Enabling Energy Efficient and Resilient Networks using Dynamic Topologies

Abdelnour Aldraho
University of Southern Queensland
Toowoomba
Australia
Email: abdelnour.aldraho@usq.edu.au

Alexander A. Kist
University of Southern Queensland
Toowoomba
Australia
Email: kist@ieee.org

***Abstract*—Communication networks are dimensioned for peak loads, but load fluctuates and networks are often underutilised. Energy consumption of routers and communication links is largely load independent, i.e. lightly loaded network have similar power consumptions than highly loaded networks. In light of the global focus on energy efficiency and greenhouse gas emissions, ideally, networks exhibit load proportional energy profiles. At the same time it is imperative that communication networks are resilient to link failure. This paper introduces the concept of resilient dynamic topologies where networks have the ability to reduce the number of active nodes and links during lightly loaded periods while maintaining resilience to link failures. Mathematical programming models are presented that result in reduced optimal network configurations. Using a sample network, energy consumption of networks with dynamic topologies and resilient dynamic topologies are compared. It is concluded that reduced resilient topologies with smaller energy footprints are feasible.**

I. Introduction

Energy efficiency and network resilience have become a key focus of the commercial as well as research networking community. In this work we explore traffic engineering and routing of IP networks to overcome energy wastage by employing dynamically changing network topologies with a strong emphasis on network resilience. With the prevalence of IP networks, resilience, i.e. the ability of a network to sustain faults or disruptions, has become a major concern for service providers. With reference to industry surveys, Kalmanek and Yang [1] highlights that reliability, network usability and network fault protection are key concerns of major network carriers. Failure can be caused by fibre cuts, for example, but also by natural disasters, such as floods, cyclones and bush fires in Australia.

ICT infrastructure has also received growing attention in regards to greenhouse gas emissions, both as an enabler for energy reductions but also as an industry with its own footprint [2]. It has been estimated that network related emissions account for 30% [3] to 37% [2] of the sector's contributions. Efforts to improve energy efficient in telecommunication systems spans from hardware developments [4, e.g.] to traffic engineering [5, e.g.].

To reduce energy consumption of ICT systems on a large scale, a number of areas have to be addressed, including infrastructure, system hardware and software; and intelligent resource management. In this work, we target intelligent resource management. Specifically, we investigate the possibility of reducing the number of active routers at times of lesser load and therefore the overall power consumption. In this scenario, networks change their topology actively to adapt to current traffic conditions and at the same time ensure that networks remain resilient to faults.

Our vision is to investigate and develop innovative self configuring communication networks with load proportional power profiles that reduce the carbon footprint of ICT network infrastructure. It is a building block in the larger endeavour of the research community and industry to develop sustainable systems that enable major GHG reductions in communication networks and other sectors. However, any practical network reduction to account for energy consumption has to take resilience into account.

At a basic level, the target is to design a network consisting of nodes and links that exhibit traffic proportional power profiles. More specifically, we propose models that result in resilient and energy efficient topologies and demonstrates that impact such changes have on the power consumption of the network. Previous work [6] has introduced a network transformation that adds a number of artificial links to determine which routers can be turned off.

In contrast to most other publications, [7, e.g.] and [8, e.g.] this research does not only allow for a reduced number of links, but also reduced the number of nodes that originate or terminate traffic. If nodes are turned off, it is important that local demands are bridged to neighbouring routers. To be able to do this it is assumed that routers support a standby state where all local access links are bridged to one neighbouring node.

The key contributions of this paper include the ILP formulation of a resilient, energy efficient dynamic topology problem, the initial evaluation and the insight that resilient, energy efficient dynamic topologies are feasible.

The remainder of this paper is organised as follows: Section II discusses related work and shows how this proposal fits into the current body of work, Section III introduces the underlying power model that is used for this project and Section IV explains the network model and notation. Section

V presents two Integer Linear Program (ILP) formulations that solve the problem of energy efficient, resilient networks. Evaluation results are introduced in Section VI and the paper concludes with Section VII.

