Energy Efficient Protocols in Self-Aware Networks

Toktam Mahmoodi

Centre for Telecommunications Research
King's College London, London WC2R 2LS, UK

Stanford NetSeminar
13 December 2011

Why Energy Efficiency in ICT is Important.

- Information and Communication Technology is a major consumer of energy
  - ICT carbon imprint $\simeq$ air travel
- Smart devices + communicate more and travel less $\rightarrow$ rise in energy consumption by communication services.

[1] Global Sustainability Initiative (GeSI), "SMART 2020: Enabling the low carbon economy in the information age", June 2008.

Potentials for Energy Optimisation

1. Admitting the incoming connections to the network
2.
3.

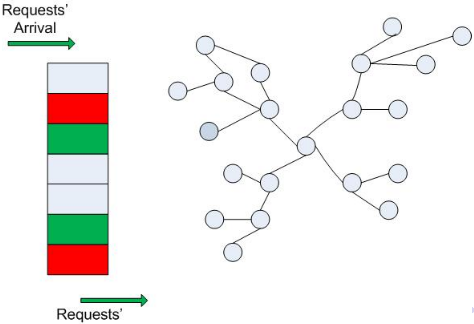

Potentials for Energy Optimisation

1. Admitting the incoming requests to the network
2. Routing the admitted flows
3.

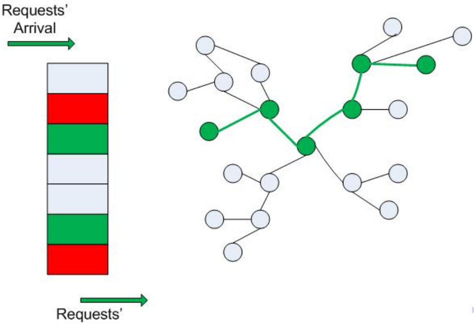

Potentials for Energy Optimisation

1. Admitting the incoming requests to the network
2. Routing the admitted flows
3. Network as involved in the cloud computing

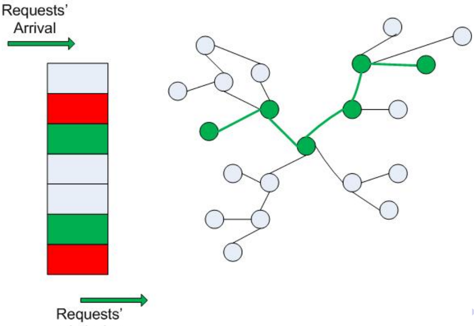

Self-Aware Networks

- Self-Aware Network (SAN) describes a system consisting of nodes that can
  - join and leave the network autonomously
  - sense the status of other nodes, connectivity conditions, traffic level and congestion
  - update their own relevant information about the paths they need to use based on criteria specific to their own needs
- Large amount of information in the network on one hand and different service requirements by users on the other hand, motivate the communication of **local** and just **needed** information rather than communicating updates after any change in the network.

Self-Aware Network Vs. Traditional Network

Locality of Information

Information update can be initiated by the node that need the information rather than throughout the network

On-Demand Information

Information update runs at the time it's needed rather than when change occurs

Outline

- Energy-Aware Routing Protocol
  - Path Selection Criteria
  - Energy Awareness in the Network
  - Experimental Results
  - Further savings can be achieved by switching off nodes.
  - Minimum Energy Topology
- Energy-Aware Admission Control
  - Prioritising Network Admissions based on Energy Consumption.
  - Pricing Network Admissions based on Energy Consumption.
- Management of the Resources in the *Liquid Network*.
  - Example: Migration of Virtual Machines in the Cloud.

Outline

- Energy-Aware Routing Protocol
  - Path Selection Criteria
  - Energy Awareness in the Network
  - Experimental Results
  - Further savings can be achieved by switching off nodes.
  - Minimum Energy Topology
- Energy-Aware Admission Control
  - Prioritising Network Admissions based on Energy Consumption.
  - Pricing Network Admissions based on Energy Consumption.
- Management of the Resources in the *Liquid Network*.
  - Example: Migration of Virtual Machines in the Cloud.

Energy-Aware Routing Protocol

- Routing protocols per se could be optimised based on energy consumption, mainly because nodes may have various,
  - energy related characteristics.
  - sources of power generation at different cost and/or GHG emission.

