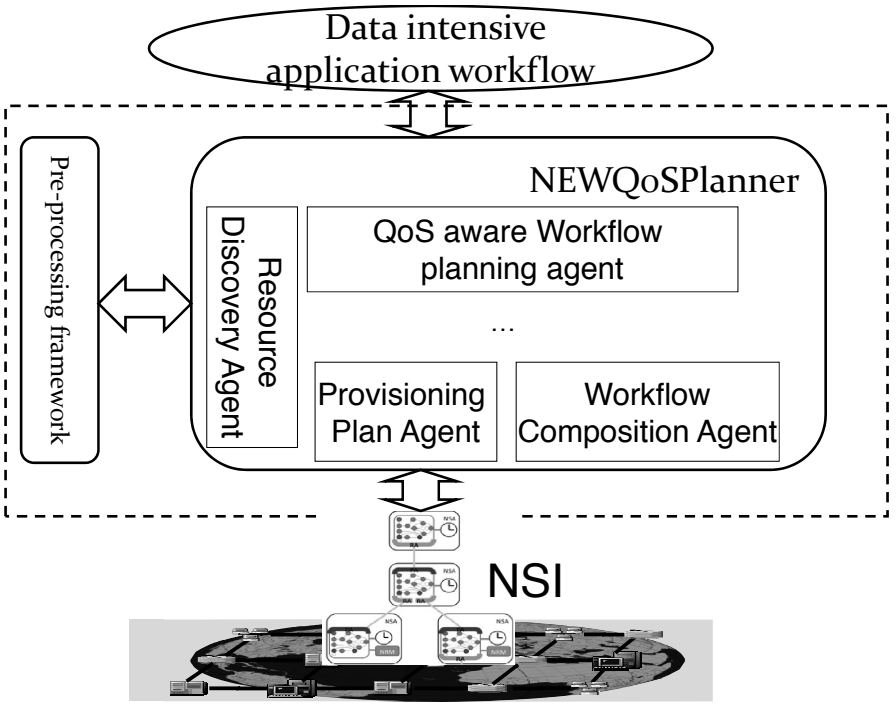

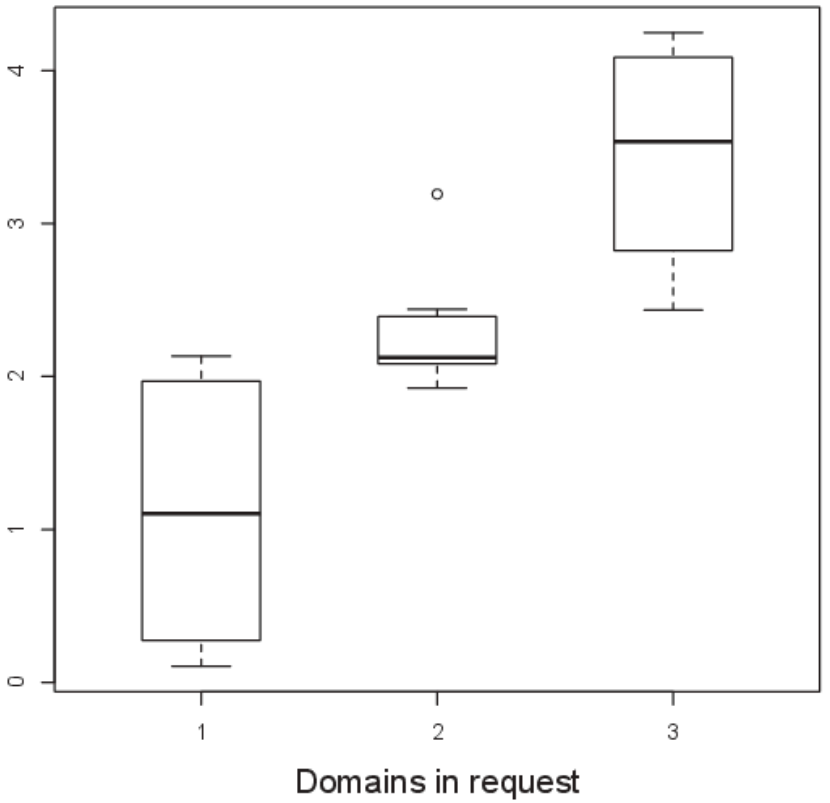

Z. Zhao, J. v/d Ham, A. Taal, R. Koning, P. Grosso and C. de Laat
Planning data intensive workflows on inter-domain resources using the Network Service Interface (NSI)
In: WORKS 2012