II. Related Work

For server and router systems, energy consumption has been a concern, in regards to the ability of supplying sufficient power, heat dissipation and cooling. Energy efficiency in the context of networking has been targeted by a number of studies. For example, power savings in LAN switches are addressed by [4]; and rate adaptation and sleeping stages in networks are discussed by [9]. Baliga et al. in [10] compare energy consumption of various technologies in the access network and concludes that optical-access is the most energy efficient. Kist [11] proposes a mechanism that allows load distribution for large scale server clusters with load proportional GHG emission footprint and includes mathematical performance analysis of a simple routing mechanism to manage server workload.

The term dynamic topology has been used by the networking community in the past, in particular, in context of circuit switched networks. Noakes et al. [12], for example, present "... an adaptive link assignment algorithm for distributed optimization of dynamically changing network topologies..." and a related routing algorithm is discussed by [13]. These studies address different technologies and optimisation problems and are not directly applicable to energy efficient network configuration in IP networks.

A number of authors have identified power consumption as a main issue in high performance router design [14, e.g.]. Solutions include the use of optics in routers [15] and energy efficient switching fabric design [16]. Power awareness in network design and routing to mitigate power consumption of high speed network equipment is also advocated by [17] who undertook benchmarking of two routers to estimate power use and developed a general model for router energy consumption and formulated the network design problem as a power aware mixed integer program. Optimisation focuses on allocation of line cards per chassis and chassis over the target network. Chiaraviglio et al. [7] introduces a network design problem with the aim to reduce the total power consumed by the network. It outlines the optimisation problem as a linear program and proposes heuristics to solve the problem. The study in [7] focused on a specific star network topology with three different aggregation levels, while the work in this paper and in [6] looks at a generic network topology also, the node capability is different.

In this particular scenario, only nodes and links that are redundant can be turned off. The problem formulation is similar to classic network design problems with an energy consumption objective. Chiaraviglio el at. [18] uses the results of [7] and evaluates an operator topology with realistic power usage figures for devices and proposes a new algorithm that accounts for power consumption of devices. Chiaraviglio el at. [19] focuses on a much larger scale and uses analytical and simulation results to estimate redundant network resources worldwide, that could be turned off to save power. All these studies do not address two important aspects that are subjects of this work: network nodes do not originate or consume traffic; and network resilience is not considered.

A number of studies target specific applications or techniques. Chiaraviglio et al. [5] targets energy aware network management in radio access networks. Vasic el at. [20] presented an "online energy-aware traffic management technique" that assumes equipment is able to adapt its power use to utilisation using rate adaptation and sleep states. Wu et al. [21] discusses the routing and wavelength assignment problems with the new objective of reducing network energy consumption. The authors introduce an ILP formulation and a number of heuristics to solve the problem. Node and link power down for unused connections in IP over WDM networks is also discussed by [22]. These approaches provide solutions for energy efficient traffic management; they do not address the dynamic topology problem and resilience. In an attempt to address resilience and energy efficiency, [8] introduces the BlueGreen network design philosophy; however, the efforts are limited to links. Kist and Aldraho [6], have introduced an optimaisation problem formulation that are reduces the power consumption of networks in case were nodes originate and terminate. This paper extends this work by taking resilient constraints into account.

Resilience in IP networks has been an important research area [23] and a number of mechanisms have been suggested that focus on reachability and provide fast rerouting around failed network devices [24, e.g.]; congestion and performance are not widely addressed. However, these aspects are particularly important as hotspots and overloaded links are largely caused by link failures [25, e.g.]. Overlapping failures [26] can lead to a large number of failure scenarios, mechanisms that rely on their enumeration are therefore not practical. Wang et al. [27] proposes Resilient Routing Reconfiguration (R3) that overcomes most of the limitations of the reactive approaches by employing Multiprotocol Label Switching (MPLS) to setup a number of alternative paths between origin–destination pairs sharing traffic load. None of the studies addresses energy efficient technologies.

