Sustainable Operationalization of Smart City Network Fabric

Kim Khoa NGUYEN
Ecole de technologie superieure
kim-khoa.nguyen@etsmtl.ca

ÉTS
Le génie pour l'industrie

Outline

- Smart city projects in Canada & Montreal
- Sustainable smart city network fabric
- Energy-aware data plane
- Energy-aware control plane
- Network data analytics and Energy efficiency
- Conclusion

Why Smart Cities?

- The Challenge: Urbanization
  - By 2050
    - Over 70% of world population will live in cities
    - Occupy 2% of landmass
    - Consume 75% of resources
- The Opportunity
  - To enable ***livable and sustainable*** cities and urban regions
    - *economic, environmental, social*
- The Focus: Technology to address this opportunity
  - Platforms to enable Smart City applications
  - …and be sustainable!

Canadian Smart City Labs

Toronto Waterfront: waterfrontoronto.ca

- **Partners**
  - **Google (Sidewalk Lab – 12acre), city of Toronto, Federal government**
  - **Focus on: smart transport, smart living, smart building, digital government**

Le génie pour l'industrie

Canadian Smart City Labs

Montreal Lab-VI: labvi.ca

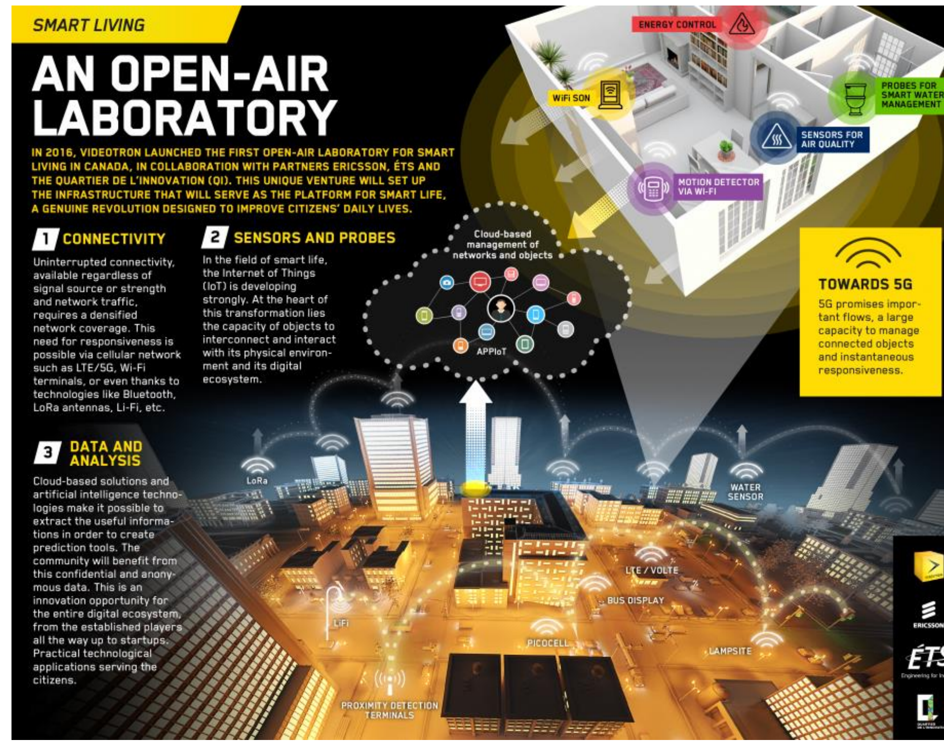

❑ Partners

- ➢ **Ericsson, Videotron, ETS Montreal, Quartier de l'Innovation de Montreal (2 km2)**
- ➢ **Focus on: 5G smart city applications, ecosystem of startups**

ÉTS
Le génie pour l'industrie

Sample applications

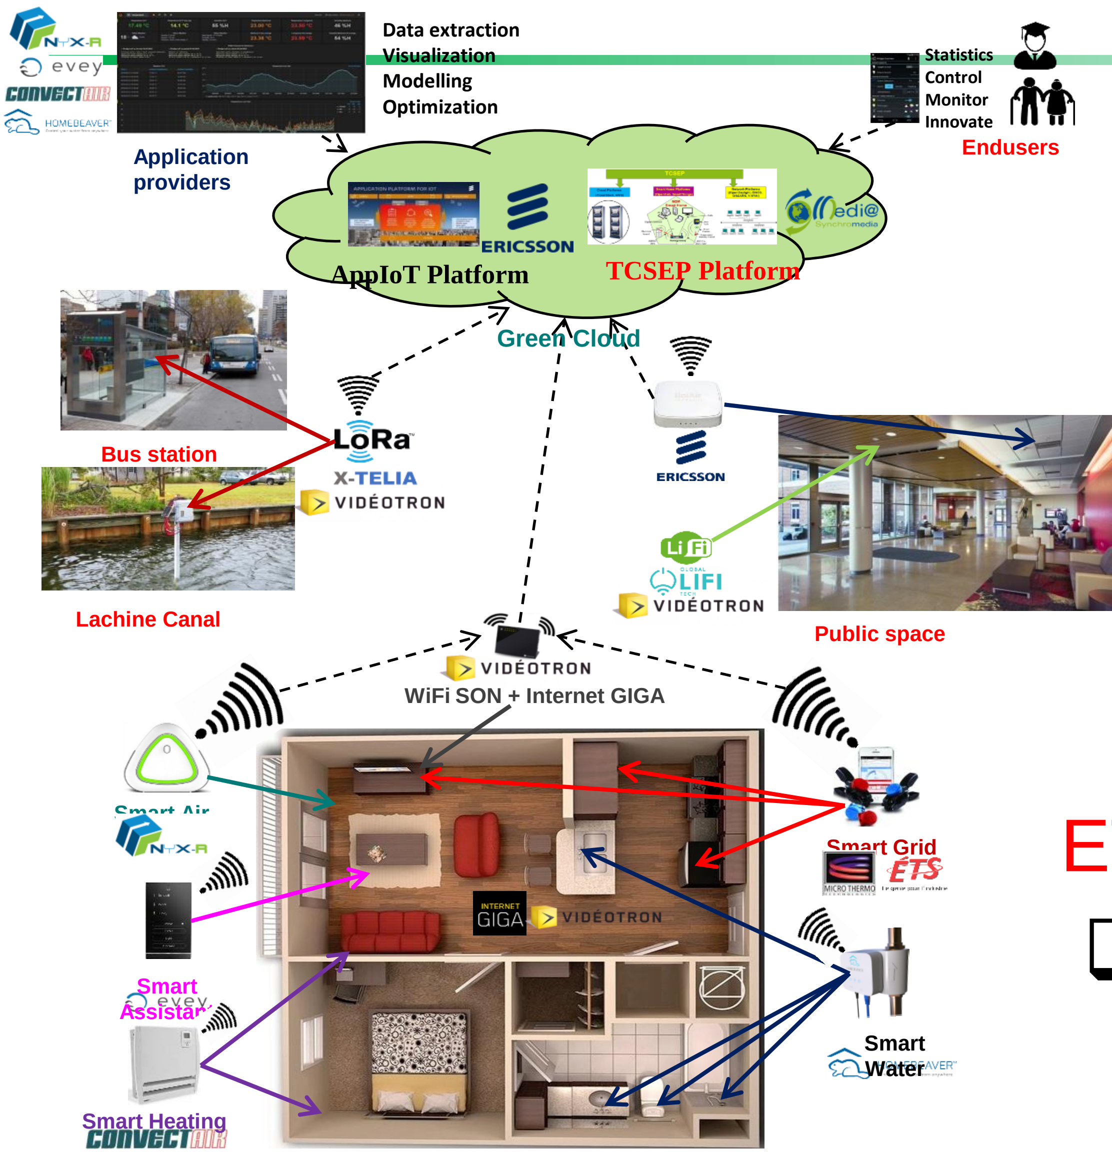

ETS Smart Residence

- ❑Gigabit university campus

Sample applications

Smart Bus Shelter

- ❑ Videotron + ETS Montreal
- ❑ Optimization of public transit resources
- ❑ Targeted advertising

Sample applications

Self-driving Shuttle

- ❑ Ollibus + ETS Montreal
- ❑ Optimize autonomous, electric vehicles
- ❑ Develop traffic accident reduction technologies
- ❑ Improve user experience, including interactive service programs

Sample applications

Autonomous Inspection

- ❑ Definitechs + ETS Montreal
- ❑ Detect and isolate the anomalies and faults
- ❑ AI-based pattern recognition

Sample applications

Smart BeeHive

- Nectar + ETS Montreal + INRS + McGill
- Bee-friendly monitoring
- Control swarms, manage pests, plan harvest

Sample applications

Rooftop Agriculture

- AU Lab + ETS Montreal
- Smart farm on top of Montreal Palais de Congrès
- Save building energy consumption

ÉTS Le génie pour l'industrie

Smart is all about Data!

- Real-time Situational Awareness
  - Continuous Monitoring
  - Data Collection & Historical Record
- Learning and Intelligence
  - Analytics & Machine Learning as-a-Service
- Visualization
    - System Dashboard Available on Demand—Anytime, Anywhere, Authorized
    - Trending & Forecasting
    - KPI Analysis
- Smart Applications Enablement
  - APIs provide real-time, historical, value-add data for applications:
    - Traffic Control, Planning, Optimization
    - Emergency Response
    - Event Planning & Management

ÉTS
Le génie pour l'industrie

Goal

- Develop a *Smart City Network Fabric*
  - integrate Internet of Things (IoT), computing, and networking technologies
  - empower individuals, organizations, and society to realize the benefits latent
  - in systems and data at scale in future smart cities.
- IoT Gateways
- Software-Defined exchanges
- Applications with smart city partners
  - Transportation, air quality, carbon footprint, smart buildings, etc.