HyperFlow

Encoding times improve as the end nodes are connected via dynamic lightpaths

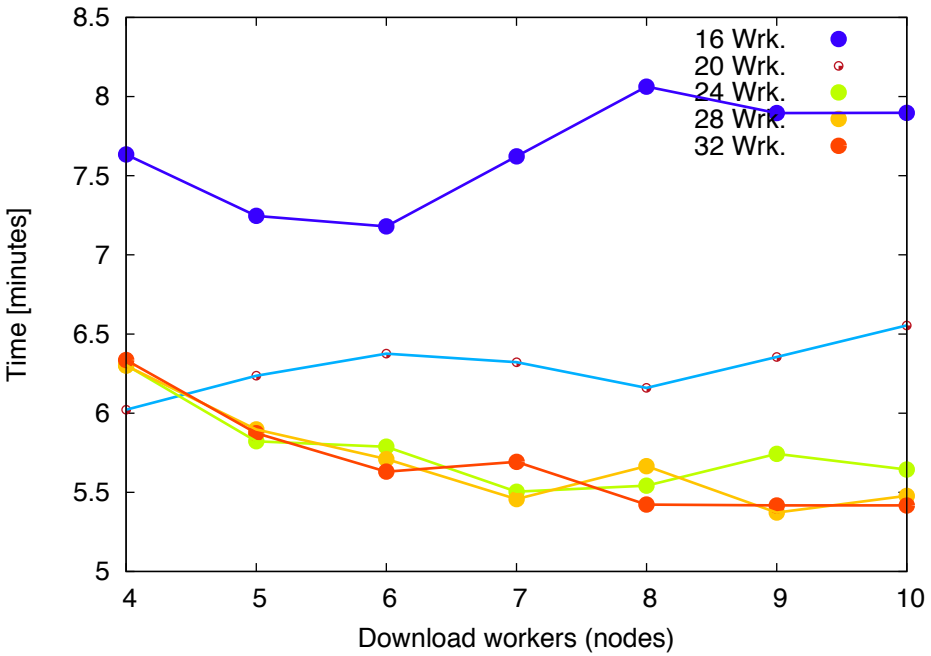

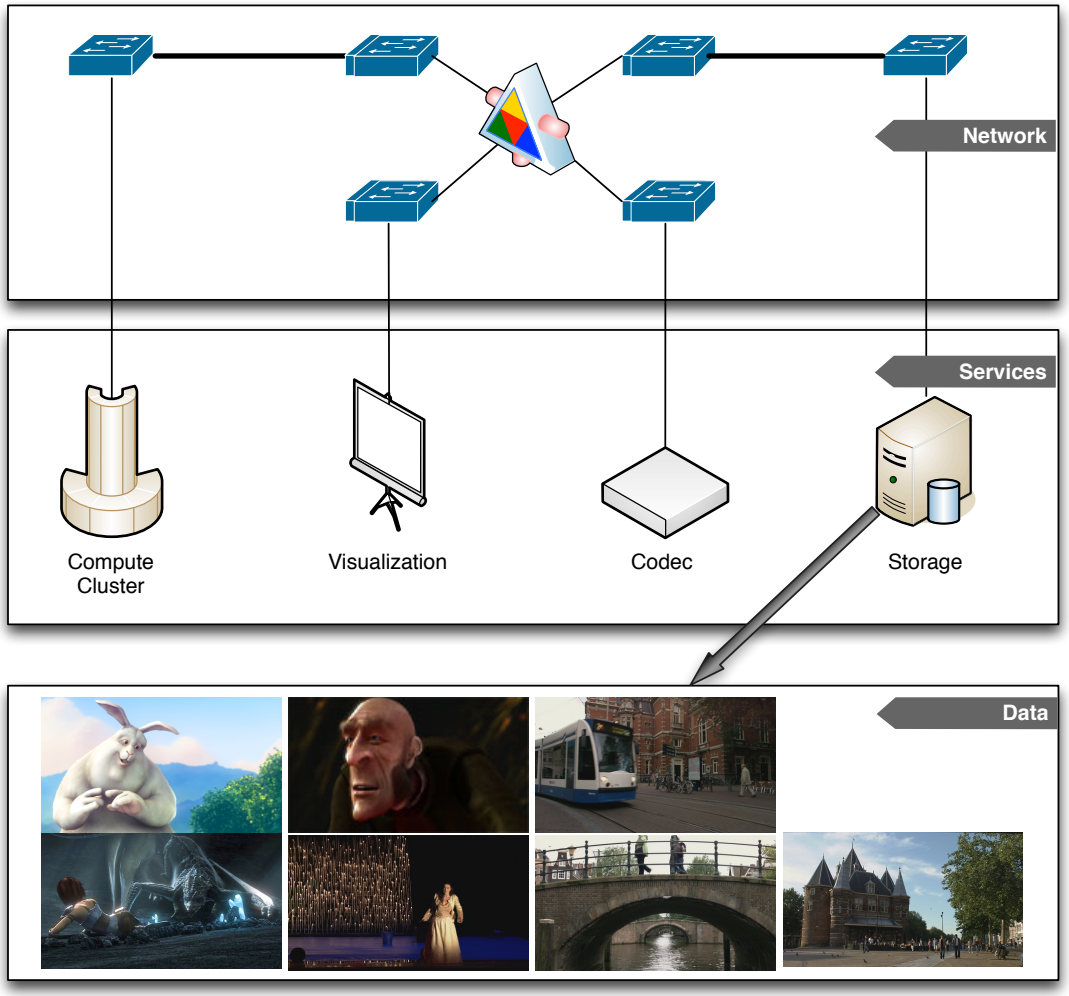

C. Dumitru, Z. Zhao, P. Grosso and C. de Laat
HybridFlow: Towards Intelligent Video Delivery and Processing Over Hybrid Infrastructures
(In CTS 2013))

CineGrid

- http://www.cinegrid.org
- http://cgdev.uvalight.nl/home/

A community that uses our tools!