III. Power Model

The power consumption of routers varies with capacity, vendor and features. To demonstrate the impact of the models, arbitrary, but realistic values for a generic router and links are required. The following set of basic assumptions is used: 10% of the router energy consumption is load dependent; i.e. caused by function, such as routing table lookups, queuing, and forwarding; and 90% of the total power consumption is load independent. These assumptions are supported by measurements undertaken by [17].

Furthermore, it is assumed that routers have two modes of operation, independent of line card configurations: A fully functional mode and a standby mode that does not support routing functionality. It is assumed that in standby mode, the router consumes 10% of the maximum power. Load dependent power consumption in low power modes is assumed to be negligible.

Power consumption of line cards is attributed to links; i.e. two line cards are required per link. Each line card has only one port and line cards can be individually activated and deactivated. As the utilisation of network processors depends on traffic load, a higher variability of 20% is assumed.

It is also assumed that link capacity is not a major factor in power consumption, i.e. technology at a similar level of maturity, have similar power requirements. These assumptions are supported by [28] and [29]. Numerical values are only indicative as changes in the absolute values are inconsequential to the overall trends.

Furthermore, it is assumed that routers without line cards consume a maximum of 600 Watts, i.e. 540W are independent of load[1]. A single, fixed power ratio has been used as variations have no direct impact on the principal outcome. For the calculations in this paper, it is assumed that links have a maximum power consumption of 80 Watts[2]. Variable power consumption is scaled to link speed. An unloaded link consumes 80% of the total power and a fully loaded link consumes 100% of the link power. These assumption are the same that have been used by [6]. The numerical values are only indicative and changes in the absolute values are inconsequential to the overall trends that are discussed in the remainder of the paper. If these models are applied to specific scenarios and technologies, assumptions can be revised accordingly.

IV. Network Model and Notation

This study focuses on IP backbone networks at the node level. The optimisation problems address individual nodes, i.e. routers or sites; however, it does not address individual chassis components. For optimised topologies, nodes can represent individual routers or Point of Presence (PoP) sites with a number of routers. It is assumed that traffic is either routed by OSPF or MPLS.

For the mathematical modelling, the following notation is used: A network $G(N,M)$ consist of N nodes and M directed arcs. The flow of commodity k on arc (i,j) is denoted as x_{ij}^k and the unit costs of commodity k using arc (i,j) as c_{ij}^k. Arc (i,j) has also a fixed cost $g_{i,j}^k$. This cost is encountered if link (i,j) is active and it is independent of the traffic x_{ij}^k. The capacity of arc (i,j) is denoted by u_{ij}. Similarly, nodes have costs c_i^k, g_i^k and capacities u_i[3].

The constant $b^k(i)$ denotes supplies or demands at node i. The variables δ_{ij} and δ_i are Boolean values that indicate if arc (i,j) and node i are in use, respectively. Integer constants γ_i and β_i indicate the maximum out-degree and in-degree of nodes, respectively. $y_{i,j}^{m,n}$ denotes the flow of commodity (m,n). (m,n) indicates the link between nodes n and m that is protected by the flows $y_{i,j}^{m,n}$. $\gamma_{i,j}$ is a Boolean constant that indicates if link (i,j) is protected. $v_{i,j}$ is the maximum flow that needs to be reserved on link (i,j) to protect the network against single link failure.

To account for node costs and node capacities in the problem formulations, the node splitting transformation is applied [30]. Nodes i are replaced by a set of additional nodes i' and i''. These are connected by an arc (i',i''). Node cost and capacity are assigned to arc (i',i''). Terminating arcs are connected to i', and emanating arcs are connected to nodes i''. Using this transformation, node constrains will appear as link constraints in the problem formulation.

To be able to develop problem formulations that have the ability to remove nodes, an extension to the node splitting transformation is necessary [6]. A set of a set of duplicated emanating and terminating arcs connected to i' and i'' are added. These bypass the node's routing functionality, i.e. link (i',i''). Traffic that is emanating and terminating at nodes that are turned off is bridged via these links to neighbouring nodes. Constraints in the problem formulation ensure that either the original arcs or one artificial link are active. The transformation is discussed in [6] in detail.