- Large part of the power consumption is associated to the idle power.
  - Traffic can be distributed such that more routers will be left unused and they can be switched off.
  - The extreme case would be to keep the minimum spanning tree and switch off the remainder of routers that can results in increasing the packet loss and packet latencies.

Energy-Aware Routing Protocol: Path Selection

- Each router along the path route traffic via the outgoing link such that the best value of *Goal* is achieved.
- Every router stores a specific Random Neural Network for each flow traveling through that router.
    - Neurons represent the links.

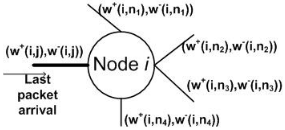

Energy-Aware Routing Protocol: Path Selection

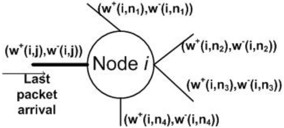

- Reinforcement Learning algorithm is used for updating the weights (w^+ for rewarding and w^- for punishing)
  - weights are changed to reward or punish a neuron according to the level of goal satisfaction measured on the corresponding output.

Power Cost Function

- The *Power Cost* of the $k-th$ flow at node i can be defined as a combination of the flow's own power consumption, and of the impact it has on other flows which are using the node.

$$m_i^k(t_k, T_i) = c \cdot p_i(t_k + T_i) + d \cdot [p_i(t_k + T_i) - p_i(T_i)]. \quad (1)$$

- $c,\ d \geq 0$
- If $c = d = 1$ both of these elements have an equal weight
- If $d = 0$ then we are ignoring the effect on the other flows that are using the node.

Definition of Goal function

- Since EARP is expected to minimise the overall cost of power while satisfying the requested QoS, the goal G_i to be optimised will combine the power consumption with the QoS constraint.

$$G_i = m^k_{\pi(i)}(t_k, \overline{T_{\pi(i)}}) + \beta \cdot 1[Q^k_{\pi(i)}(t_k, \overline{T_{\pi(i)}}) - Q^k_o > 0](Q^k_{\pi(i)}(t_k, \overline{T_{\pi(i)}}))^{\nu}, \quad (2)$$

  - $m^k_{\pi(i)}(t_k, \overline{T_{\pi(i)}})$ and $Q^k_{\pi(i)}(t_k, \overline{T_{\pi(i)}}))$ are the total power cost function and the total QoS value measured on the path going from the $i - th$ node to the destination of flow k.
  - Q^k_o is the QoS value that *should not be exceeded* for flow k.

Energy Awareness in the Network

- Header of data packets should carry the needed information for EARP.
- Each node engaged in forwarding packets to the destination(s) sends out *probing packets* that,
  - gather desired data about these paths (e.g. energy information).
  - help the network to train faster.

Energy Awareness in the Network

- Header of data packets should carry the needed information for EARP.
- Each node engaged in forwarding packets to the destination(s) sends out *probing packets* that,
  - gather desired data about these paths (e.g. energy information).
  - help the network to train faster.

On-demand **Local information**

Networking Testbed

- Our experimental testbed consists of 46 nodes,
  - nodes are Pentium IV-machines with up to fifteen Ethernet interfaces running Linux Kernel 2.6.15.
  - nodes are connected with the full-duplex links at 10 Mbps.

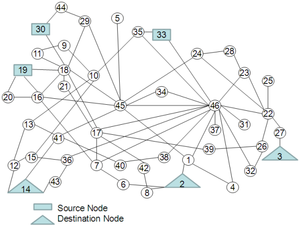

[2] This topology resembles that of the Swiss Education & Research Network (SWITCH)

Experimental Results: Power Consumption

- Nine source-destination pairs are selected such that,
  - three flows are initially activated in the network,
  - during the experiment three additional connections are launched, followed by three more.
  - 20% savings can be observed.

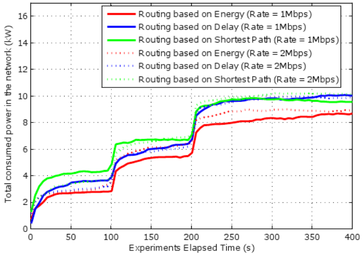

EARP in Summary

- EARP major advantages
  - Distributed Protocol
  - Supporting multiple QoS criterion for each/multiple flows
- EARP major drawbacks
  - Potential increase in Packet Latency, and Packet Loss
  - Extra Computational Complexity
- Conclusion
  - Further savings can be achieved by switching off some nodes.