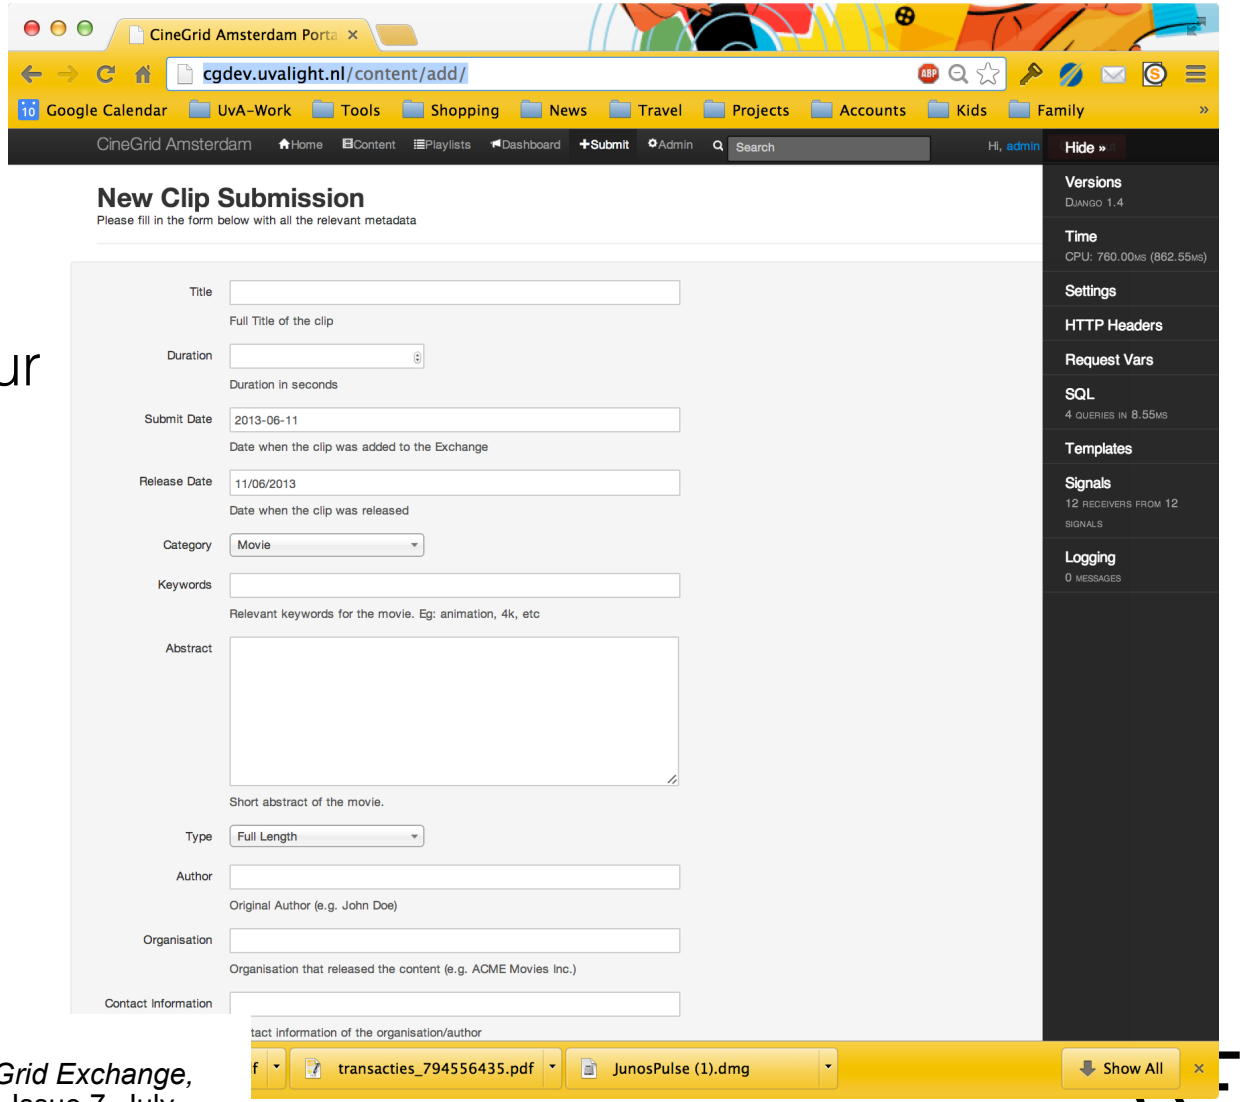

Ralph Koning, Paola Grosso and Cees de Laat.
Using ontologies for resource description in the CineGrid Exchange,
In: Future Generation Computer Systems,Volume 27, Issue 7, July 2011