Smart city network fabric

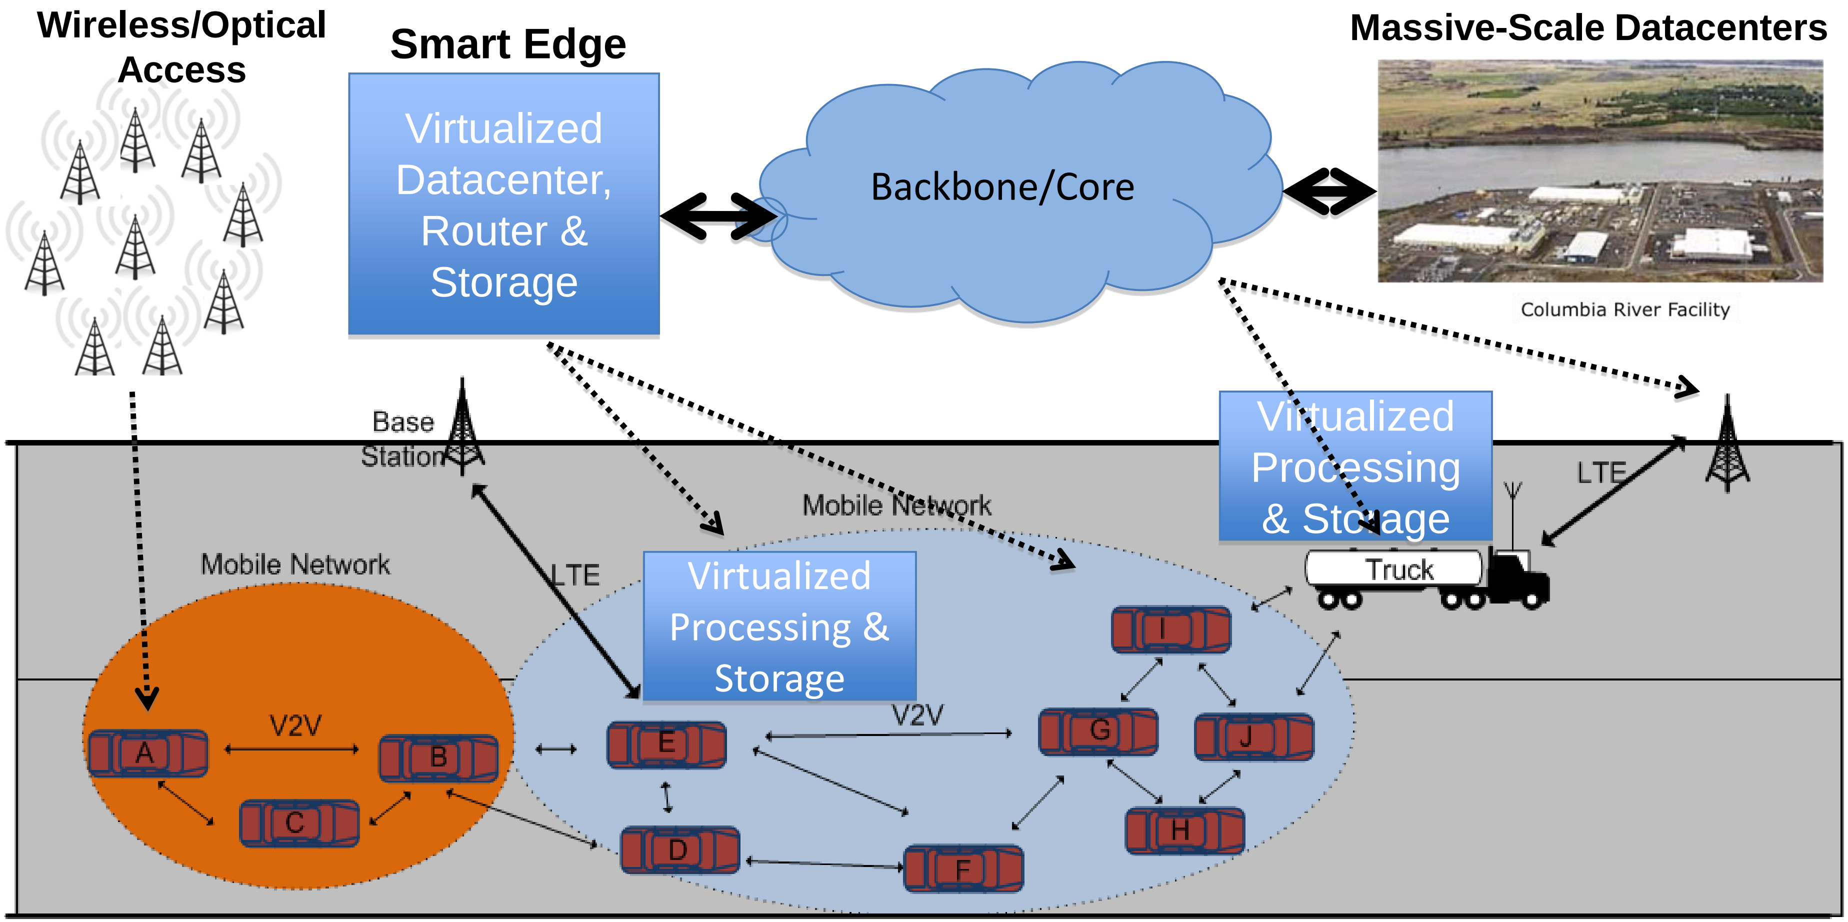

Issues in the network fabric

- Currently focus on bandwidth provisioning
- Issues:
  - Use a large variety of hardware appliances, dedicated to specific tasks, which typically are
    - *Inflexible,*
    - *Energy inefficient*

ÉTS
Le génie pour l'industrie

Keywords

- The main tasks of the network is allocating resources
  - How to make it more dynamic, performance-optimized and cost-effective?

- Current keywords are
  - ***Flexibility***
  - ***Programmability***
  - ***Energy-efficiency***

Energy inefficiencies in networks

- ❑ Current network infrastructures, technologies and protocols are designed to be extremely over-dimensioned and available 24/7
- ❑ Links and devices are provisioned for rush hour load.
- ❑ The overall power consumption in today's networks remains more or less constant even in the presence of fluctuating traffic loads

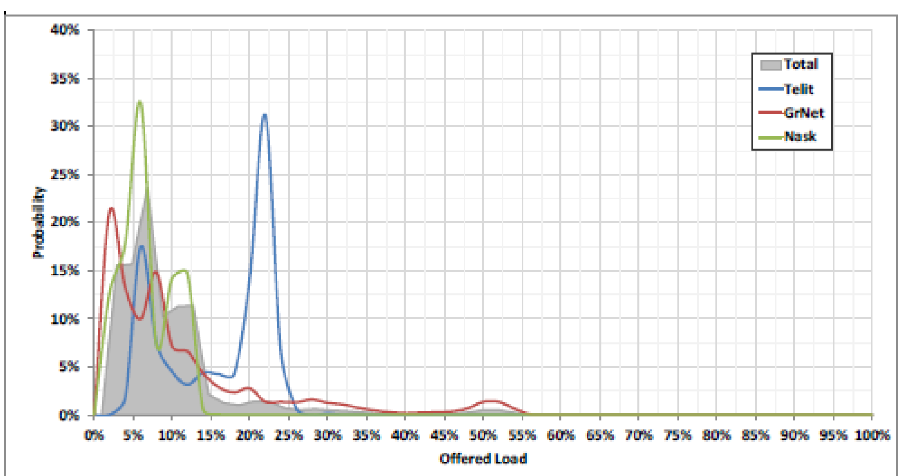

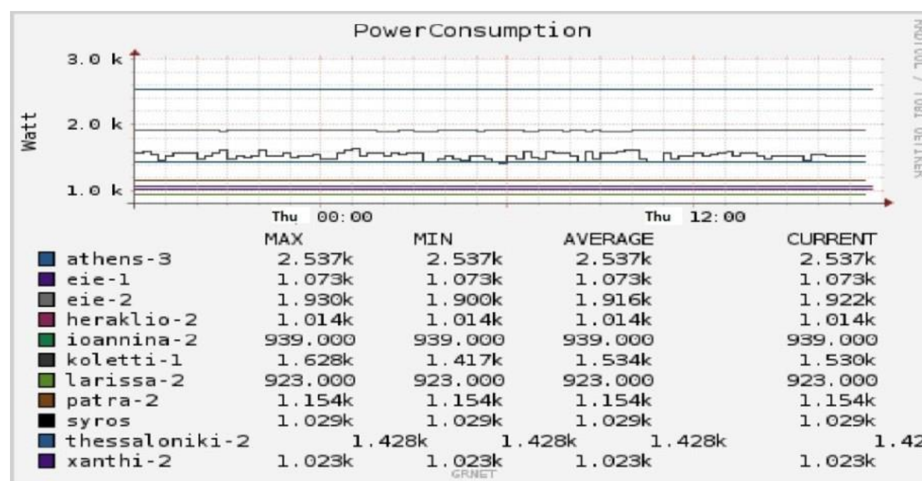

Energy inefficiencies in networks

- In reality:
  - Percentage w.r.t. peak level
  - The profiles exhibit regular, daily cyclical traffic patterns with Internet traffic dropping at night and growing during the day

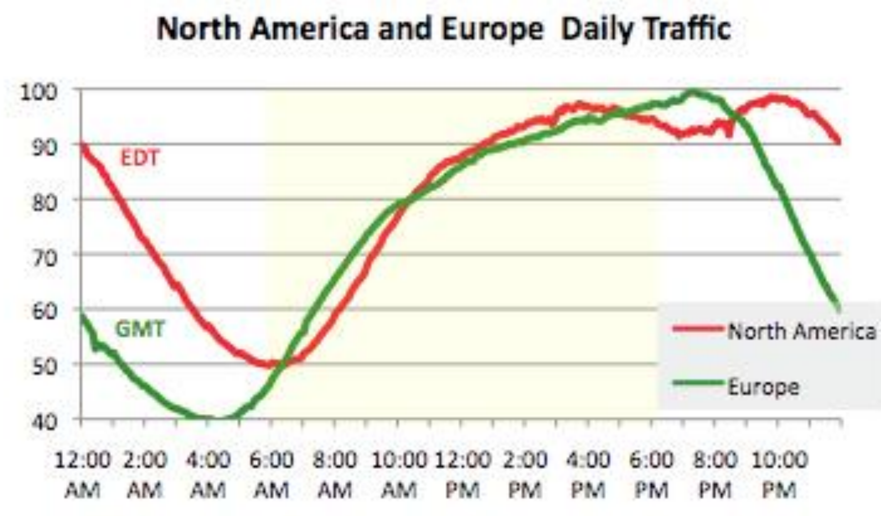

Traffic load fluctuation at peering links for about 40 ISPs from USA and Europe

Energy inefficiencies in networks

- Traffic within data centers
  - Time dependent (day vs. night time)
  - Lot of un-utilized links
    - Back-up/redundant links
  - Aggregation layer has the most used links
  - Very few links with losses

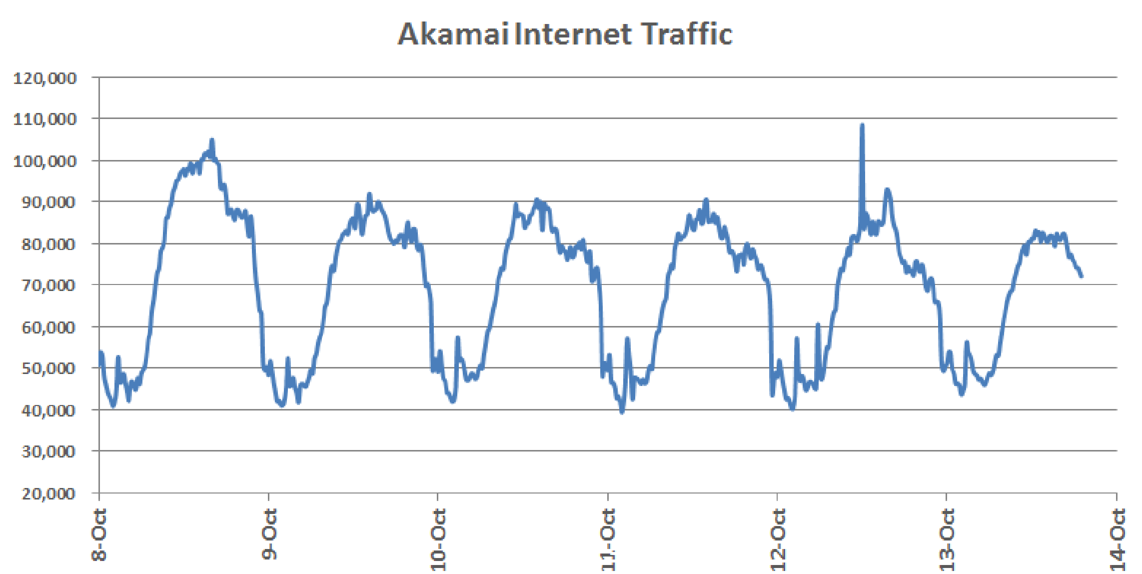

Breakdown the energy consumption in network

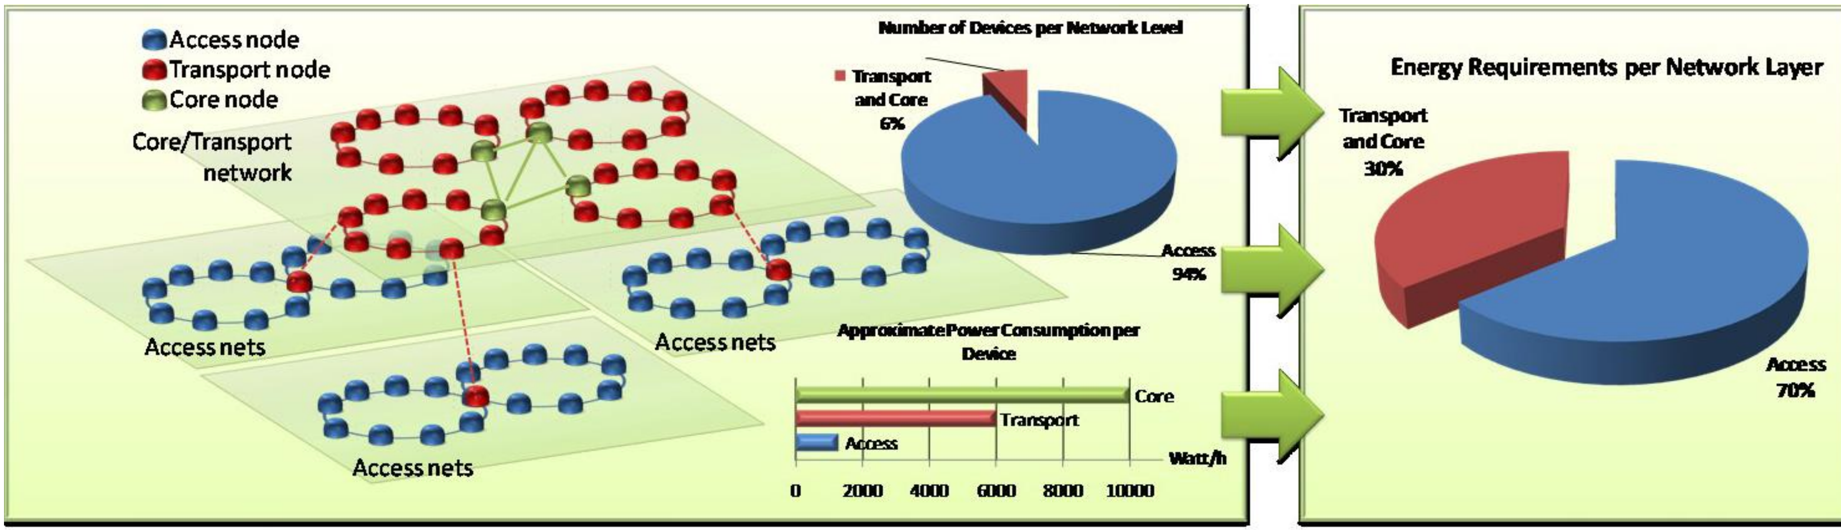

❑ Typical access, metro and core device density and energy requirements in today's typical networks deployed by Telcos, and ensuing overall energy requirements of access and metro/core networks

