

GreenSonar

Cees de Laat & Paola Grosso & Inder Monga

EU
COMMIT
UvA

DoE
LBNL

NWO
PID/EFRO


SURFnet
NLESC

TNO
NWO/ncf



Need for GreenIT

Global Average Temperature and Carbon Dioxide Concentrations, 1880 - 2004



Data Source Temperature: ftp://ftp.ncdc.noaa.gov/pub/data/anomalies/annual_land_and_ocean.ts

Data Source CO2 (Siple Ice Cores): <http://cdiac.esd.ornl.gov/ftp/trends/co2/siple2.013>

Data Source CO2 (Mauna Loa): <http://cdiac.esd.ornl.gov/ftp/trends/co2/maunaloa.co2>

Graphic Design: Michael Ernst, The Woods Hole Research Center





Turn Green Tech into Greenbacks

IT Certifications for Jobs That Make a Difference



Uptime Institute Accredited Tier Designer

The Uptime Institute has long been a proponent for green data center design and implementation. Its certification course on data center design embeds green principles into the curriculum.



SUSTAINABILITY
Your Career



Greening the Processing System

Positive proof of global warming.



ECO-Scheduling



The Ten Problems with the Internet

1. **Energy Efficient Communication**
2. Separation of Identity and Address
3. Location Awareness
4. Explicit Support for Client-Server Traffic and Distributed Services
5. Person-to-Person Communication
6. Security
7. Control, Management, and Data Plane separation
8. Isolation
9. Symmetric/Asymmetric Protocols
10. Quality of Service

Nice to have:

- Global Routing with Local Control of Naming and Addressing
- Real Time Services
- Cross-Layer Communication
- Multicast
- Receiver Control
- Support for Data Aggregation and Transformation
- Support for Streaming Data
- Virtualization

GreenSonar

- Sustainability requires Green-IT in all kinds of resources, compute, network, storage, sensor, ...
 - > System wide approach!
- Measure to know; information needed for smart infrastructure.
- Basic idea: been there done that in Networking!
- Why not apply iNDL/NML & PerfSonar methods to provide Green & Energy information?
- Need also application information.
- Big hurdle: energy metrics on heterogeneous resources
- OGF BOF/Charter meeting at OGF36, Chicago.


Approach

- Define metrics
- Extend xDL with GreenDL
- GreenSonar
- Measure & publish
- Build data base -> information base
- Machine Learn good/bad practices
- Reason & Select & Schedule