A queueing model approach

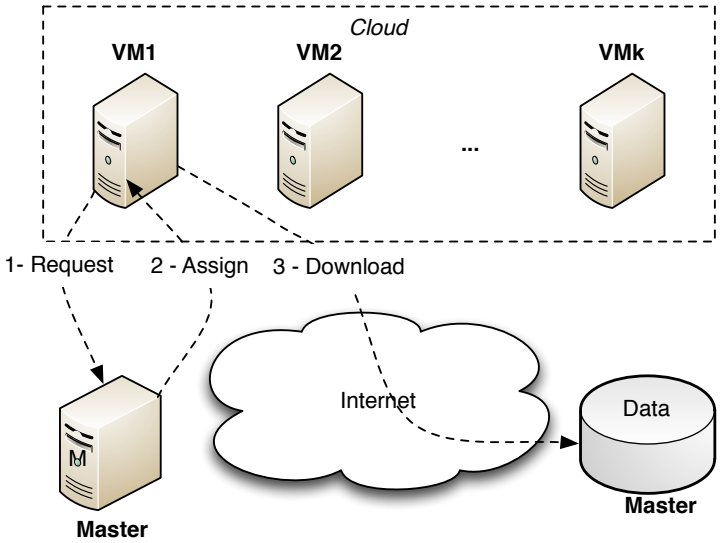

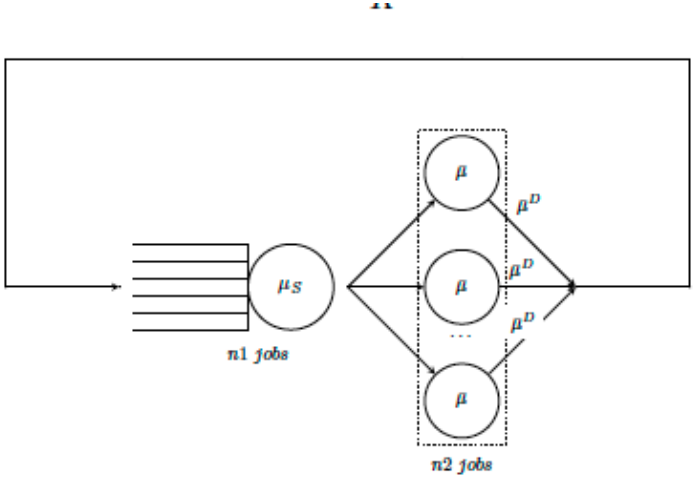

Pareto fronts

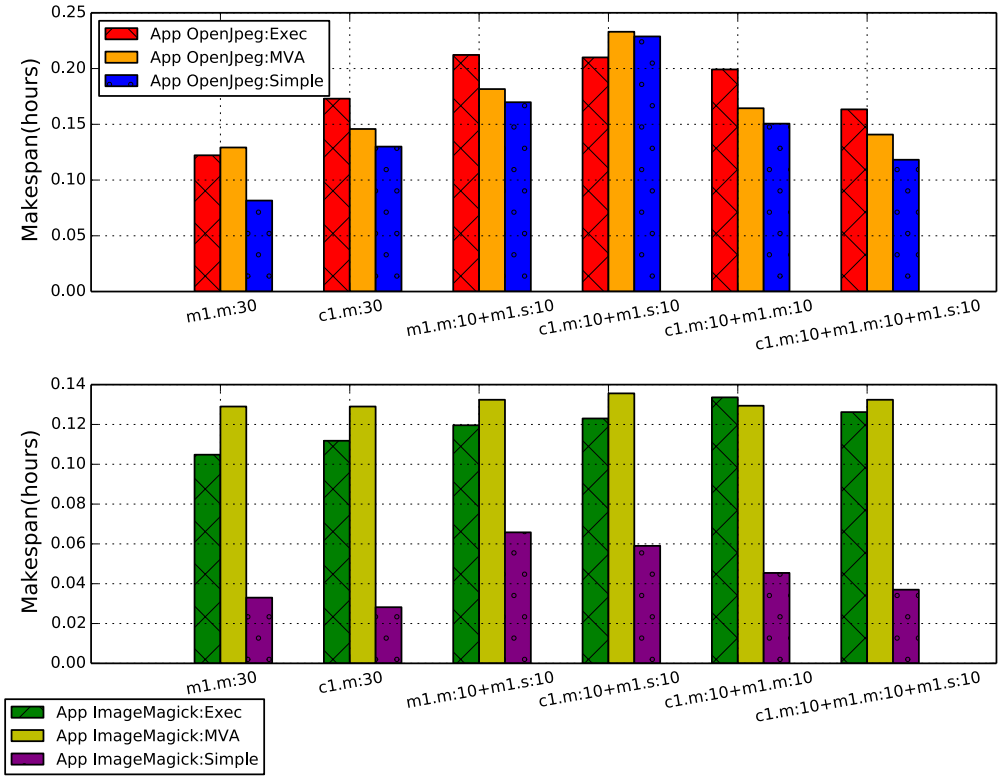

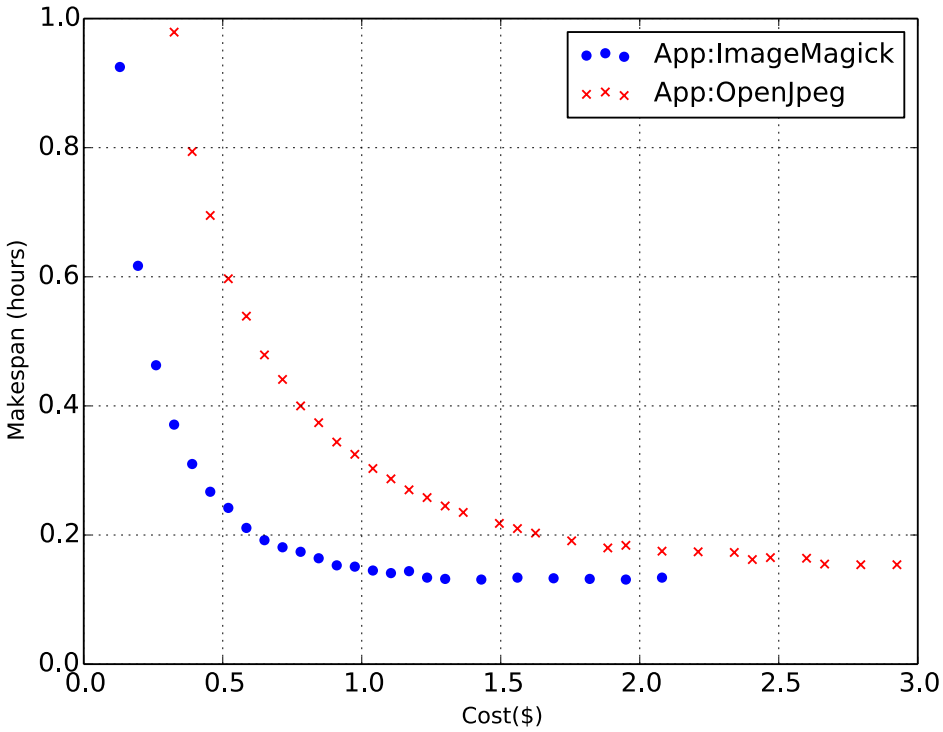

Publications (II)

- *A Queueing Theory Approach to Pareto Optimal Bags-of-Tasks Scheduling on Clouds*

C. Dumitru, A. Oprescu, M. Zivkovic, R. v/d Mei, P. Grosso and C.de Laat
Submitted to Europar2014

- *HybridFlow: Towards intelligent video delivery and processing over hybrid infrastructures*

C.Dumitru, Z. Zhao, P. Grosso and C.de Laat
In: 2013 International Conference on Collaboration Technologies and Systems (CTS), 2013

- *An agent based network resource planner for workflow applications*

Z. Zhao, P. Grosso, J. van der Ham, R. Koning and C.de Laat.
In: IOS Press Journal Multiagent and Grid Systems, Volume 7, Issue 6, Jan 2011, Pages 187-202.