[3] T. Mahmoodi, "Energy-aware routing in the cognitive packet network," Performance Evaluation-Elsevier, vol. 63, no. 4, pp. 338346, April 2011.

[4] E. Gelenbe, T. Mahmoodi, "Energy-aware routing in the cognitive packet network," In Proc. Int'l Conf. on Smart Grids, Green Commun. & IT Energy-Aware Tech., Venice, Italy, May 2011.

Minimum Power Topology

- Modeling the network with $G = (V, E, w)$ and assuming w_{ij} is the power consumption of packet transmission between two nodes of i and j, computing topology that consumes the least power can be narrowed down to finding the Steiner Minimal Tree.
- The proposed algorithm here is based on the minimum spanning tree and has relatively straight forward implementation (not necessarily low in computational complexity)

Minimum Power Topology

1. Set L as a subset of nodes that are source or destination.
2. Construct the complete graph over L with edge weights equivalent to the smallest sum on G (metric closure G_L)
3. Find the minimum cost tree T_L on G_L.
4. Replace each $e = (u, v) \in E(T_L)$ with the equivalent on G .

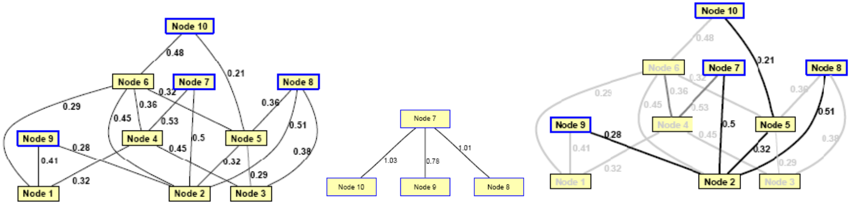

(a) Total power consumption = 6.46 kw, (b) Minimum cost tree T_L, (c) Total power consumption = 2.19 kw

QoS Bounded, Minimum Power Topology

- In addition to the savings in energy, increasing heterogeneity in the current and future networks, also increase the range of QoS requirements by different users.
- Network energy optimisation should also consider QoS parameters such as delay, jitter or packet loss.
  - In this case, we may need to keep more than one path to each of destination node in order to for example keep packet latency, or packet loss below certain threshold.
- We choose delay as the QoS parameter for the example given here.

QoS Bounded, Minimum Power Topology

1. Compute minimum power tree T_p, minimum delay tree T_d.
2. Start from the source node u and traverse T_p in the depth first search order.
3. When destination node v is visited,
   3.1 replace the $u - v$ path with its equivalent in T_d if delay constraint is violated.

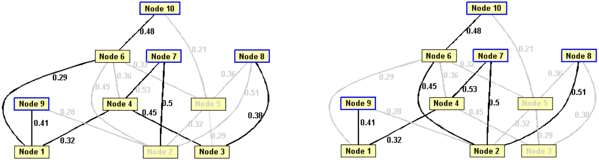

(a) T_d, (b) Total power consumption = 3.14 kw.

Distributed Algorithm

- For the distributed implementation of such an algorithm, nodes can decide to be switched off based on,
  - detecting the active flows and their possible paths (we can assume topology is available offline at all nodes).
  - observing their neighbouring environment, e.g. either their neighbours are on or off.

Distributed Algorithm

- For the distributed implementation of such an algorithm, nodes can decide to be switched off based on,
  - detecting the active flows and their possible paths (we can assume topology is available offline at all nodes).
  - observing their neighbouring environment, e.g. either their neighbours are on or off.

On-demand **Local information**

[5] Preliminary observations are reported in "Delay Bounded Green Network Design", which is submitted to the *IEEE INFOCOM Workshop on Green Networking*.

Outline

- Energy-Aware Routing Protocol
  - Path Selection Criteria
  - Energy Awareness in the Network
  - Experimental Results
  - Further savings can be achieved by switching off nodes.
  - Minimum Energy Topology
- Energy-Aware Admission Control
  - Prioritising Network Admissions based on Energy Consumption.
  - Pricing Network Admissions based on Energy Consumption.
- Management of the Resources in the *Liquid Network*.
  - Example: Migration of Virtual Machines in the Cloud.

Energy Aware Admission Control : Motivations

- The main rationale behind this is that studies show the network power profile has a step-like behaviour. Thus, adding certain traffic to the network,
  - may not change the power level (or apply slight increase).
  - may increase the power consumption largely by moving the power level from one step to another.