V. Problem Formulation

This section introduces two ILP models that result in reduced topologies that are resilient against single link failures of active links. To protect against link failures, network can either protect the capacity of individual links; or protect all demands. The former leads to simpler optimisation problems, the latter require less additional capacity.

The aim of these problem formulations is to turn off as many nodes as possible to reduce the fixed power consumption of nodes. Local demands of nodes that are turned of are bridged to neighbouring (active) nodes. At the same time the aim is to protect the remaining network against single link failures. This excludes links to nodes that are operating in

[1]The magnitude of this value is based on [17]. The authors report absolute power values of unloaded systems of 430W. As nodes in this model inject and consume traffic, node power consumption has to also account for an access link.

[2][17] reports 26W for a 1 port Fast Ethernet line card, 30 W for a 1 port Gigabit Ethernet line card and 92W for a 4 port Gigabit card (23W per port). Further discussions in the remainder of this paper assume a total of 32W fixed power consumption per port (40W total) for generic links. The power consumption of a link is the sum of the power of the two line cards at the link endpoints. The power consumption of network interfaces is decreasing as technology advances, these values reflect conservative assumptions.

[3]In telecommunication networks, costs are the same for all k. For this application, different commodities correspond to alternate traffic flows, between other origin destination pairs and the cost is the same for all pairs.

standby. A link failure for access links to these stub nodes requires the router to power up.

A. *Problem (1) - Links are protected*

The following optimisation problem protects the capacity of all links. The objective function is given in Equation (1).

$$\text{Minimise } Z = \sum_{kij} c^k_{ij} x^k_{ij} + \sum_{ij} \delta_{ij} g_{ij} + \epsilon \sum_{ijnm} y^{nm}_{ij} \quad (1)$$

The aim is to minimise energy consumption, i.e. proportional costs $c^k_{i,j}$ and fixed costs g_{ij} of active nodes and links. Backup traffic y^{nm}_{ij} does not cause any power consumption in operational networks. Only when links fail these flows are used. However, additional traffic in the case of a failure should be minimised. Hence the scaling factor ϵ, a small number, is used.

The following constraints are required. Links have to accommodate normal traffic and reserve same capacity for traffic that is caused by failing links. The balance constraints for both traffic sets are independent. Equation (2) shows the mass balance constraint for normal traffic.

$$\mathcal{N}\boldsymbol{x}^k = \boldsymbol{b}^k \quad \forall k \in K \quad (2)$$

$\mathcal{N}$ is the node-arc incidence matrix and $\boldsymbol{b}$ is the right hand side vector, the vector that specifies supplies and demands. Equation (2) expresses the conservation of flow, as the sum of all elements $b(i)$ in $\boldsymbol{b}$ must be equal to zero. In the case of network traffic flows, there are only two non zero elements for each $\boldsymbol{b}^k$, the flow source $b(s)$ and the flow destination $b(t)$. The commodities, k, correspond to the demands between source and destination nodes.

Equation (3) depicts the mass balance constraint for backup traffic that is caused by the failure of link (n, m).

$$\mathcal{N}\boldsymbol{y}^{mn} = \delta_{mn}\boldsymbol{a}^{mn} \qquad \forall (n, m) \in A \quad (3)$$

a^m, n is a right hand side vector with a positive supply of u_{ij} in row i and a negative demand of $-u_{ij}$ in row j for all links (m, n), where,u_{ij} is the capacity of links (i, j). If the link (m, n) doesn't exists, $\boldsymbol{a}^{m,n}$ equals zero.