- *Using ontologies for resource description in the CineGrid Exchange*

R. Koning, P. Grosso and C.de Laat.
In: Future Generation Computer Systems, Volume 27, Issue 7, July 2011, Pages 960-965

Greening

"….*the cloud model provides also* **benefits** *from the* **environmental perspective**…."

Green scheduling

Network infrastructures

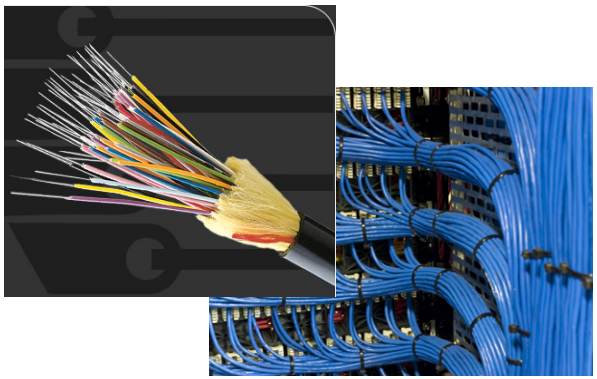

Green energy sources

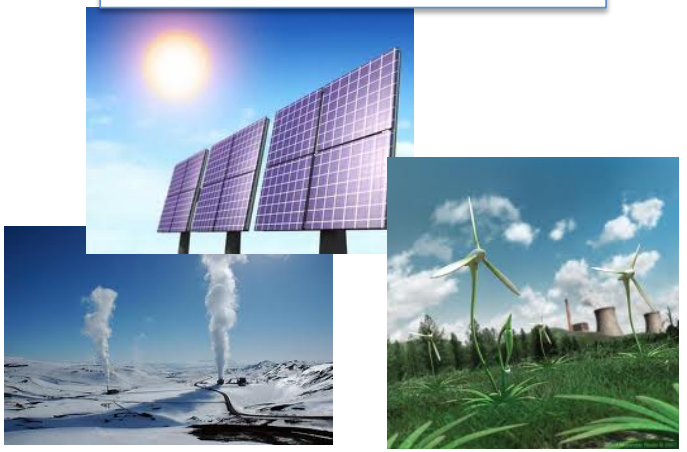

CO_2 footprint;
Energy needed and lost

Bits to energy

CO_2 footprint;
Energy needed and lost

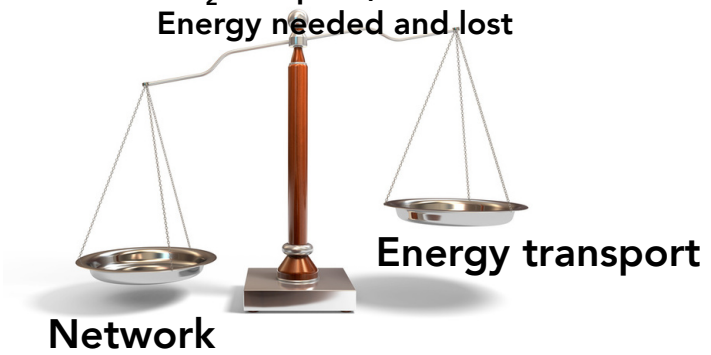

Energy to bits

Energy saving in clouds

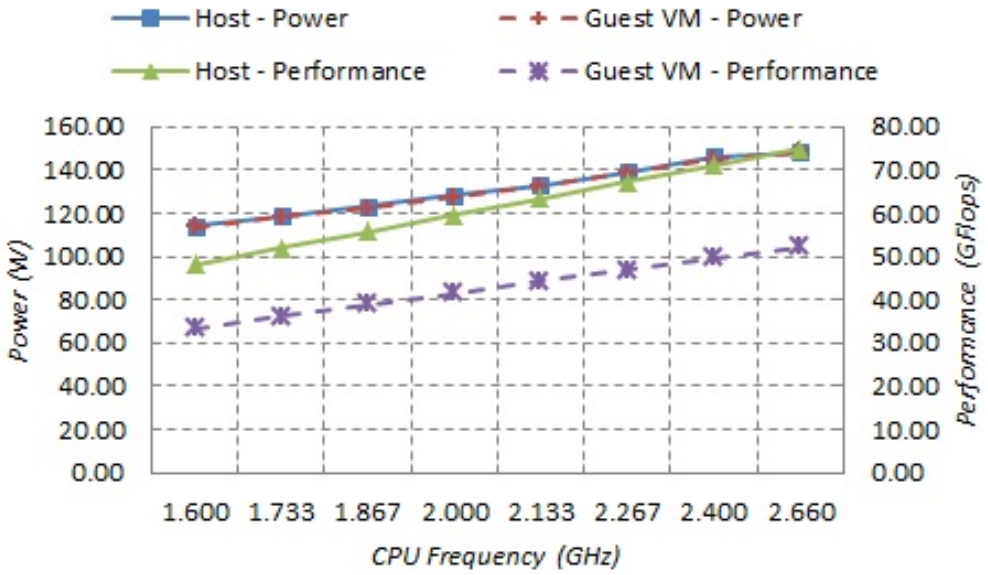

Quantifying the energy performance of VMs is the first step toward energy-aware job scheduling.