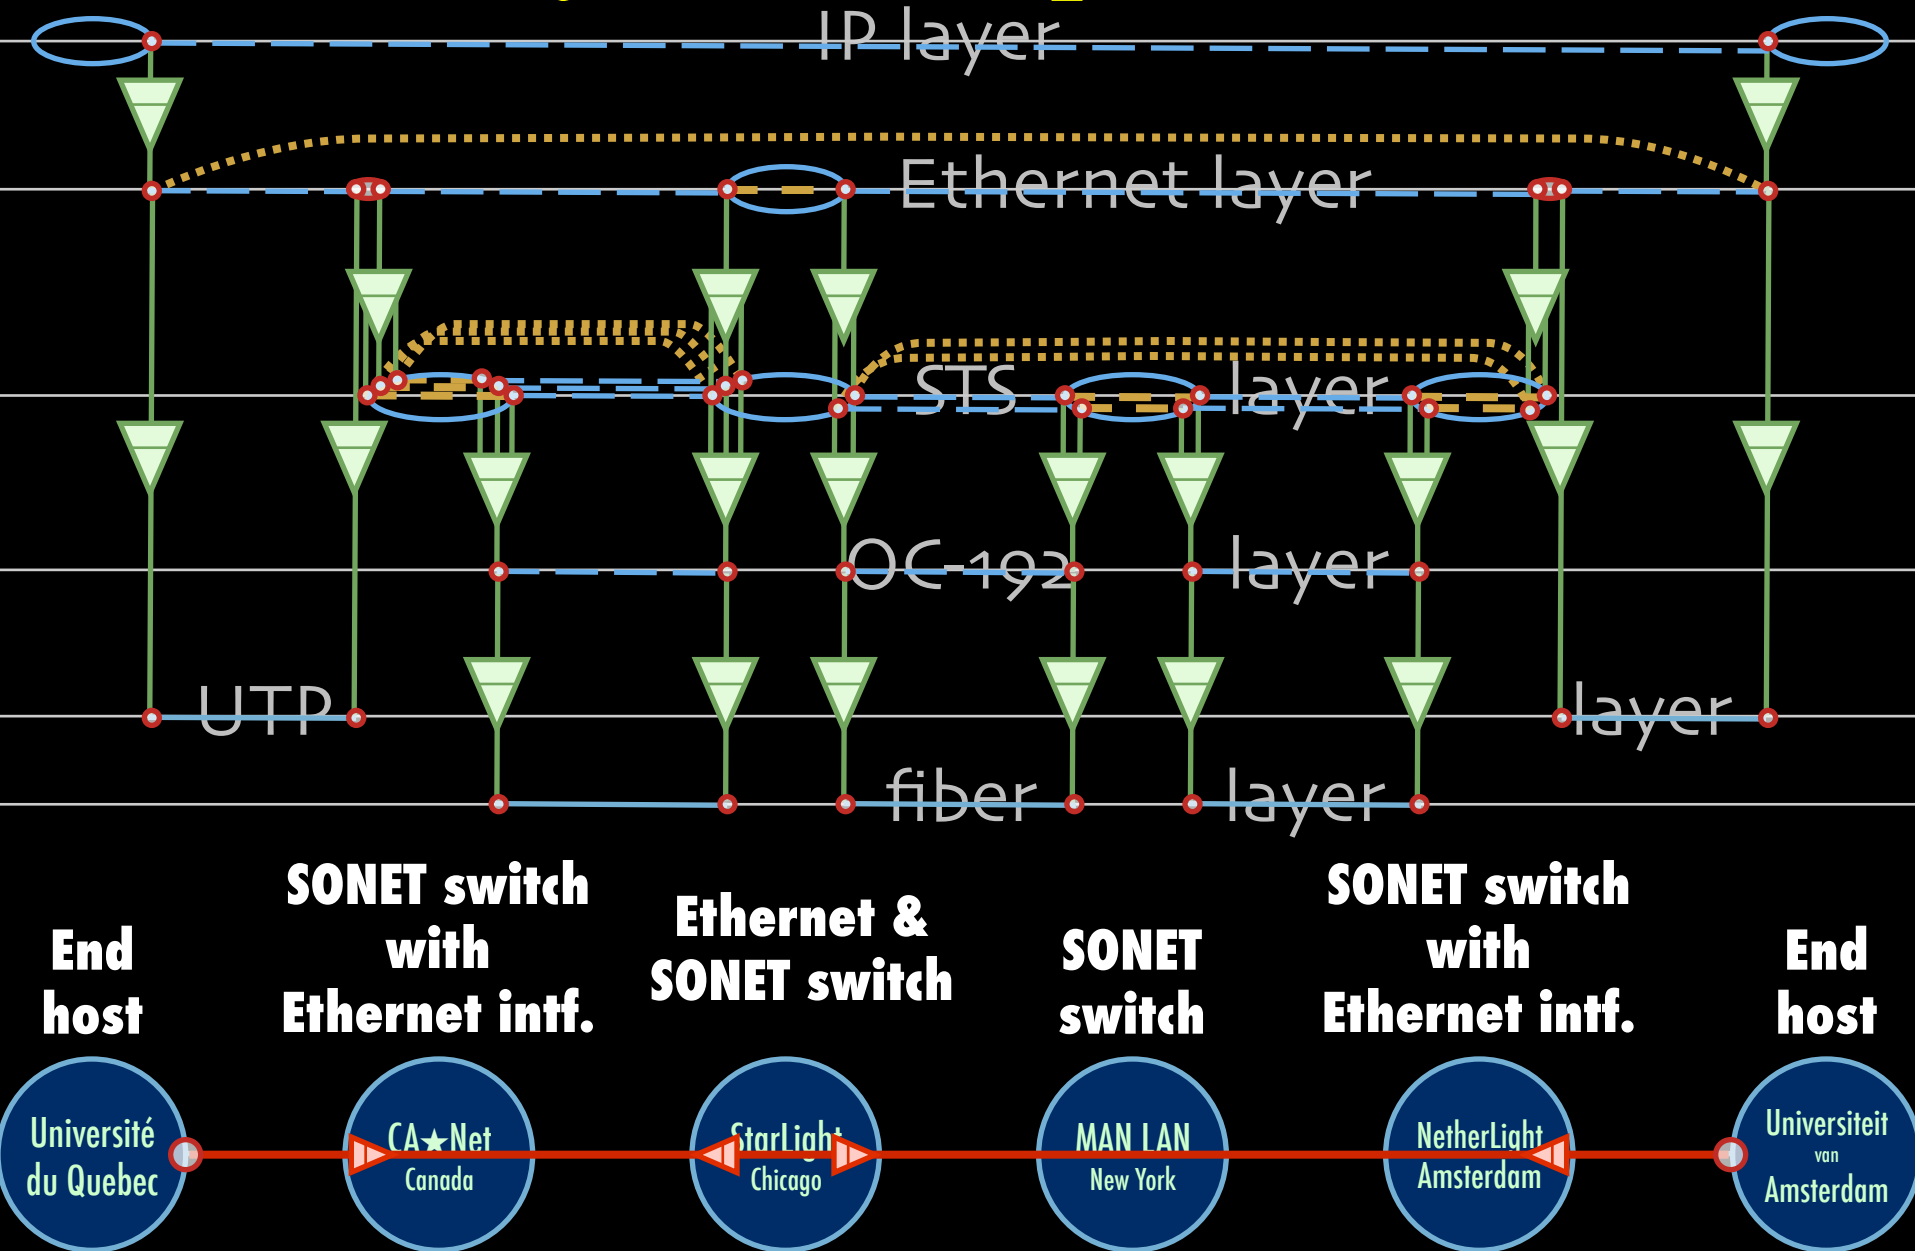
LinkedIn for Infrastructure



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



Multi-layer descriptions in NDL



CdL

Applications and Networks become aware of each other!

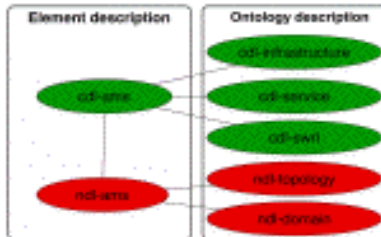
CineGrid Description Language

CineGrid is an initiative to facilitate the exchange, storage and display of high-quality digital media.

The CineGrid Description Language (CDL) describes CineGrid resources. Streaming, display and storage components are organized in a hierarchical way.

CDL has bindings to the NDL ontology that enables descriptions of network components and their interconnections.

With CDL we can reason on the CineGrid infrastructure and its services.



SQWRL is used to query the Ontology.

Which CineGrid nodes are directly connected?



```
cdl:hasElements(?node1, ?host1) ^
ndl:topo:hasInterface(?host1, ?IF1) ^ ndl:
topo:connectedTo(?IF1, ?IF2) ^
ndl:topo:hasInterface(?host2, ?IF2) ^
cdl:hasElements(?node2, ?host2) ->
sqwrl:select(?node1, ?node2)
```