Equation (4) show the bundle constraint. Links have to accommodate all normal flows x^k_{ij} and the backup capacity v_{ij}. the variable δ_{ij} are boolean value that indicate if arc (i, j) in use.

$$\sum_K x^k_{ij} + v_{ij} \leq \delta_{ij} u_{ij} \qquad \forall (i, j) \in A \quad (4)$$

The backup capacity constraint is depicted in Equation (5).

$$y^{nm}_{ij} \leq v_{ij} \quad \forall (n, m), (i, j) \in A \quad (5)$$

v_{ij} denotes the upper limit for any backup flow on link (i, j). A failing link (n, m) can not be part of the backup topology. This is enforced by the disjointness constraint in Equation (6).

$$y^{m,n}_{i,j} \gamma_{i,j} = 0 \quad \forall (i, j) \in A, (n, m) \in A | m = i, n = j \quad (6)$$

Flow y on link (i, j) protects link (m, n) and therefore is not allowed to use link (m, n). γ_{ij} indicates if a link is protected. If γ is equal to zero, the link can not accommodate backup flows. It is one for all links, but artificial links.

The constraints in Equation (7) and (8) enforce the restrictions that are required for the extended transformation.

$$\sum_j \delta_{ij} = \alpha_i \quad \forall\ i \in N \quad (7)$$

$$\sum_j \delta_{ij} = \beta_j \quad \forall\ j \in N \quad (8)$$

α_i and β_j limit the number of links that emanate and terminate at nodes and enforce constraint that are important for extended network transformation.

B. *Problem 2 - Demands are protected*

This ILP protects all demands not individual links. It adds another dimension to the backup flow variable y. The new objective function is depicted in Equation (9).

$$\text{Minimise } Z = \sum_{kij} c^k_{ij} x^k_{ij} + \sum_{ij} \delta_{ij} g_{ij} + \epsilon \sum_{ijnm} y^{nmk}_{ij} \quad (9)$$

Variable y that represents the backup traffic is the only change. Equations (2), (4), (7) and (8) remain the same. Equation 10 shows the updated mass balance constraints for backup traffic caused by the failure of link (n, m) for demand k.

$$\mathcal{N}\boldsymbol{y}^{mnk} = \delta_{mn}\boldsymbol{b}^k \qquad \forall (n, m) \in A \quad (10)$$

Disjointness is enforced by Equation (11).

$$y^{m,n,k}_{i,j} \gamma_{i,j} = 0 \quad \forall (i, j) \in A, (n, m) \in A | m = i, n = j \quad (11)$$

Equation (12) ensures that v_{ij} marks the upper bound for all backup demands. v_{ij} corresponds to the capacity on link (i, j) reserved for backup flows.

$$\sum_k y^{nmk}_{ij} \leq v_{ij} \quad \forall (n, m) \in A \quad (12)$$

This problem is very similar to Problem (1). The key difference is the size: y is increased by an additional dimension.

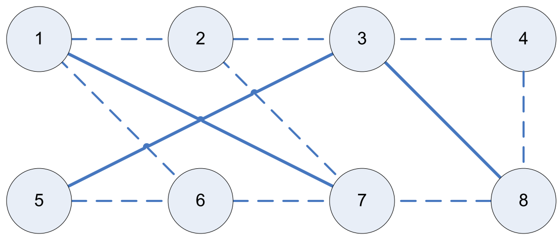

Figure 1. Test topology – 8 nodes.

VI. Evaluation

This section discusses numerical examples demonstrating how resilience constraint and energy efficiency impact on practical network configurations. Results for different optimisation problems are compared:

1) *links only* – all nodes are active only links can be turned off;
2) *links only (resilient)* – only links can be turned off, the remaining network is resilient to single link failures;
3) *nodes and links* – both links and nodes can be turned off; and,
4) *nodes and links (resilient)* – both links and nodes can be turn of, the network can cope with single link failures.

Results for (1) and (3) are obtained using the optimisation problems introduced in [6]. Results for options (2) and (4) are based on Problem (1) discussed in Section V. For option (3) the ILP is used with an untransformed network, for option (4) the transformed network is used. The mathematical programs were solved using the commercial IBM ILOG CPLEX Optimizer 12.4.