Q. Chen, P. Grosso, K. van der Veldt, C. de Laat, R. Hofman and H.Bal.
Profiling energy consumption of VMs for green cloud computing
In: International Conference on Cloud and Green Computing (CGC2011), Sydney December 2011

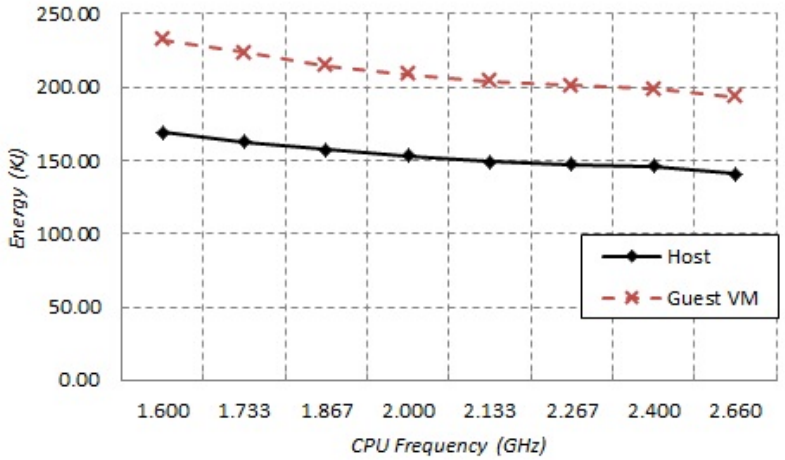

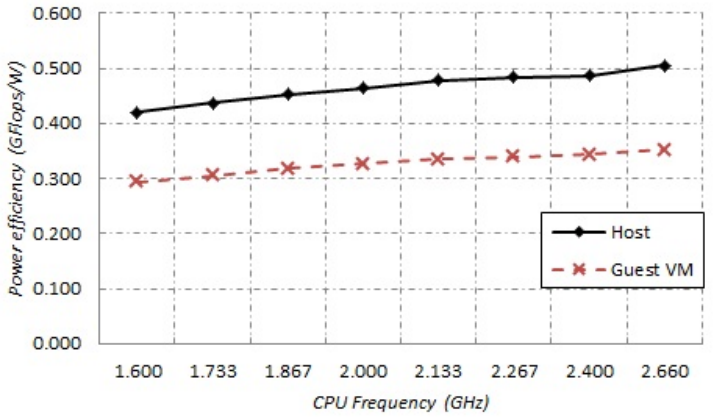

COMMIT/

Energy Efficient Ethernet (802.3az)

Power savings techniques in hardware can be leveraged in architecturing communication patterns in data centra