How to manage this trend

- Today's (and future) network infrastructures characterized by:
  - **Design capable to deal with strong requests and constraints** in terms of resources and performance (large loads, very low delay, high availability, ...)
  - **Services characterized by high variability of load and resource requests** along time (burstiness, rush hours, …)

- The current feasible solution:
  - **Smart power management**: energy consumption should follow the dynamics of the service requests.
  - **Flexibility in resource usage**: virtualization to obtain an aggressive sharing of physical resources

ÉTS Le génie pour l'industrie

Network energy-efficiency and QoS

- Making the network energy-efficient ("Green") cannot ignore Quality of Service (QoS) / Quality of Experience (QoE) requirements

- Much higher flexibility are required to effectively deal with the performance/power consumption tradeoff
  - A new dimension of energy-awareness is taken into account in all phases of network design and operation

Network energy-aware approaches

- Data plane (micro level)
  - Dynamic adaptation
  - Smart stand-by
- Control plane (macro level)
  - Consolidation
  - Traffic engineering and routing

Energy-aware data plane

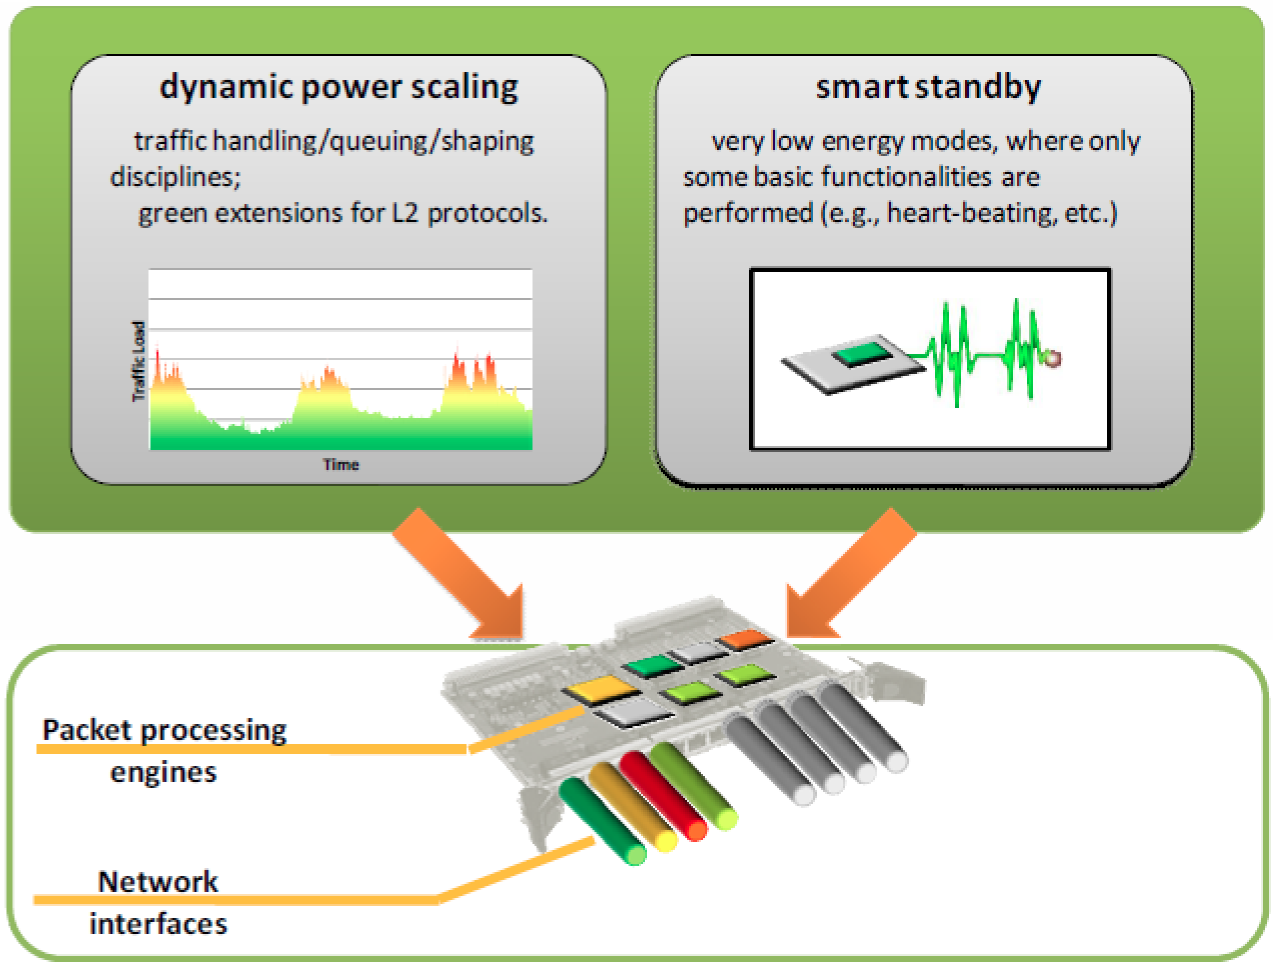

Dynamic Adaptation: QoS vs Power Management

- The maximal power saving is obtained when equipment is actually turned off
  - However, under such condition the performance is actually zero
- The best performance equipment may provide is under no-power-limit mode
- We look into intermediate possibilities between these two extremes

Energy-aware network element

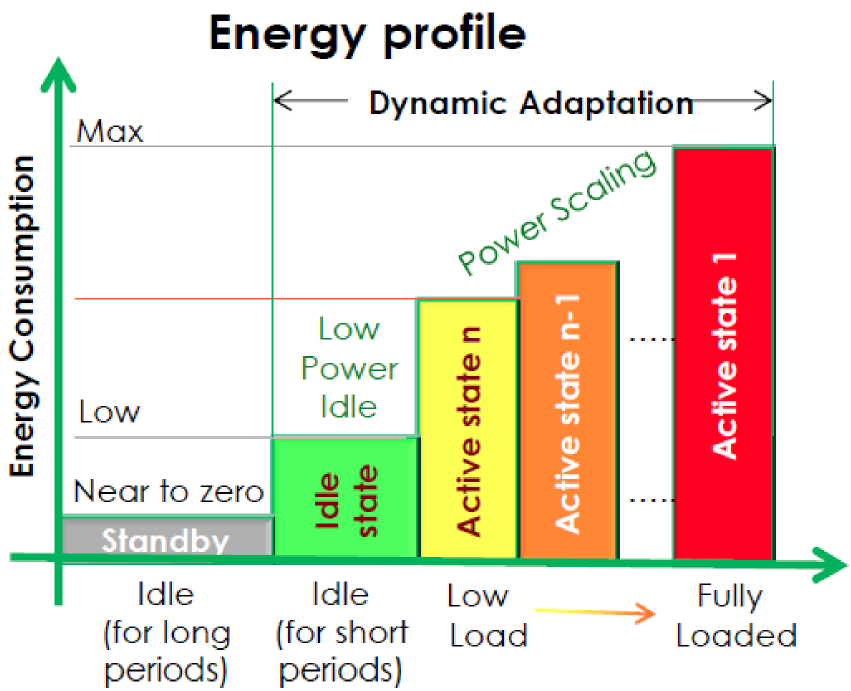

QoS vs Power Management

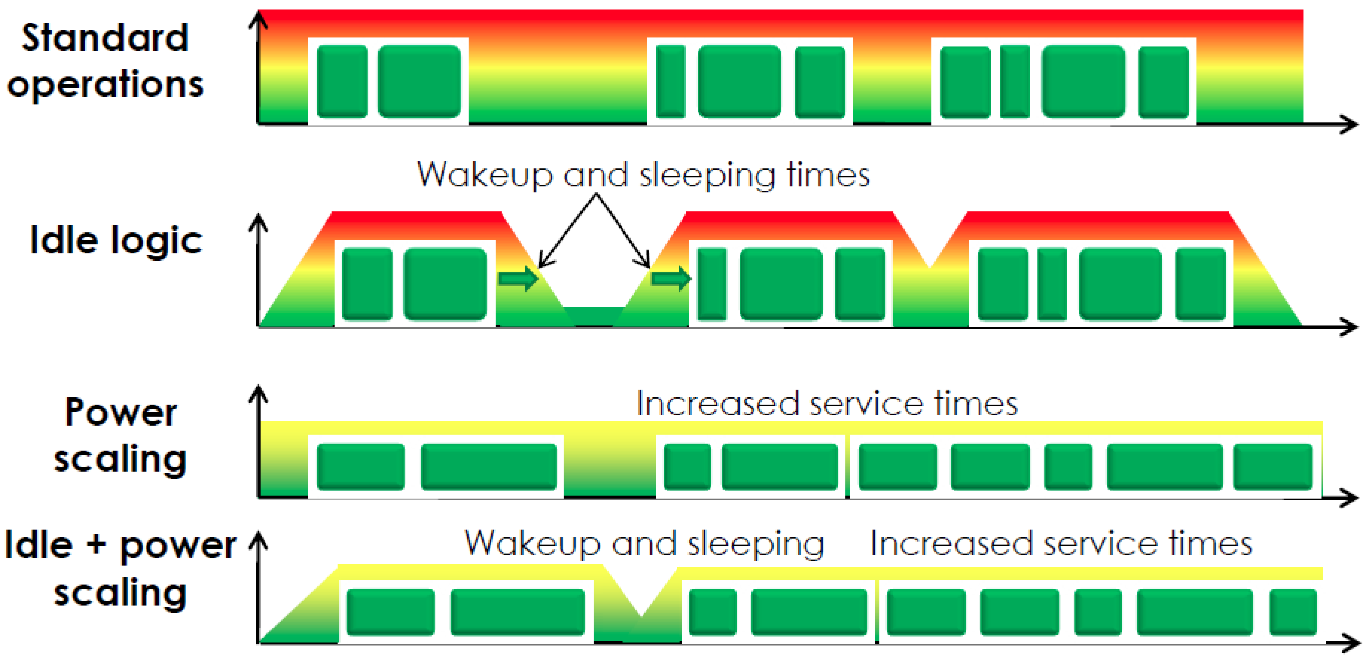

Dynamic Adaptation

Link-level example: Green Ethernet (IEEE 802.3 az)