A. Network and Traffic

To limit the computing time, a small network with 8 routers and 24 uni-directional links was chosen. The topology is depicted in Figure 1 and inspired by the Australian Telstra network (AS1221). Larger network topologies examples can show the ability to apply the proposed mathematical approach in network resilience more than the smaller topologies, however, the larger network topologies is required higher specification machine and longer time than that is consuming by the smaller topologies to solve the problem and get the results. Links in this network have nominal capacities of 1 Gbps (dashed lines) and 10 Gbps (full lines), respectively. If link (ij) exists, link (ji) also exists. Nodes in this network have a capacity of 100 Gbps and do not pose a bottleneck. Traffic matrices include 56 demands, i.e. traffic between all routers. The experiment uses 3327 instances of traffic matrices which have been generated randomly. The matrices include traffic demands between origin and destination nodes.

Realistic traffic matrices feature a demand distribution that reflects the size of the nodes in terms of connected link capacities. These effects have been captured by gravity models such as [31] and [32]. Uneven traffic distributions are advantages for dynamic topologies. Traffic is accumulated at fewer, highly loaded nodes and therefore results in more unloaded nodes that can be put into standby modes. For this study, random traffic matrices were chosen as a worst case scenario.

B. Resulting Topologies

This section presents a comparison of topologies for one lightly loaded traffic instance. The maximum total load of the network is approximately 10.29 Gbps. The traffic instance used, loads the network with 2.64 Gbps – 26% load. The examples show how the constraints impact on the network topology.

The original network is depicted in Figure 1, the topologies for *links only* and *links only (resilient)* are depicted in Figure 2 and 3 respectively. The arrows in Figure 2 indicate that all but one link in this case are unidirectional. The network forms two overlapping rings. As the links are directional, one disconnected link will impact on a number of nodes. This solution is not fault tolerant. In all other cases all links are bidirectional and arrows are omitted.

The solution in Figure 3 for *links only (resilient)* results in a classic ring network. Without capacity limitations, this is the expected result.

Figure 4 and 5 depict the topologies for case were both, nodes and links can be turned off. The former depicts the energy efficient, the latter the resilient solution. Orange (dark) nodes indicate active routers; lightly shaded nodes indicate in-active routers in standby. In-active nodes are connected to active nodes via one hop. In-active links have been omitted from the diagram.

Option (3) results in multiple connected star topologies that do not include meshing links. As this topology does not include any redundant links, it is vulnerable to link failure. However, links failures have a lesser impact than in the case of option (1). For the resilient case, active routers form a ring network, the most efficient, resilient configuration. Related power consumption and the number of active nodes and links are discussed in the next section.

C. Energy Consumption

Figure 6 depicts the total network power consumption in Watt against the total demand in Gbps for the optimisation problems were only links can be activated or deactivated – options (1) and (2). Results for the unmodified network are given by the cloud on top – (red) square markers (■); for the reduced topology by the lower cloud – (green) triangles (▲) and the cloud in-between indicates the results for the resilient network with (blue) lozenges (♦). The energy consumption increases near linear with traffic load.

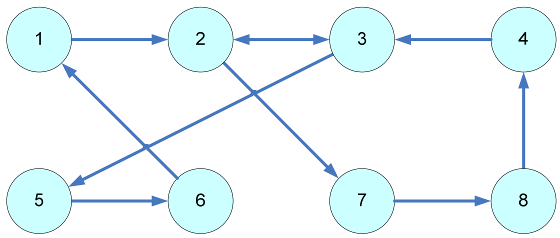

Figure 2. Reduced topology – *links only* - (1).

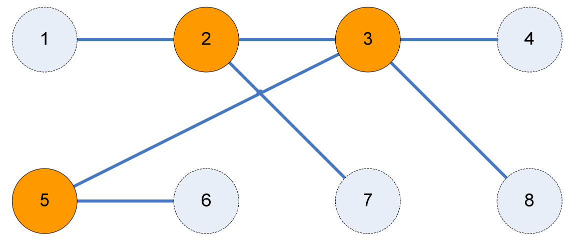

Figure 4. Reduced topology – *nodes and links* - (3).