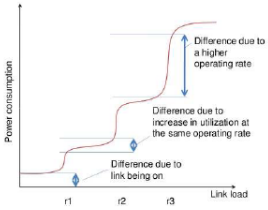

[6] N. Vasic and D. Kostic, "Energy-aware traffic engineering," EPFL Tech. Rep., 2008.

Illustration of the Main Idea

- By postponing admission of some requests and/or reordering the requests, network energy optimisation can be achieved.
- Trade off in this problem is between admission delay and power consumption.
  - In an extreme case, one connection at a time can be admitted in the network to keep its power level minimal.
- Two different solutions are examined here.
  - Prioritise admissions to the network based on their potential energy consumption.
  - Price admissions to the network based on their potential energy consumption.

Prioritising Admissions to the Network

1. Assuming the power consumption behaviour is known, the increase in power consumption upon admitting the new flow can be estimated depending on its bandwidth and the path it travels.
   - computations are based on the shortest path between source and destination nodes.
2. If the estimated wattage increase, δP, is smaller than a fixed value Δ, the flow is accepted and admitted into the network.
   2.1 Unadmitted flows are sent to the waiting queue.
3. New arrivals join the service queue, if at their arrival time server is busy serving other flows.

Prioritising Admissions to the Network

4. In the waiting queue, flows are stored in a ascending order of their remaining waiting time.
   4.1 Admission requests are submitted to the network with the delay value as their maximum accepted delay, which is assigned as their waiting time right after their request.
5. Admission control mechanism serves the service queue first and the waiting queue afterwards.
6. If the waiting time of a flow is due to expire, it is admitted into the network, irrelevantly of its estimated power increase.
   6.1 All flows are guaranteed admission within their predefined waiting time.

Preliminary Experiments in the Networking Testbed

- Accepted admission delay for all flows (four) are 30 sec.
- Average power consumption is reduced by 17%.

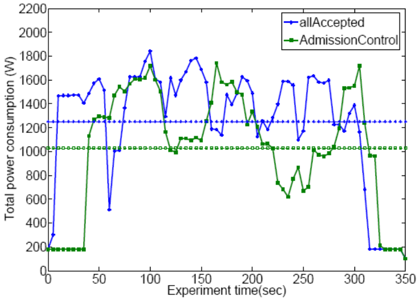

[8] G. Sakellari, C. Morfopoulou, T. Mahmoodi, E.Gelenbe, “Energy-Aware Admission Control for Wired Networks,” submitted to IEEE INFOCOM Workshop on Green Networking.

Distributed Admission Control

- In the distributed mechanism, admission requests are submitted to the source node of that connection.
- Nodes are getting aware of the power consumption of different paths in the network via similar method to the one used in ERAP.
  - Source node generate a probing packet and get all the information upon return of the ACK.

Distributed Admission Control

- In the distributed mechanism, admission requests are submitted to the source node of that connection.
- Nodes are getting aware of the power consumption of different paths in the network via similar method to the one used in ERAP.
  - Source node generate a probing packet and get all the information upon return of the ACK.

On-demand **Local information**

Pricing Admissions to the Network

- Differentiated pricing for the Internet services (rather than flat pricing) has been addressed as a more efficient solution for congestion control.
- We address differentiated pricing for admitting flows to the network.
  - Users define their value for the service depending on their acceptable admission delay.
  - As different users will have different requirements, e.g. different admission delay requirement, they have the incentive of revealing their true value.
  - Those users who value the service less will potentially suffer from larger admission delay.

Pricing Admissions to the Network

- Service provider defines the payoff depending on the increase in the power consumption, δP, users incur to the network.
  - For example if network remain in its current step on the energy curve, service will be priced lower than otherwise.
- The main advantage of pricing rather than prioritising the incoming flows, is that,
  - applications that are more sensitive to delay can potentially get a better service.
  - there will be more incentive for the users to be part of the green infrastructure of the network.

[9] T. Mahmoodi, "Pricing Network Admissions based on the Energy Consumption," submitted to IEEE INFOCOM 2012 Workshop on NetSciCom.

Outline

- Energy-Aware Routing Protocol
  - Path Selection Criteria
  - Energy Awareness in the Network
  - Experimental Results
  - Further savings can be achieved by switching off nodes.
  - Minimum Energy Topology
- Energy-Aware Admission Control
  - Prioritising Network Admissions based on Energy Consumption.
  - Pricing Network Admissions based on Energy Consumption.
- Management of the Resources in the *Liquid Network*.
  - Example: Migration of Virtual Machines in the Cloud.