cdl-ams.owl

cdl-ams:Amsterdam cdl-ams:Prague
cdl-ams:Prague cdl-ams:Amsterdam

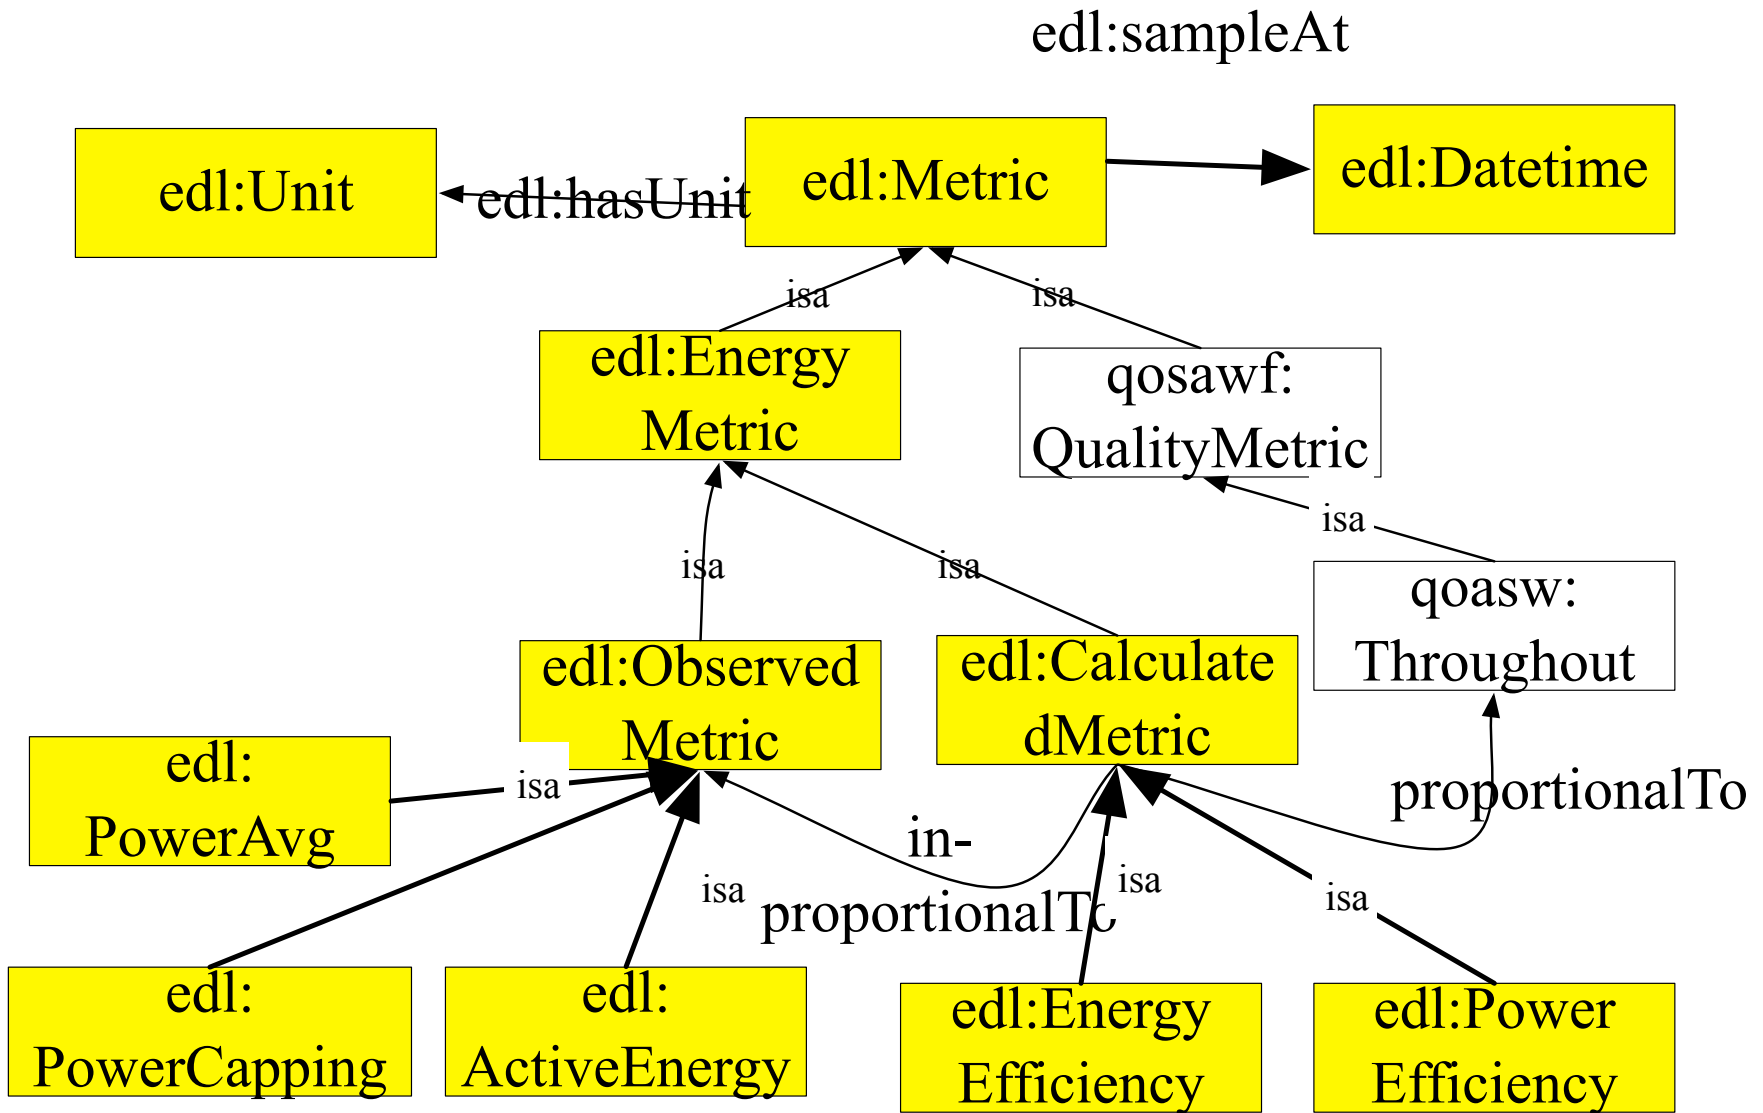
UML representation of CDL




CDL links to NDL using the **owl:SameAs** property. CDL defines the services, NDL the network interfaces and links. The combination of the two ontologies identifies the host pairs that support matching services via existing network connections.



EdL (have already CdL)








I want to



“Show Big Bug Bunny in 4K on my Tiled Display using green Infrastructure”

- Big Bugs Bunny can be on multiple servers on the Internet.
- Movie may need processing / recoding to get to 4K for Tiled Display.
- Needs deterministic Green infrastructure for Quality of Experience.
- Consumer / Scientist does not want to know the underlying details.
➔ His refrigerator also just works.





Why GreenSONAR?

- ICT today is responsible for ~2% of the global carbon emissions, and is expected to overtake the airline industry in the near future.
- There is no systematic approach to measuring where that energy is spent, what type of energy is spent, or if it is spent well.
- Analyzing the data may lead to better scheduling and routing algorithms geared towards increasing system-wide energy efficiency.

What is GreenSONAR?

- Proposed architecture to monitor and give insight into the energy efficiency of ICT infrastructure.
- Holistic approach.
- Requires monitoring both power consumption and performance.
- Potentially appealing to researchers, networking engineers, system administrators, maybe even CEOs, politicians.
 - Potential to save money, chart impact of ICT on environment, help make decisions on where and how to apply green energy to ICT.

What is PerfSONAR?

- Service-oriented architecture to monitor performance of networking infrastructure.
- Specifically meant to make life easier for networking engineers.
- Is the PerfSONAR architecture suitable for GreenSONAR?