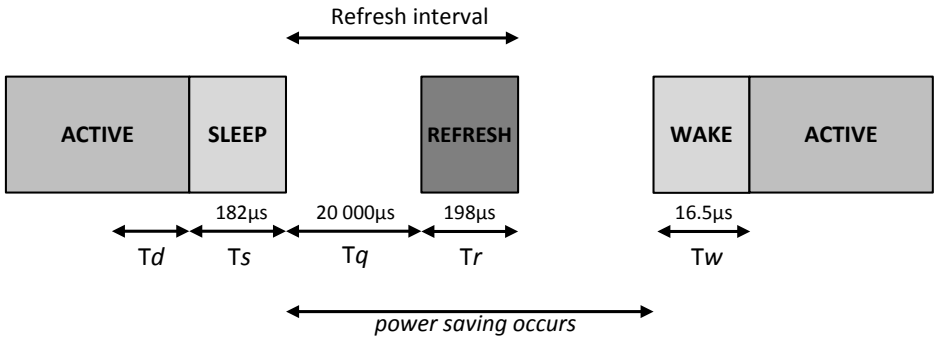

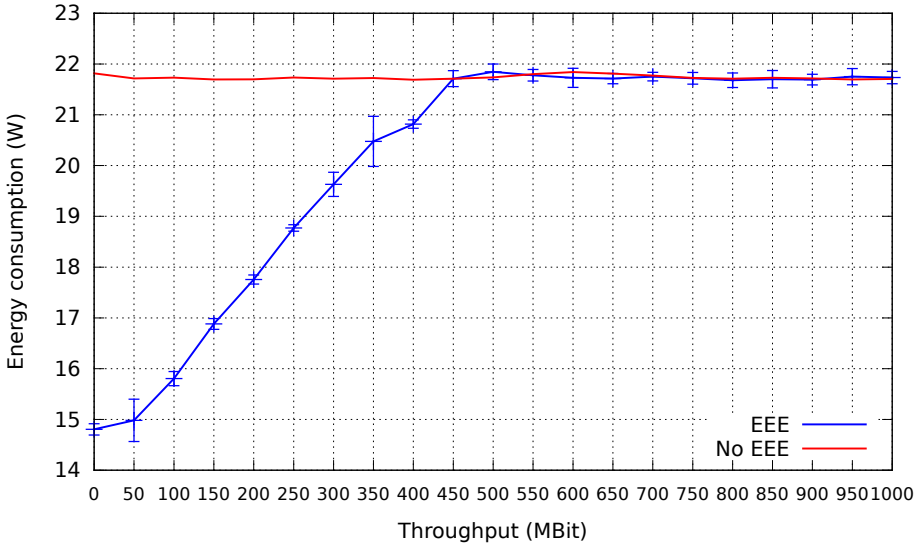

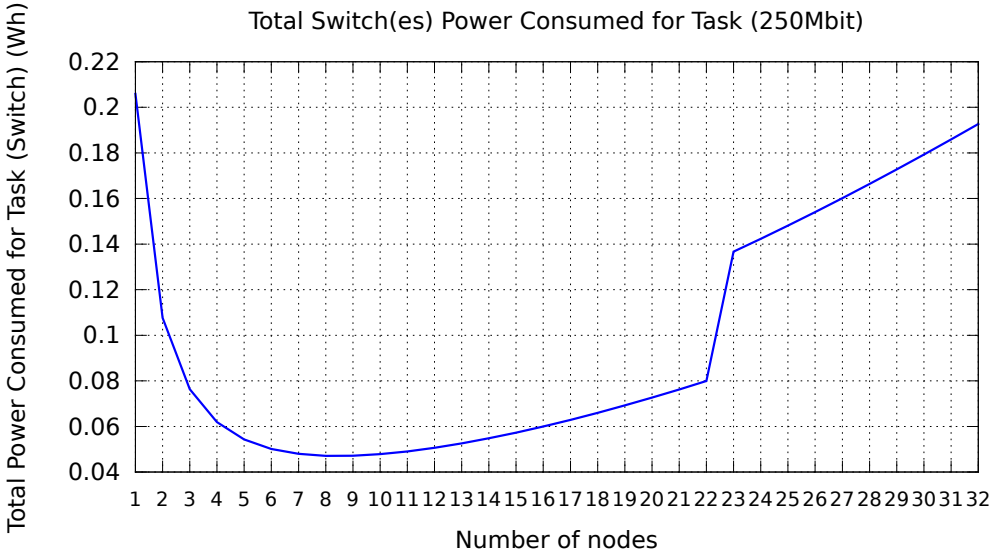

D. Pavlov and J. Soert and P. Grosso and Z. Zhao and K. van der Veldt and H. Zhu and C.de Laat
Towards energy efficient data intensive computing using IEEE 802.3az
In: DISCS 2012 workshop - Nov 2012

Networks and CO2

- Take a network (Esnet, working on using SURFnet data)
- Define the traffic model running on it
- Use the energy monitoring information and energy costs data
- Compare path selection strategies : shortest, cheapest and greenest