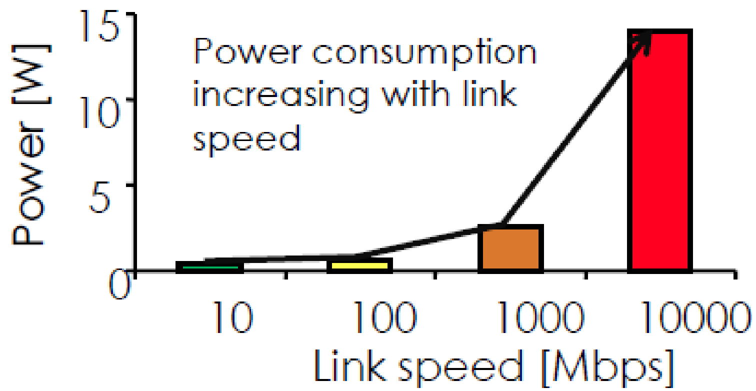

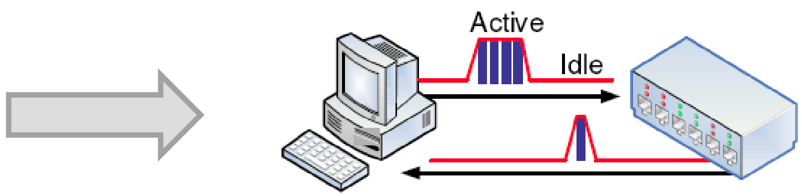

- ❑ Based on the “low power idle” concept
- ❑ **Idea:** transmit data at the maximum speed, and put the link to sleep when it is idle
- ❑ **Effect:** LPI has two transitions for each packet (or block of packets) : Link wake-up and sleep
  - ➢ LPI can possibly be asynchronous (one direction awake, the other asleep)
  - ➢ Retraining can be done via periodic on intervals (if no packets are being sent)
  - ➢ LPI requires no complicated handshaking

Dynamic Adaptation

Dynamic Voltage & Frequency Scaling

- DVFS scales the operating voltage of the processor along with the frequency
- Processors are based on CMOS logic
  - Static power + Dynamic power
- Dynamic power: $P_{dyn} = CfV^2$
  - C: constant, f : frequency, V: voltage
- Total power consumption: $P = P_{dyn} + P_{static}$

Dynamic Adaptation
Simple DVS-Scheme

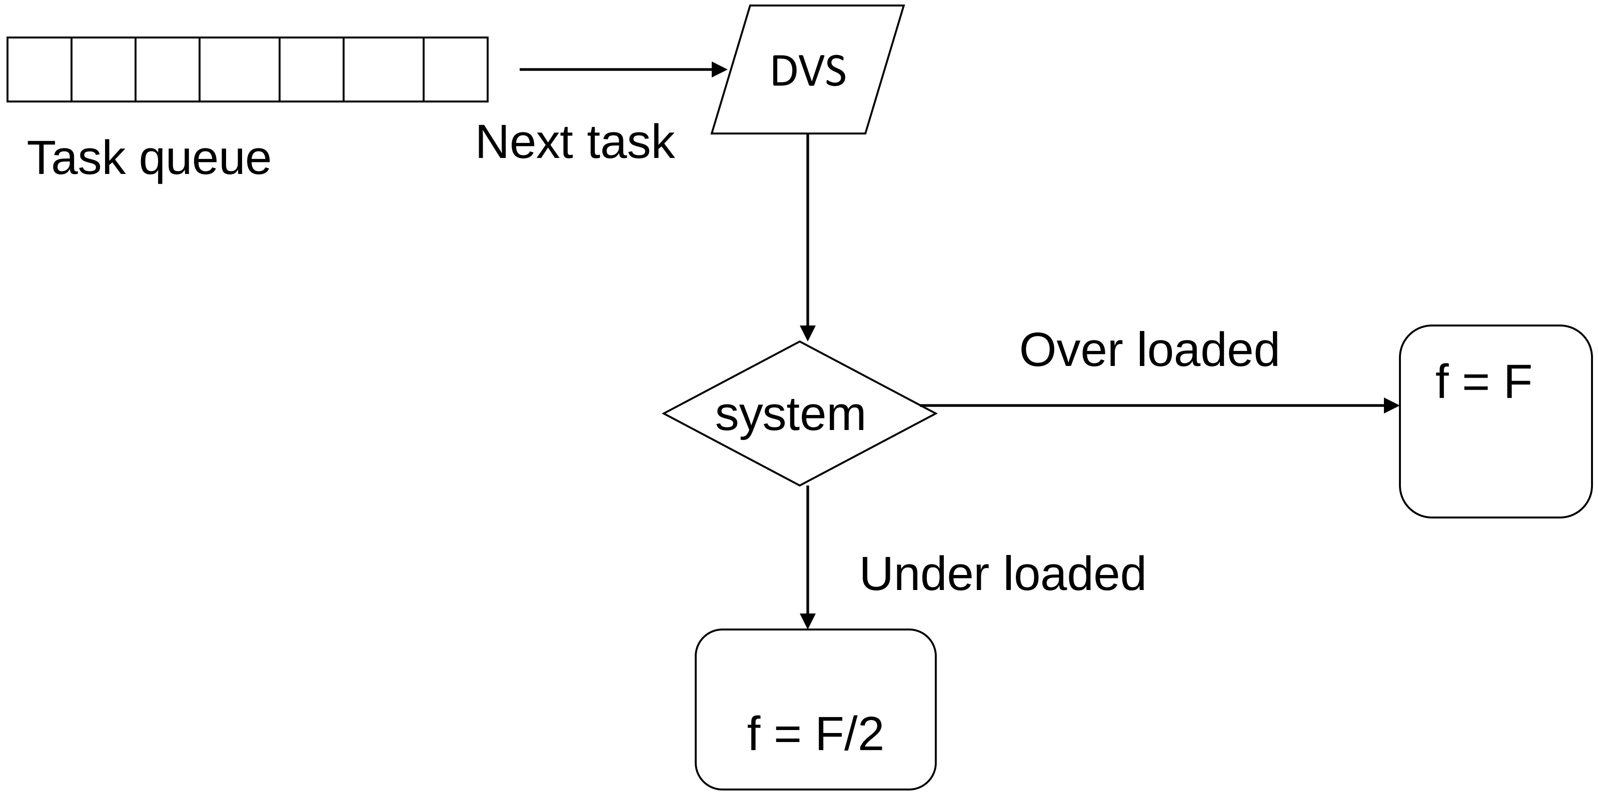

Dynamic Adaptation

Advanced Configuration and Power Interface (ACPI)

- ACPI provides a standardized interface between the hardware and the software layers.
- ACPI introduces two power saving mechanisms, which can be individually employed and tuned for each core:
  - Power States (C-states)
    - C_0 is the active power state
    - C_1 through C_n are processor sleeping or idle states (where the processor consumes less power and dissipates less heat).
  - Performance States (P-states)
- while in the C0 state, ACPI allows the performance of the core to be tuned through P-state transitions
- P-states allow modify the operating energy point of a processor/core by altering the working frequency and/or voltage, or throttling the clock

Dynamic Adaptation

- ❑ If we change the consumption we change also the performance

=>Need to model a device or system in terms of consumption and performance versus loads and configurations

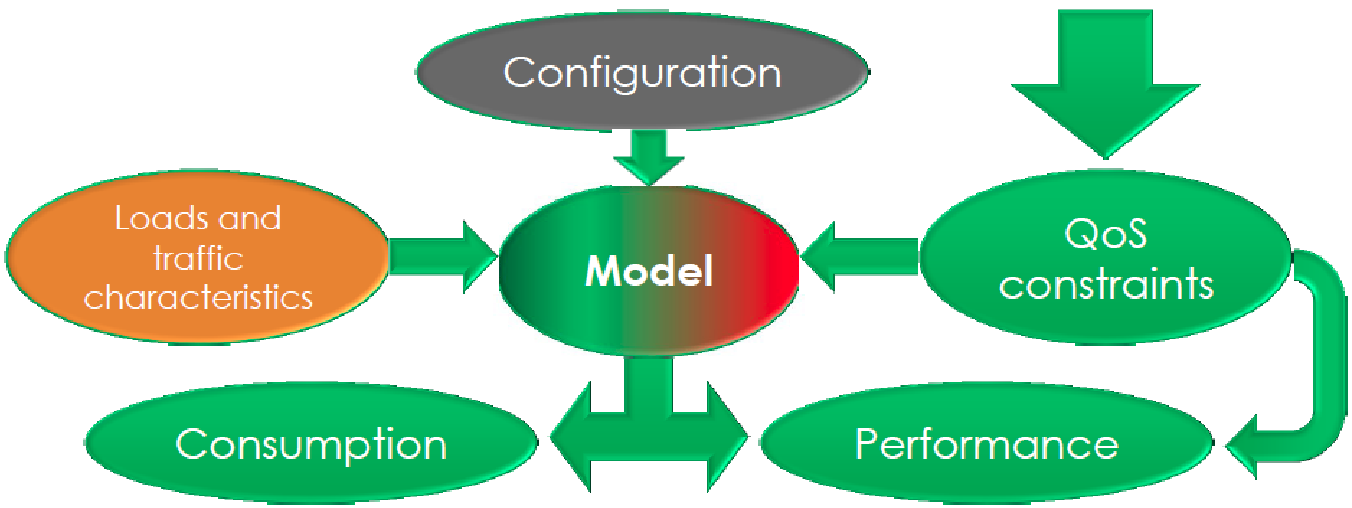

Sleeping/standby

- Smartly and selectively drive unused network/device portions to low standby modes, and to wake them up only if necessary.

- However, since networks and applications are designed to be continuously and always available
  - Standby modes have to maintain the “network presence” of sleeping elements:
  - **Solution: Proxying the network presence**

- Additional techniques should be added to enlarge as much as possible the number of “sleeping” parts or elements
  - **Solution: Network virtualization**
  - **Solution: Energy-aware traffic engineering and routing**

ÉTS
Le génie pour l'industrie

Sleeping/standby

Proxying the Network Presence

- Proxying at the application layer

Current situation
Everything active
(and consuming)

Router

Home
gateway

P2P application
active -> PC is
maintained active all
time long

Proxying state
Home gateway can manage the
interactions on behalf of the
application