Holistic approach

- Have to look at all aspects of ICT infrastructure, not only networking.
- How easily could each aspect be monitored with GreenSONAR/PerfSONAR?
 - Networking: Relatively straightforward.
 - Grids (and supercomputers): Difficult.
 - Virtualized environments: Not possible?

GreenSONAR vs Networking

- Relatively straightforward for the following reasons:
 - Networking elements have a single well-defined purpose.
 - Performance metrics are always the same.
 - Means correlating performance and power metrics is straightforward.

GreenSONAR vs Grids

- Want to measure energy efficiency of a given job running on a grid, as well the grid as a whole.
- This is difficult, because:
 - Requires provenance on many characteristics: How to identify an application between job runs? On which node was it running? When did the job start and finish? etc...
 - How to systematically gather performance statistics on a job? Which metrics?
 - Heterogeneous hardware adds more parameters to be tracked.

GreenSONAR vs Grids

- All in all lots of meta-data needs to be tracked and stored, just to know what you are monitoring.
- Practical problems: Grid engines may not maintain all the required information.

GreenSONAR vs Virtualization

- Want to monitor VMs and clouds in the same way as grids.
- Currently not possible, because:
 - VMs are black boxes, the cloud system has no knowledge of what goes on inside them, so you cannot take meaningful measurements.

GreenSONAR vs Virtualization

- Fundamentally there is a paradox in monitoring a system that is designed to be opaque.
- May require a new kind of cloud architecture or platform in which the system knows what is happening **inside** a VM.

On metrics

- There are many ways to measure energy efficiency, including but not limited to: joule/bit, GFLOPs/Watt, PUE, DCiE, TEEER, and so on.
- Essentially they all represent the same thing: the amount of energy dissipated per unit of work done, but they are not comparable.
- Is there a better metric?

Conclusions

- PerfSONAR architecture is a good starting point for GreenSONAR, but may not be sufficient.
- Need provenance in grid and virtualized (computing) environments.
- Not clear which metrics are most useful to monitor.

Links

- <http://ext.delaat.net/>
- <http://ext.delaat.net/smartgreen/index.html>
- <http://ext.delaat.net/news/2012-03-23/index.html>
- http://wiki.cs.vu.nl/greenclouds/index.php/Main_Page

- Thanks:
 - Karel van der Veldt
 - Hao Zhu
 - Todor Yakimov
 - Lutz Engels

Intermezzo

Food for Thought

Or

Thoughts for after Food.