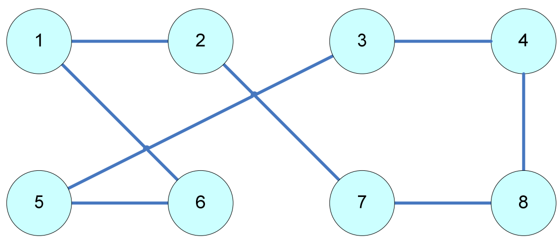

Figure 3. Reduced topology – *links only (resilient)* - (2).

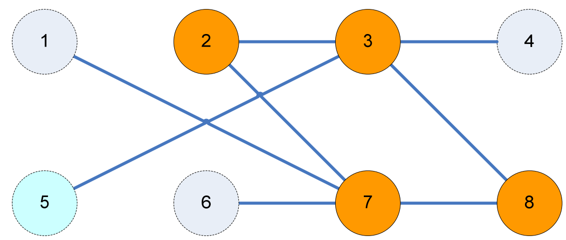

Figure 5. Reduced topology – *nodes and links (resilient)* – (4).

The small steps are cause by the fixed power component of links that are being activated. The power consumption for the resilient problem is less than the unmodified network, but more than the reduced topology. The resilient solution requires more active links. Above a total load of 6 Gbps total load resilient solutions no longer exist.

Figure 7 depicts the total network power consumption in Watt versus the total demand in Gbps for options (3) and (4). The consumption of the unmodified network is indicated by the top cloud of (red) square markers (■), the reduced topology by the lower cloud of (green) triangles (▲) and the cloud in-between show the results for the resilient network with (blue) lozenges (⧫). The original network consumes the most energy and the use increased linearly with load. Option (4) uses the least amount of energy. The graph shows distinct steps where the number of nodes changes. For higher loads more active nodes are necessary. The power consumption for the remaining option (3) lies is between these two extremes.

This graph also exhibits step changes when the number of active nodes change. For most instances above 60 Mbps total traffic, no feasible solutions exist. The optimisation problem (4) does no longer yield feasible solutions. This indicates that these traffic matrixes result in problems that do not allow for additional backup flows.

Figure 8 depicts the number of active nodes versus the total traffic demand. Square markers (■) indicate the energy efficient option, lozenges (⧫) the resilient option. Instances where the number of active nodes is zero indicate that the instance has no feasible solution for the resilient optimisation problem. For infeasible solutions, normal traffic demands use the links and sufficient backup capacity is not available. The total load this network can carry is approximately 10.29 Gbps, the highest load the resilient topology can accommodate is 6.86 Gbps - about 67%. Infeasible instance begin to occur at about 60% load. This agrees with the rule of thumb to reserve 50% of link capacities in case of link failures.

The graph also shows that for reduced topology, for all but one instance, seven nodes are sufficient. For the resilient case, eight nodes are only necessary in a few instances.

VII. Conclusion

Power consumption and reliability of communication networks are important aspects. This work has introduced a mathematical model to demonstrate energy savings that can be achieved using dynamic topologies, i.e. changing the network topology according to load. Earlier work has shown that dynamic topologies can lead to considerable energy savings, this work has shown that this is also possible while a resilient active core of nodes is maintained. Further work will explore a number of avenues . It is potentially possible to develop more efficient mathematical models than the current ILP formulations. This might also include appropriate heuristics. More efficient formulations allow for larger networks. Practical application and the potential implementation of dynamic, resilient topologies will also be addressed as future work.