Router

Smart Home
gateway
with proxy
capabilities

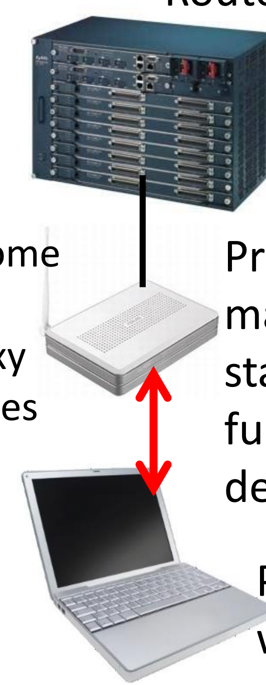

Protocols for
managing the
standby state and
functionality
delegations

PC enters standby
with very low
consumption

ÉTS
Le génie pour l'industrie

Sleeping/standby

Supporting End-System Standby

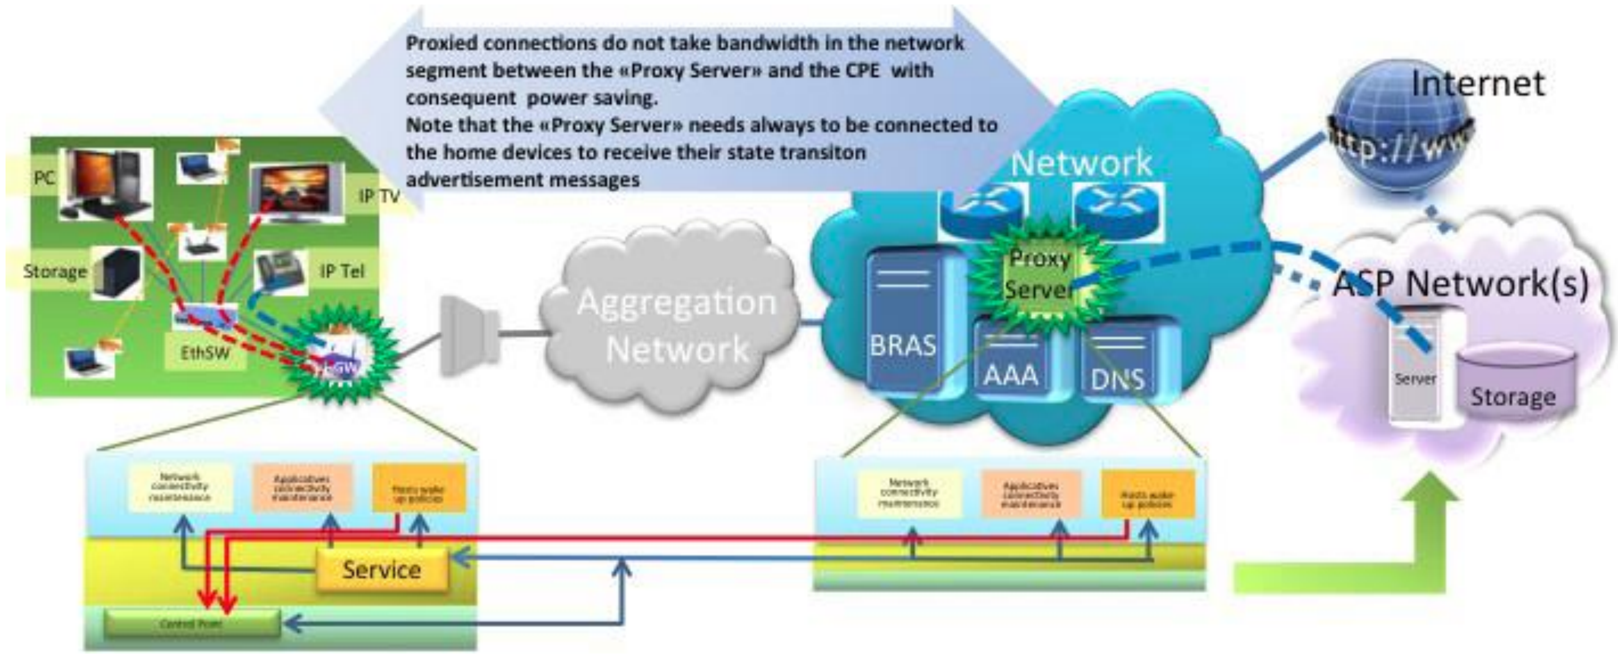

- Advertises itself as NCP to the devices in the home network
- Supports proxying of connections within the home network
- Exports towards the mate «external» NCP
  - requests for actions received from hosts related to connections over the «external» network
  - Hosts' power state
- Receives from mate «external» NCP requests for hosts' wake up

- Supports proxying of connections over the external network
- Collects from «internal» NCP
  - Requests for actions received from hosts related to connection over the «external» network
  - Hosts' power state
- Forwards to «internal» NCP requests for hosts' wake up

Sleeping/standby
Network-wide Control

- Give the possibility of moving traffic load among network nodes
- **When a network is under-utilized, we can move network load on few "active" nodes, and put all the other ones in standby**
- Different network nodes can have heterogeneous energy capabilities and profiles
- Network-wide control strategies could cut the overall energy consumption by more than 23%.

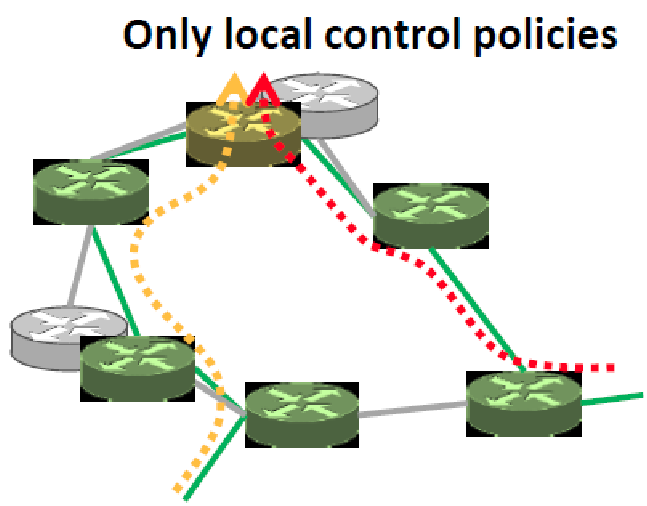

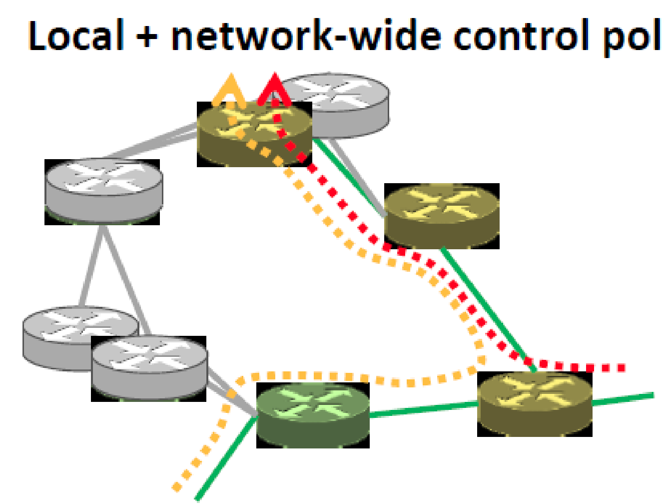

Le génie pour l'industrie

Green network-wide control: Traffic engineering and routing

- Current research in this area mainly focus on:
  - Putting links in standby modes
    - calculating the minimal sub-topology for meeting QoS constraints.
  - Considering the energy profile of devices or their sub-components
    - acting on routing/TE metrics in order to move flows towards "greener" alternative paths

Energy-aware control plane

Autonomic and short-term on-line optimizations

Operator-driven long-term off-line optimizations

Local Control Policies (LCPs)

- Given:
  - the actual traffic workload from input links
  - Local service requirements
- Dynamically find the best energy-aware configuration

Routing & Traffic Engineering

- Given:
  - The traffic matrix
  - Service requirements
  - The energy-aware capabilities of network nodes and links
- Dynamically move the traffic flows among network nodes in order to minimize the overall network consumption

Network-wide Monitoring

- Given the history of measurements regarding:
  - network performance
  - energy consumption
- The operator can explicitly plan and/or reconfigure the settings of:
  - single device
  - traffic engineering and routing

Network Control Policies (NCPs)

Traffic engineering and routing
Putting links in standby modes

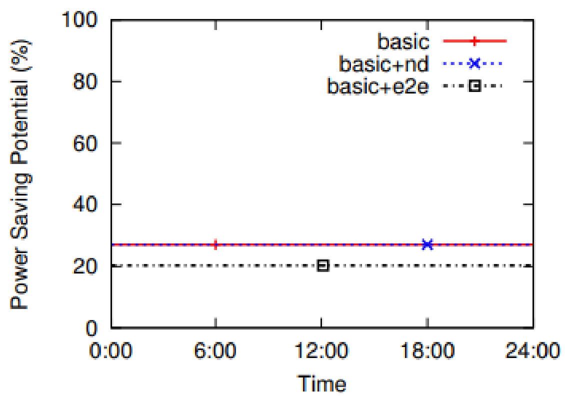

Power saving potential of Abilene

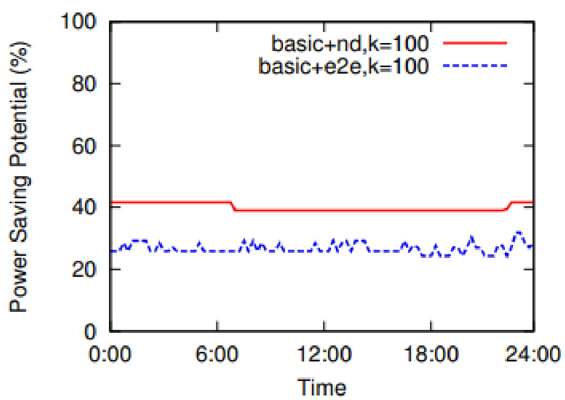

Power saving potential of GÉANT

- Given the network topology and traffic matrix
- Determine a routing solution (i.e., the links to be used and the traffic volume to be carried on each link) that
  - Maximizes the power saving from turning off line-cards
  - Satisfies performance constraints including link utilization and packet delay.
- Problem formulated based on the Multi-Commodity Flow (MCF) model
- Solving Mixed Integer Programming model

Traffic engineering and routing Considering the energy profiles

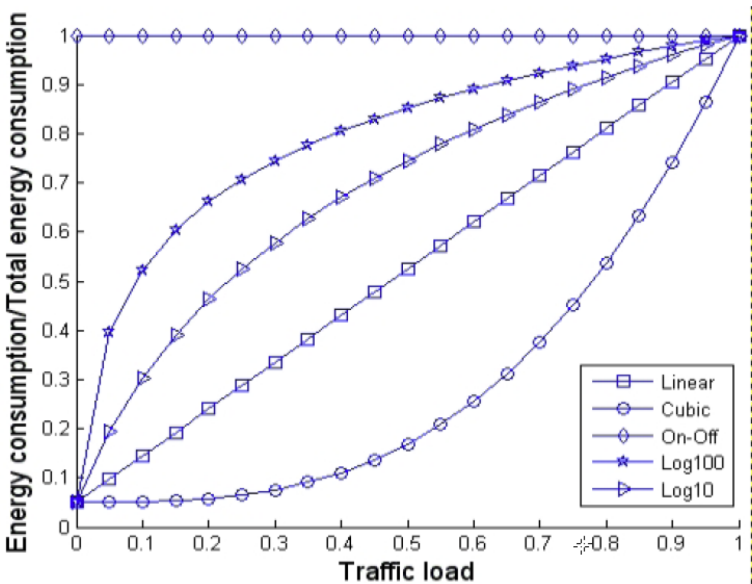

$$\text{Minimize:} \quad \sum_{n \in N} EP_{n,L(n)}$$

$$\sum_{p \in P_{s,d}} f_p = D_{s,d} \qquad \forall (s,d) \in D$$

$$UCE_e = \sum_{(s,d) \in D} \sum_{p \in P_{s,d}, e \in p} f_p \qquad \forall e \in E$$

$$UCE_e \leq C_e \qquad \forall e \in E$$

- The impacts of different router energy profiles on the energy-aware routing problem solution is shown

- $EP_{n,L(n)}$: energy-profiles (energy consumption of a component n at traffic load L)
- UCE_e: used capacity of an edge
- $P_{s,d}$: set of all possible paths between a source s and a destination d and f_p represents the flow on a path p.

Traffic engineering and routing Network data analytics

- The keyword is flexibility
- Routing is efficient only if we know network and traffic behaviours
  - Network modeling
  - Traffic forecasting

Traffic engineering and routing
Network monitoring

- What are we monitoring in the network?
  - System & Services
    - Available, reachable, energy consumption
  - Resources
    - Expansion planning, maintain availability
  - Performance
    - Round-trip-time, throughput
  - Changes and configurations
    - Documentation, revision control, logging

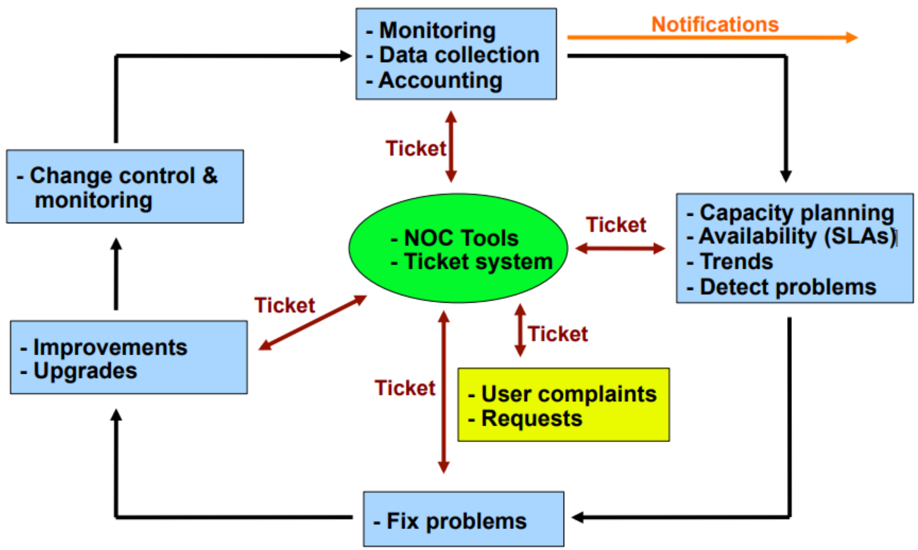

- Most used network monitoring software
  - Availability: Nagios
    - Services, servers, routers, switches
  - Reliability: Smokeping
    - Connection health, rtt, service response time, latency
  - Performance: Cacti
    - Total traffic, port usage, CPU, RAM, disk, processes