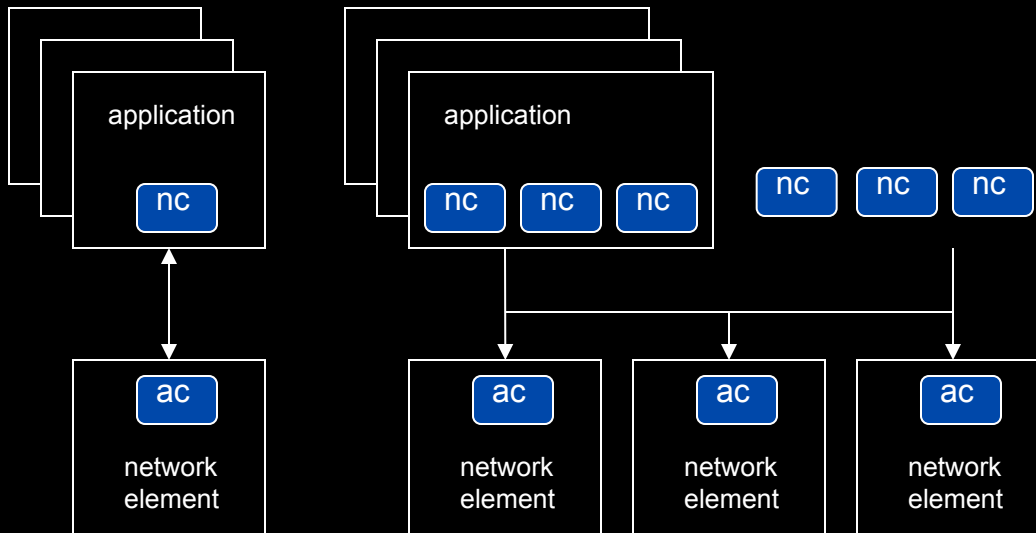
TeraThinking

- What constitutes a Tb/s network?
- CALIT2 has 8000 Gigabit drops ?->? Terabit Lan?
- look at 80 core Intel processor
 - cut it in two, left and right communicate 8 TB/s
- think back to teraflop computing!
 - MPI turns a room full of pc's in a teraflop machine
- massive parallel channels in hosts, NIC's
- TeraApps programming model supported by
 - TFlops -> MPI / Globus
 - TBytes -> OGSA/DAIS
 - TPixels -> SAGE
 - TSensors -> LOFAR, LHC, LOOKING, CineGrid, ...
 - Tbit/s -> ?

User Programmable Virtualized Networks allows the results of decades of computer science to handle the complexities of application specific networking.

- The network is virtualized as a collection of resources
- UPVNs enable network resources to be programmed as part of the application
- Mathematica, a powerful mathematical software system, can interact with real networks using UPVNs

$\text{Eigenvalues}\left[\begin{bmatrix} -1 & 0 & 2 \\ 2 & 9 & 2 \\ 3 & 1 & 4 \end{bmatrix}\right]$ <p>{9.484782381, 4.488378326, -1.973160708}</p>	$\sum_{\beta=1}^{30} \frac{1}{\beta^2}$ <p>1.612150118</p>
$\text{Plot}[\text{Sin}[13 x] + \text{Sin}[18 x], \{x, 0, 2\}]$	$\text{BesselJ}[1, 3 + i]$ <p>0.4326156394 - 0.4295057869 i</p>
	$\text{Simplify}[1 + 5 x + 10 x^2 + 10 x^3 + 5 x^4 + x^5]$ <p>(1 + x)⁵</p>
	$\text{mydata} = \{\{0.444539, 0.908491\}, \{1.4486, 1.84577\}, \{1.8734, 1.84577\}, \dots\}$ <p>$\text{Fit}[\text{mydata}, \{1, x, x^2\}, x]$</p> <p>0.2617148495 + 1.007 x - 0.0034235343 x²</p>



Mathematica enables advanced graph queries, visualizations and real-time network manipulations on UPVNs

Topology matters can be dealt with algorithmically

Results can be persisted using a transaction service built in UPVN

Initialization and BFS discovery of NEs

```
Needs["WebServices`"]
<<DiscreteMath`Combinatorica`
<<DiscreteMath`GraphPlot`
InitNetworkTopologyService["edge.ict.tno.nl"]
```

Available methods:

```
{DiscoverNetworkElements, GetLinkBandwidth, GetAllLinks, Remote,
NetworkTokenTransaction}
```

Global`upvnverbose = True;

```
AbsoluteTiming[nes = BFSDiscover["139.63.145.94"];][[1]]
```

```
AbsoluteTiming[result = BFSDiscoverLinks["139.63.145.94", nes];][[1]]
```

Getting neighbours of: 139.63.145.94
 Internal links: {192.168.0.1, 139.63.145.94}
 (...)
 Getting neighbours of: 192.168.2.3



Transaction on shortest path with tokens

Internal links: {192.168.2.3}

```
nodePath = ConvertIndicesToNodes[
ShortestPath[
g,
Node2Index[nids, "192.168.3.4"],
Node2Index[nids, "139.63.77.49"],
nids];
Print["Path: ", nodePath];
If[NetworkTokenTransaction[nodePath, "green"]==True,
Print["Committed"], Print["Transaction failed"]];
```

Path:
 {192.168.3.4, 192.168.3.1, 139.63.77.30, 139.63.77.49}

Committed



TouchTable Demonstration @ SC08