Moving Resources in the Cloud

- Increasing the cloud services introduce more liquidity in the network.
  - Allocation of services and data communication becomes increasingly complex.
- Resources could be relocated from their physical location with the aim to reduce energy consumption or the GHG.
- One example would be migrating a VM from to a physical server in another DC, when,
  - the second DC has access to the renewable energy sources.
  - the destination physical server is faster, or more energy efficient.

Migration of the Virtual Machines

- Energy optimisation policies in a federated data centres' scenario may decide to move a VM from physical server A to B.
  - Extra energy in the network will be consumed for the move.
  - An estimation of network energy is needed to confirm this migration decision.

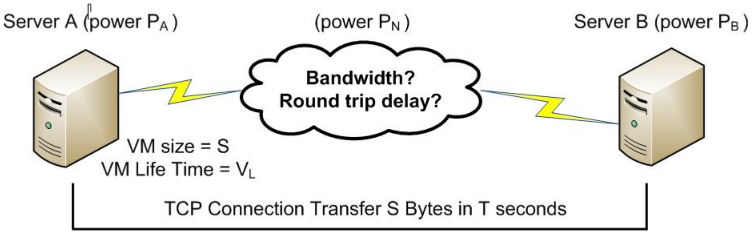

Migration of the Virtual Machines

- Information such as bandwidth, round trip delay and number of hops in the path are required for this estimation.
  - Transfer time can be estimated e.g. based on a TCP connection.

- Due to the variety of options in the networking equipments in the path, offline information can be used.
  - By sending *Probing packets* the related data can be gathered and brought back.
  - If power consumption related data is not available at each hop (which is the case in the current network), the estimation is provided based on number of hops.

[7] T. Mahmoodi et al., "Description of Energy Optimisation Mechanisms and Policies," FIT4Green WP4 Deliverable, May 2011.

Summary

- Network policies such as routing and admission control can be modified to achieve energy savings.

Summary

- Network policies such as routing and admission control can be modified to achieve energy savings.
  - Different routers and network connections can have different energy consumption characteristics.

Summary

- Network policies such as routing and admission control can be modified to achieve energy savings.
  - Different routers and network connections can have different energy consumption characteristics.
  - The existence of multiple power states in the nodes, and in the network, pose the step-like behaviour to their power consumption.

Summary

- Network policies such as routing and admission control can be modified to achieve energy savings.
  - Different routers and network connections can have different energy consumption characteristics.
  - The existence of multiple power states in the nodes, and in the network, pose the step-like behaviour to their power consumption.
  - The step-like behaviour motivates differentiated prioritising and differentiated pricing in the network's admission policies.

References

[1] Global Sustainability Initiative (GeSI), "SMART 2020: Enabling the low carbon economy in the information age", June 2008.

[2] Swiss Education & Research Network (SWITCH): www.switch.ch

[3] T. Mahmoodi, "Energy-aware routing in the cognitive packet network," Performance Evaluation-Elsevier, vol. 63, no. 4, pp. 338346, April 2011.

[4] E. Gelenbe, T. Mahmoodi, "Energy-aware routing in the cognitive packet network," In Proc. Int'l Conf. on Smart Grids, Green Commun. & IT Energy-Aware Tech., Venice, Italy, May 2011.

[5] T. Mahmoodi, "Delay Bounded Green Network Design," submitted to IEEE INFOCOM 2012 Workshop on Green Networking.

[6] N. Vasic and D. Kostic,"Energy-aware traffic engineering," EPFL Tech. Rep., 2008.

[7] T. Mahmoodi et al., "Description of Energy Optimisation Mechanisms and Policies," FIT4Green WP4 Deliverable, May 2011, www.fit4green.eu.

[8] G. Sakellari, C. Morfopoulou, T. Mahmoodi, E.Gelenbe, "Energy-Aware Admission Control for Wired Networks," submitted to IEEE INFOCOM 2012 Workshop on Green Networking.

[9] T. Mahmoodi, "Pricing Network Admissions based on the Energy Consumption," submitted to IEEE INFOCOM 2012 Workshop on NetSciCom.

Questions?

Thanks for your attention.

Contact: toktam.2.mahmoodi@kcl.ac.uk