- NOC: Network Operations Centre
  - Coordinate tasks
  - Report status of network and services
  - Process network-related incidents and complaints
  - Host tools (ex. monitoring)
  - Generate documentation

Le génie pour l'industrie

Hierarchy of Data Analytics

- Data analytics can be looked at in multiple segments
  - ***Historical Analytics***: Build data warehouses / run batch queries to predict future events / generate trend reports
  - ***Near Real-Time Analytics***: Analyze indexed data to provide visibility into current environment / provide usage reports
  - ***Real-Time Analytics***: Analyze data as it is created to provide <u>instantaneous, actionable</u> business intelligence to affect immediate change
  - ***Predictive Analytics:*** Build statistical models that can classify/predict the near future

ÉTS
Le génie pour l'industrie

Hierarchy of Data Analytics

- Each segment of analytics serves specific purposes
  - ***Historical Analytics***: Campaign & service plan creation, network planning, subscriber profiling, customer care
  - ***Near Real-time Analytics***: Network optimization, new monetization use-cases, targeted services (ex. location-based)
  - ***Real-time Analytics***: Dynamic policy, self-optimizing networks, traffic shaping, topology change, live customer care
  - ***Predictive Analytics:*** traffic demand forecasting, fault avoidance, planed service provisioning

- Data is richer when associated to context – layer, location, time of day, etc.
- For each type of data, there is a window / meaningful time period of which the data is relevant

Sample network data analytics application

Traffic differentiation and QoS provisioning with traffic analyzer –

Elephants and mice flows

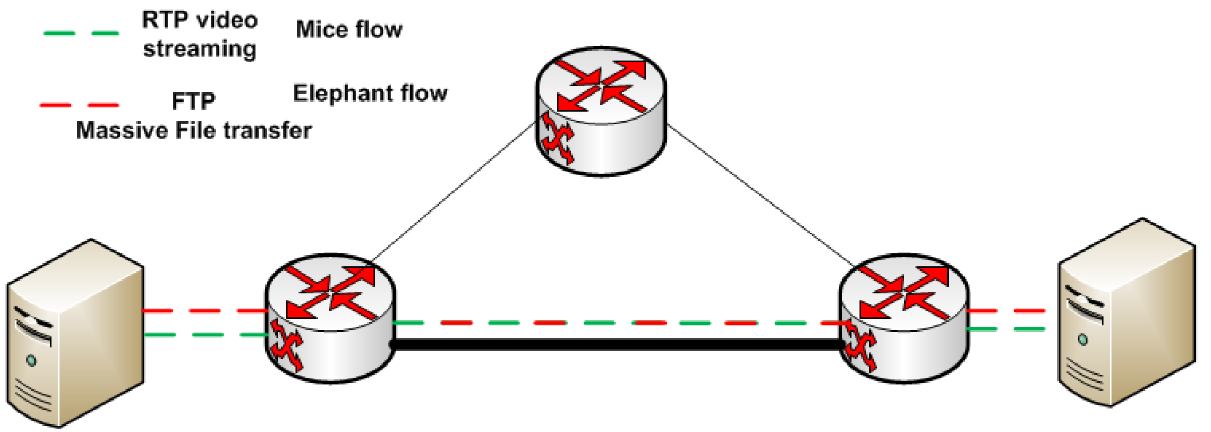

A) SPF routing

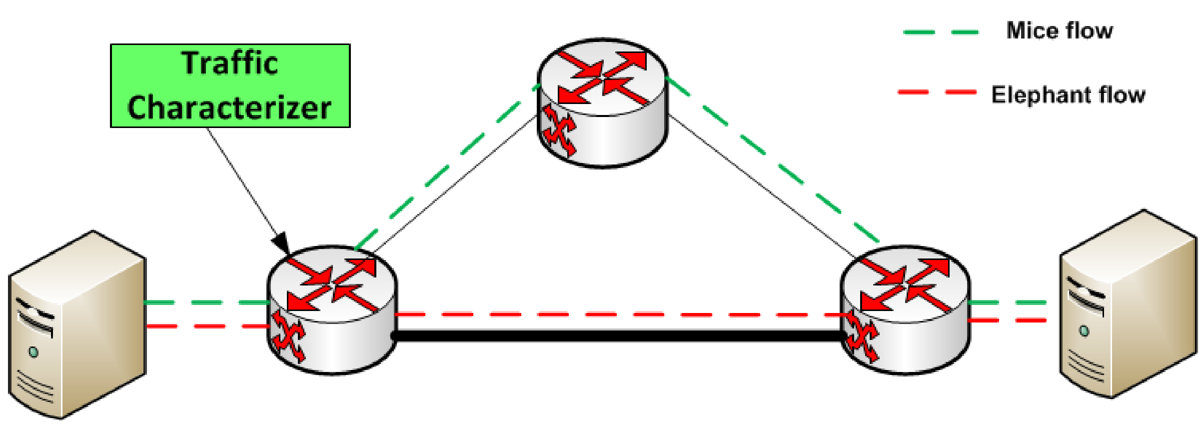

B) Traffic differentiation

Sample network data analytics application

Bandwidth defragmentation based on real-time monitoring and forecasting

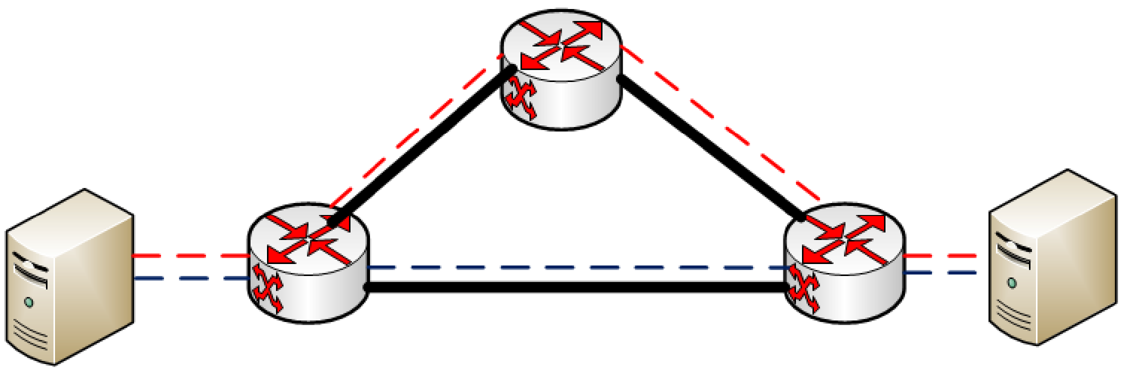

A) Random routing

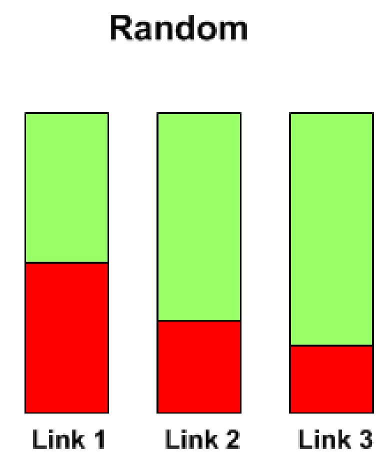

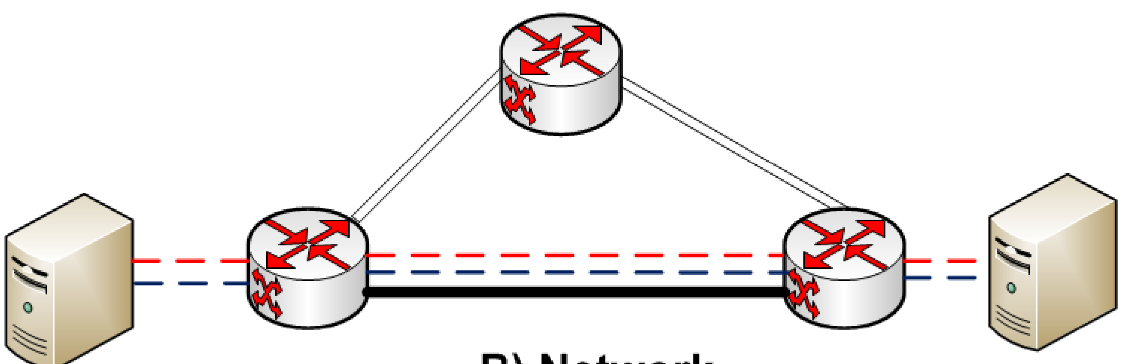

B) Network defragmentation

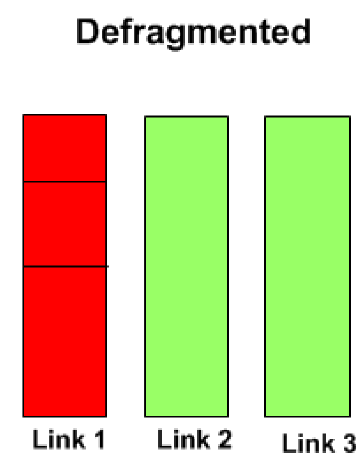

Sample network data analytics application

Prediction of faults

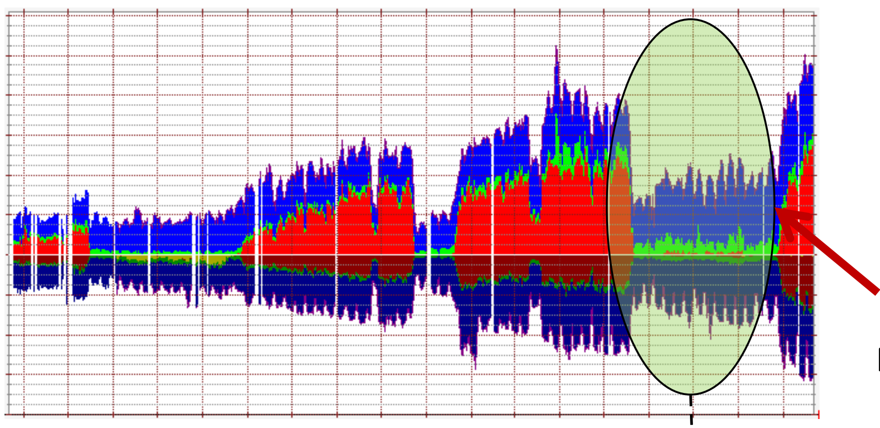

- Predicting and avoiding failures
- Load balancing
- Optimize resource allocation