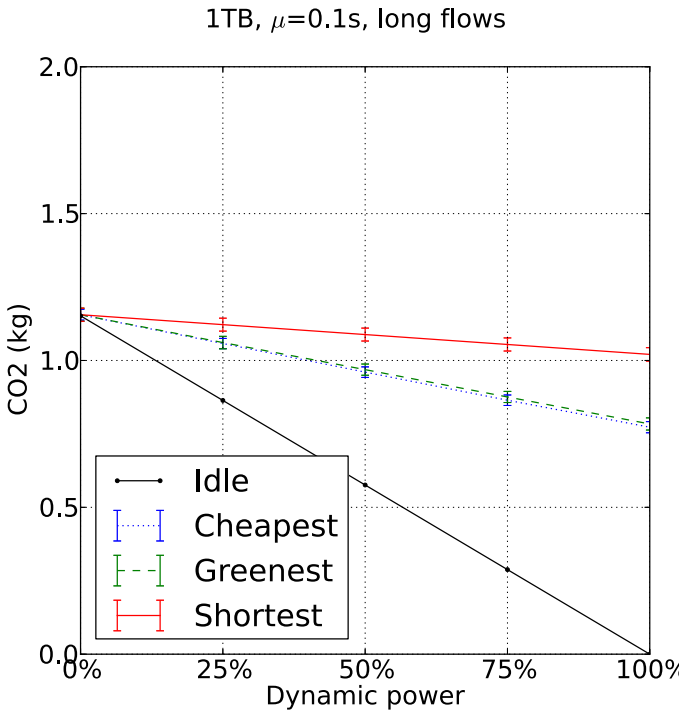

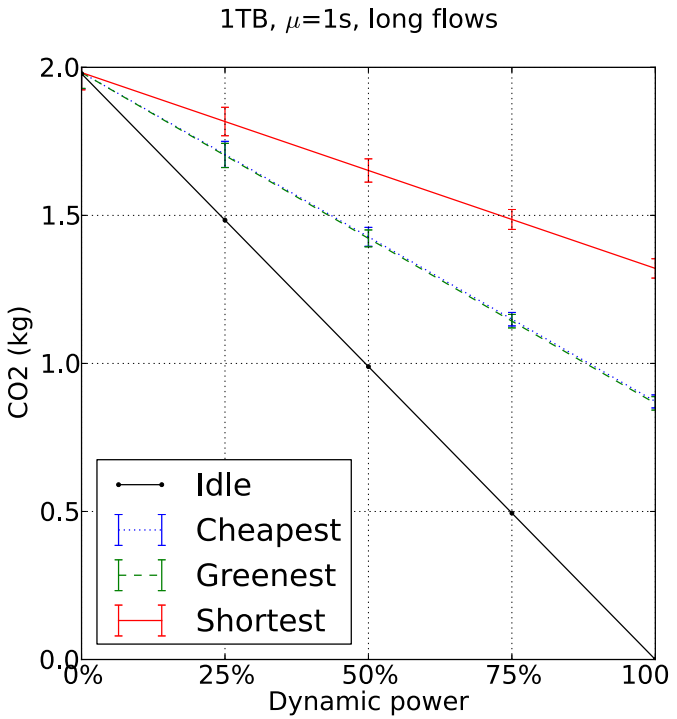

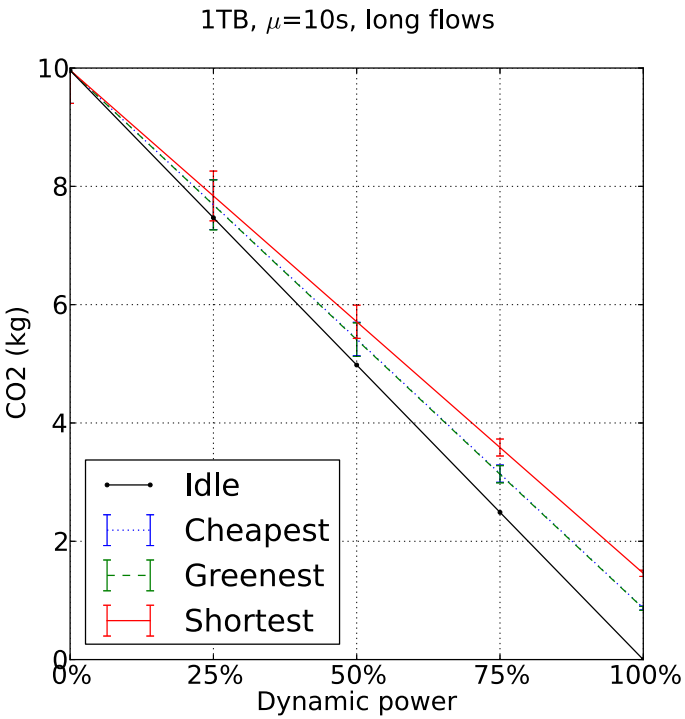

Results

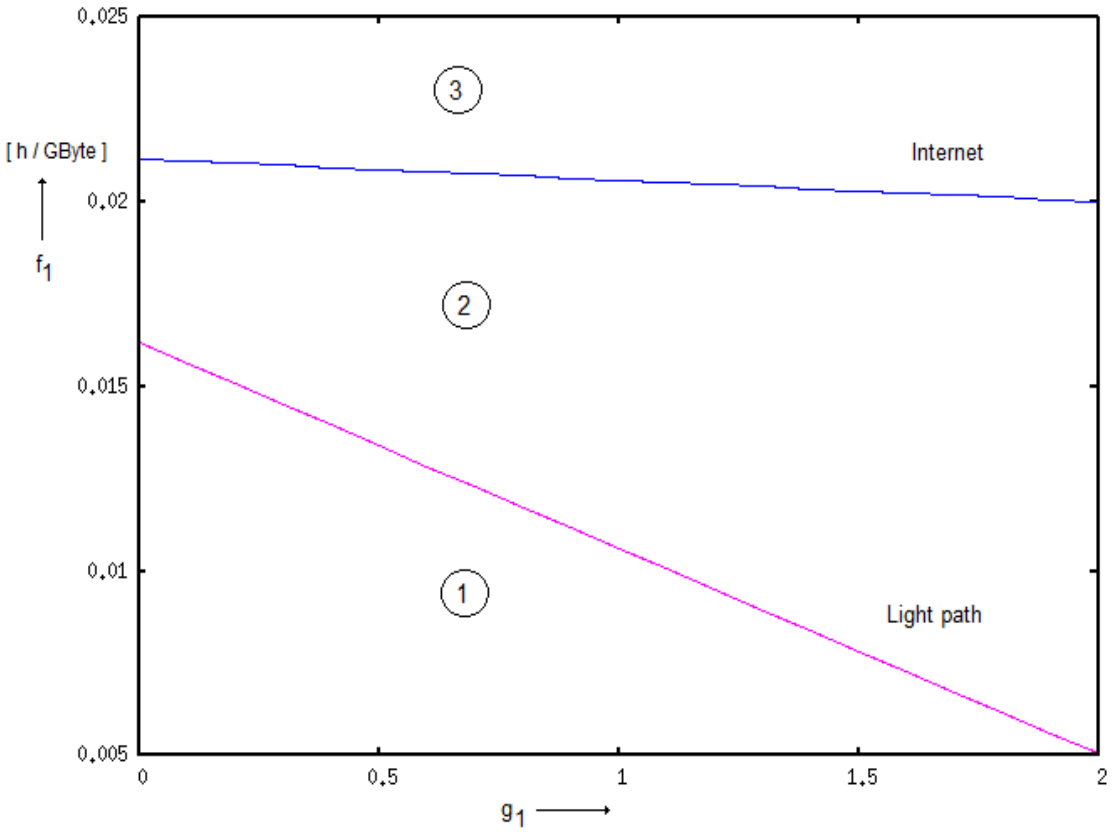

In region 1 the task should be performed locally, independently of the type of transport network.

In region 2 the task can be performed remotely provided that the connection is a light path.

In region 3 the task should be done remotely for both types of transport networks.

Given different network paths we can identify decision boundaries as function of the task complexity.

Publications (III)

- *Storage to Energy: modeling the carbon emission of storage task offloading between data centers*

A.Taal, D. Drupsteen, M. Makkes and P.Grosso
In: CCNC conference(Las Vegas Jan. 2014)

- *A decision framework for placement of applications in clouds that minimizes their carbon footprint*

M. Makkes, A. Taal, A. Osseyran and P. Grosso
In: Journal of Cloud Computing: Advances, Systems and Applications 2013, Vol.2

- *EKB: semantic information system for energy-aware monitoring in distributed infrastructures*

H. Zhu, K. v/d Veldt, P. Grosso, X. Liao and C.de Laat
In: IEEE International Conference on Cloud and Green Computing (CGC 2013) - Sep. 2013

- *Towards energy efficient data intensive computing using IEEE 802.3az*

D. Pavlov, J. Soert, P. Grosso, Z. Zhao, K. v/d Veldt, H. Zhu and C.de Laat
In: DISCS 2012 workshop - Nov 2012

Conclusions

- There is a clear need for a common information model that allows intelligent information exchange.

INDL, NML

- Such models support novel scenarios that will empower BigData transport and processing.

BaTS, CineGrid, NetQoSPlanner

- BigData can be ‘green’.