Packet drops pattern detected

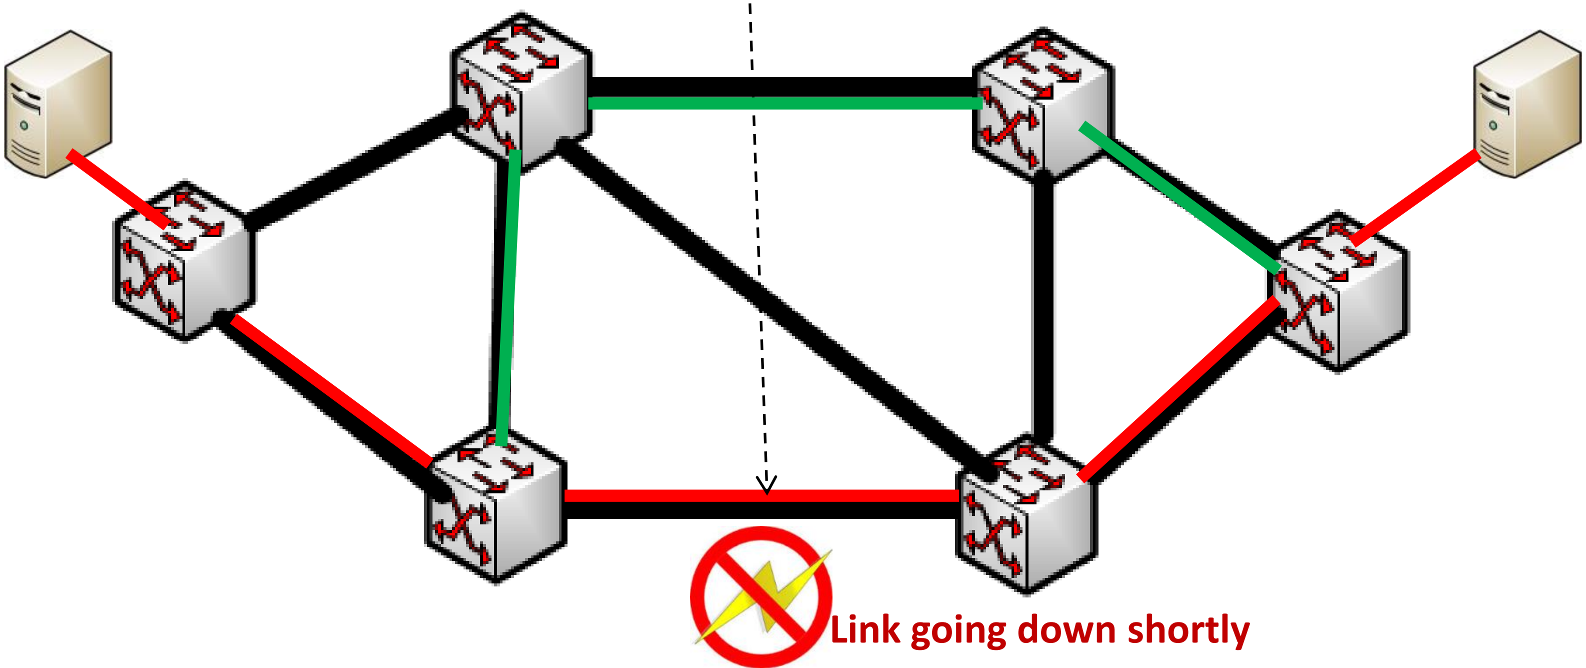

Link going down shortly

Sample network data analytics application

Network access visualization

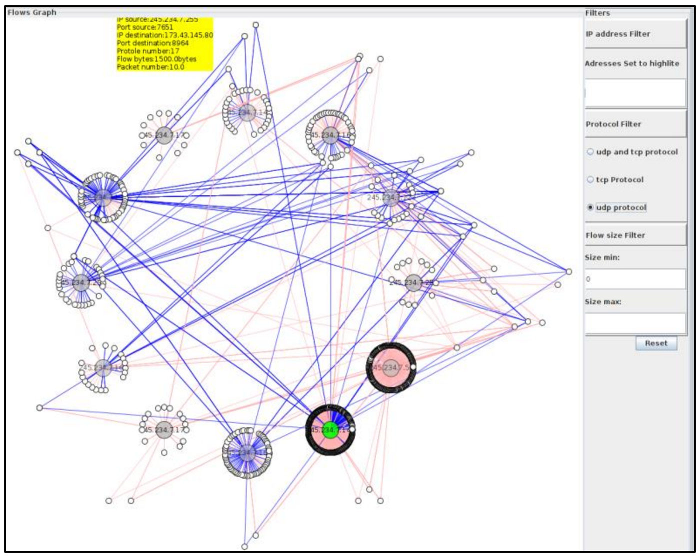

UDP traffic in network

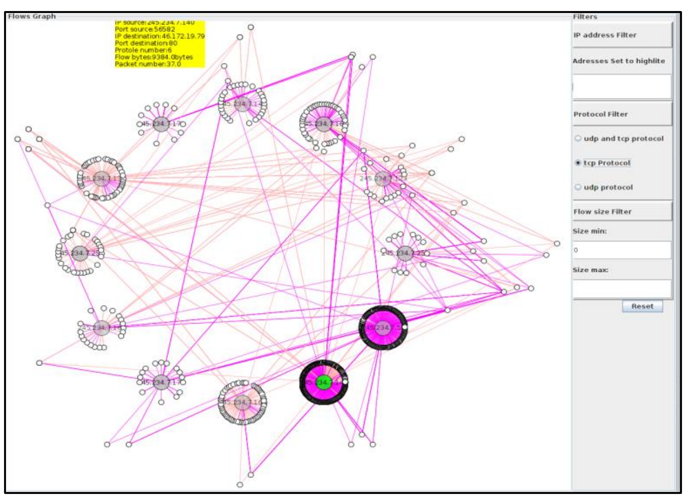

TCP traffic in network

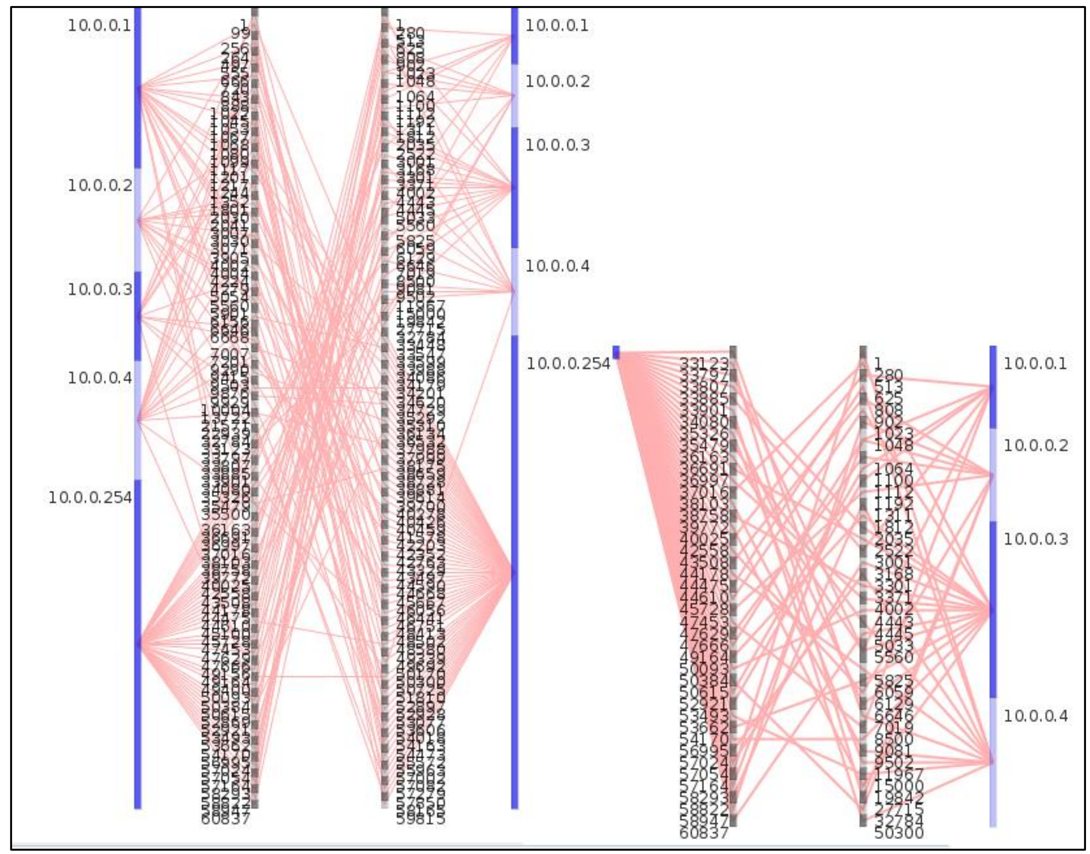

Port analyzing

Sample network data analytics application

Application traffic visualization

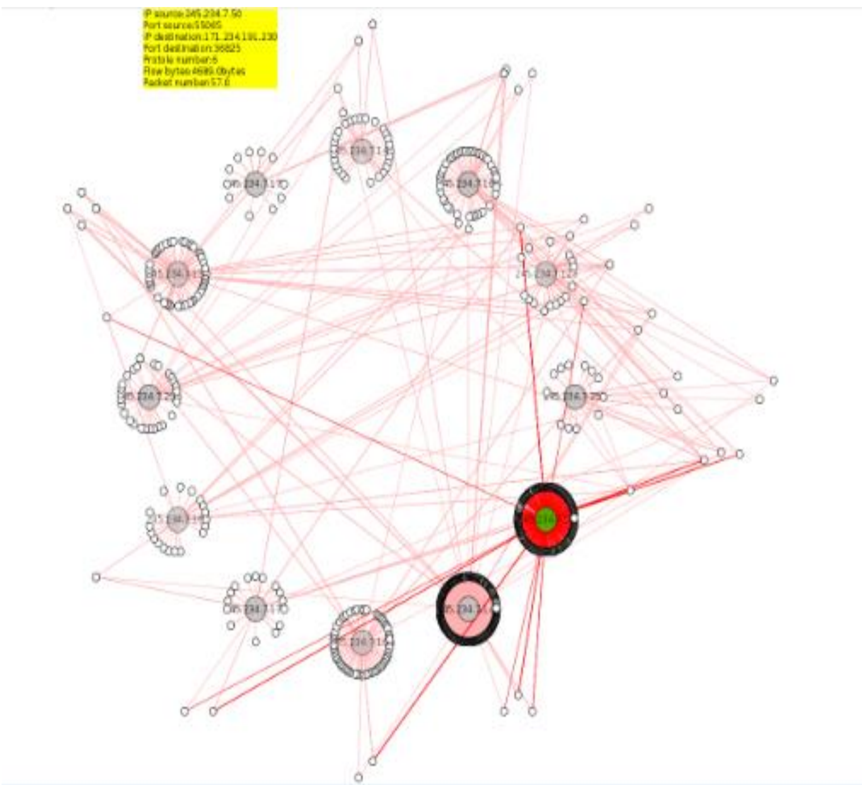

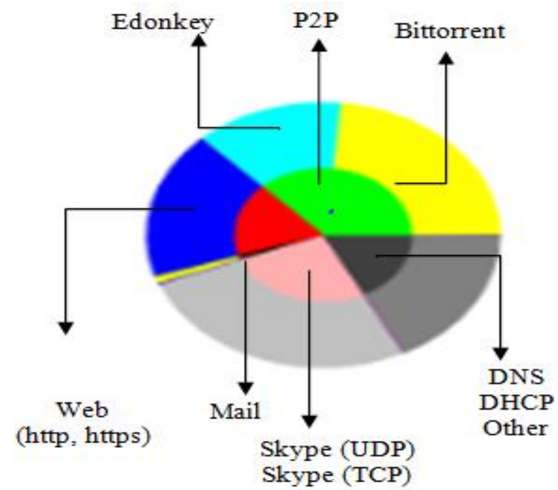

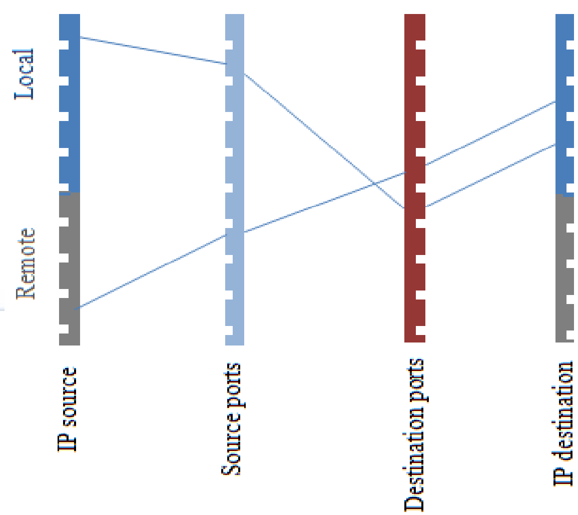

Network Traffic Analytics Framework

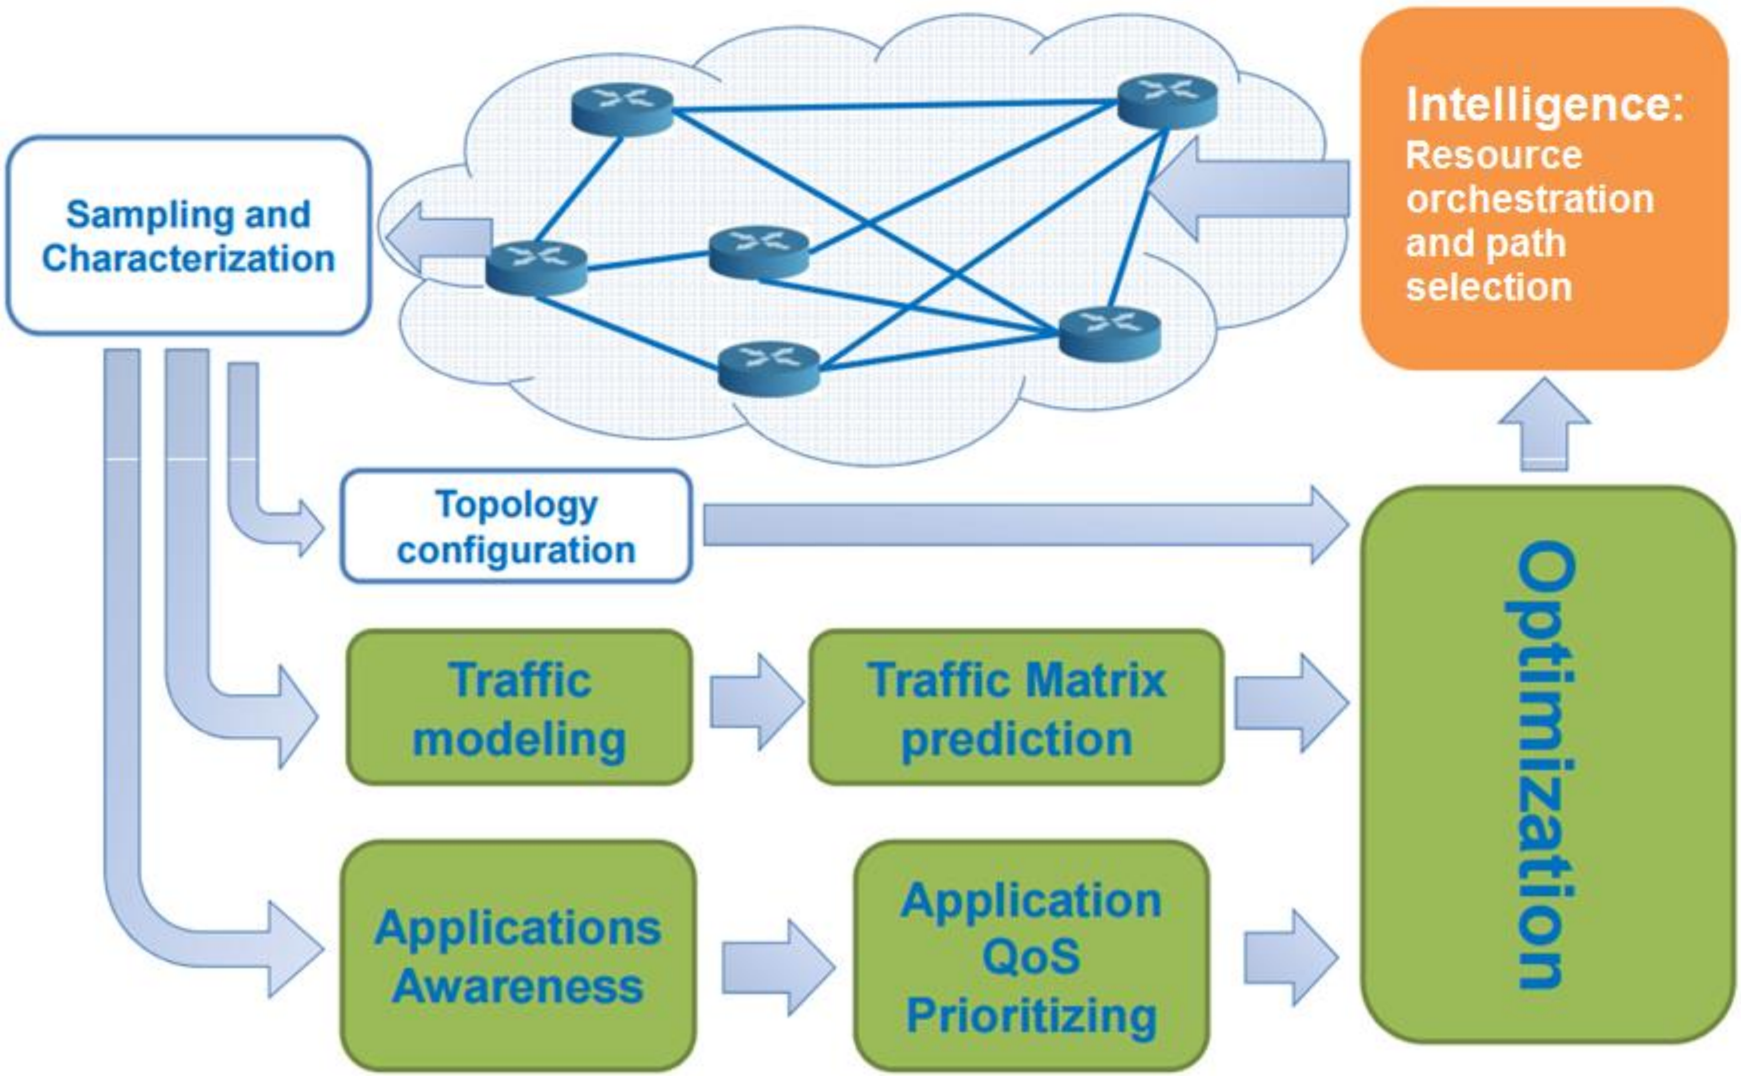

Network Traffic Analytics Framework

❑ Sampling

- ❖ Capturing all packets across the network is no longer appropriate
  - ▪ Overhead (resource consumption, computational time)
- ❖ Random, deterministic, or hash-based sampling
- ❖ Flow sampling vs. packet sampling
- ❖ Sample tools: NetFlow, sFlow

❑ Characterization

- ❖ Port-based characterization
  - ▪ Ex: HTTP, FTP ports
- ❖ Header-based characterization
  - ▪ Ex: IP packet header
- ❖ Payload-based analytics
  - ▪ Ex: Deep Packet Inspection (DPI)
- ❖ Application behaviour-based characterization
  - ▪ Ex. Video, voice, text messaging

Network Traffic Analytics Framework

❑ Modeling

- ❖ Stationary vs. non stationary
- ❖ Random model (e.g., Poisson) cannot capture traffic accurately
  - – There is self-similarity in traffic
- ❖ Two factors affecting traffic patterns
  - – Amount of *multiplexing* on the link: how many flows are sharing the link?
  - – Where flows are *bottlenecked:* Is each flow's bottleneck on, or off the link? Do all bottlenecks have similar rate?

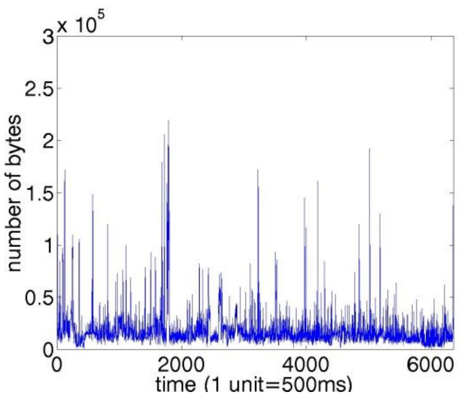

Low multiplexed traffic

- Marginals: highly variable
- Autocorrelation: low

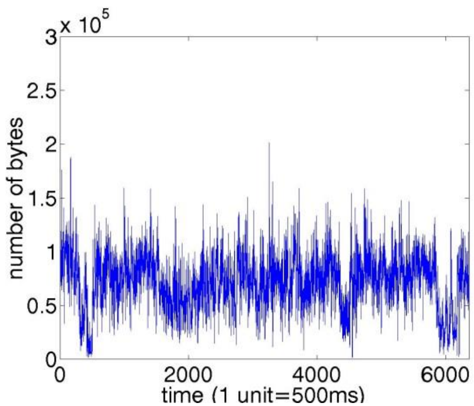

Highly Multiplexed, Bottlenecked Traffic

- Marginals: tending to Gaussian
- Autocorrelation: high

Network Traffic Analytics Framework

❑ Prediction and Machine learning

❖ Linear models

- ARMA, ARIMA (Auto Regressive Integrated Moving Average)

❖ Non-linear models

- GARCH (Generalized Auto Regressive Conditional Heteroskedasticity)
- Gaussian Regression Framework (GRP)
- Neural network: ANN (Artificial Neural Network), FNN (Feedforward Neural Network), RNN (Recurrent Neural Network), ENN (Elman Neural Network), PNN (Propagation Neural Network), MLP (Multi-Layer Perception), etc.

❖ Challenges for Machine learning

- Unlabeled vs. Labeled Data
    - Most commercial successes in ML have come with deep supervised learning
    - There is no large labeled network data sets
- Training vs. {prediction, classification} complexity
    - Stochastic (online) vs. Batch vs. Mini-batch
    - Real-time requirements

Network Traffic Analytics Framework

❑ Real-time traffic visualization

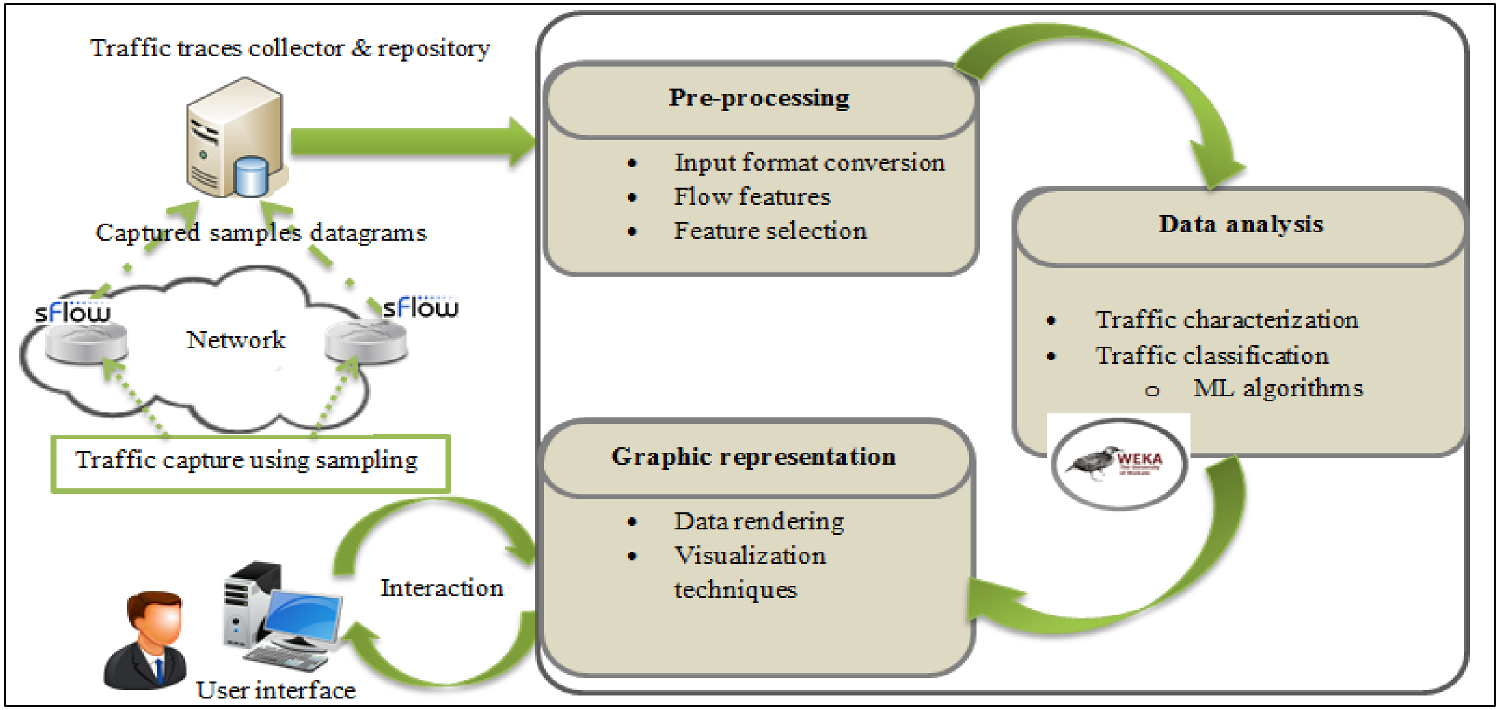

ÉTS
Le génie pour l'industrie

Conclusion

- Smart city network fabric is a complex and energy-inefficient
  - Combining access, core, and cloud
- Network energy efficiency concept is not new and ideas have been around in many different forms
  - Current advances in technology make them feasible: SDN, NFV, FlexE, DVFS, etc.
- Machine learning is a power tool for network energy efficiency
- Cloud is a critical element and consumes lots of energy
  - Optimized scheduling and management is required
  - Can be done through SDN and NFV
  - Solution is multiple, however, complexity is high