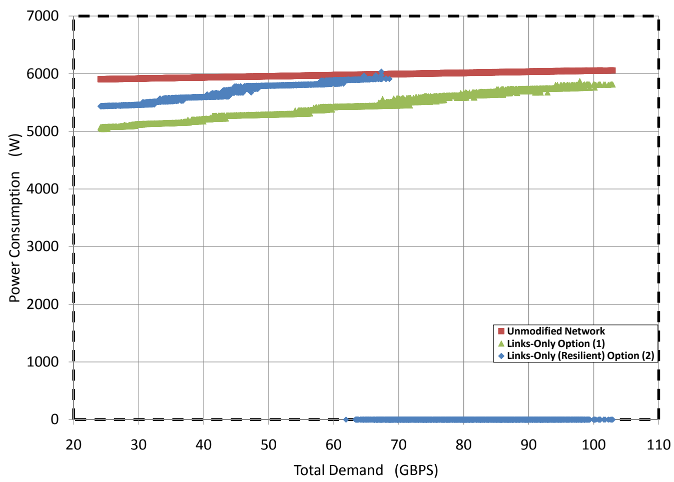

Figure 6. Power consumption versus total demand, *links-only* options.

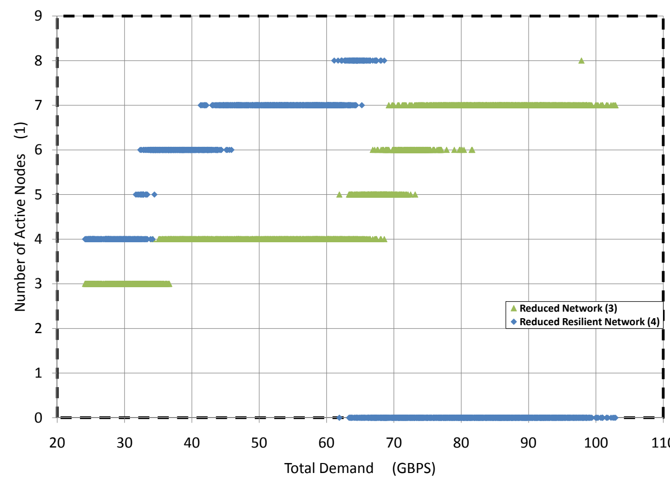

Figure 8. Number of active nodes versus total demand.

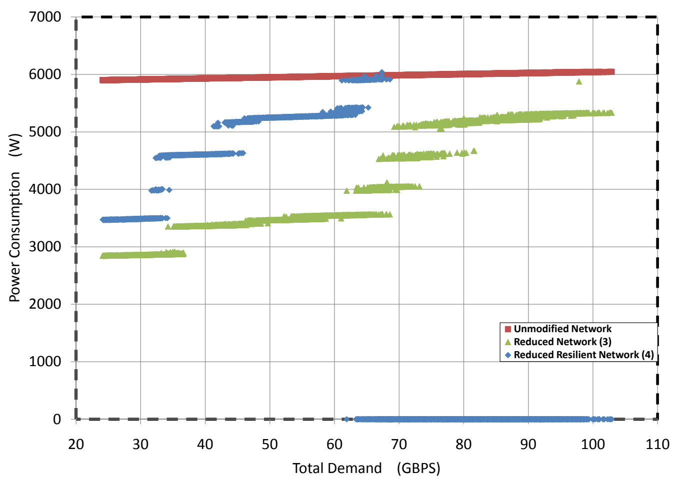

Figure 7. Power consumption versus total demand, *nodes and links* options.

From a topology perspective, energy efficient and resilient dynamic topologies are feasible and potentially enable significant energy savings. Wether such technique can be used in live networks will largely depend on implementation and operational factors.

References

[1] C. R. Kalmanek and Y. R. Yang, "The challenges of building reliable networks and networked application services," in *Guide to Reliable Internet Services and Applications*, ser. Computer Communications and Networks, C. R. R. Kalmanek, S. Misra, and Y. R. R. Yang, Eds. Springer London, 2010, pp. 3–17.

[2] T. C. group and G. e Sustainability Initiative, "Smart 2020: Enabling the low carbon economy in the information age," The Climate group and Global e-Sustainability Initiative, Tech. Rep., 2008.

[3] Gartner, "Gartner says data centres account for 23 per cent of global ICT CO2 emissions," Press release, 11 October 2007. [Online] at http://www.gartner.com/it/page.jsp?id=530912, 2007.

[4] G. Ananthanarayanan and R. H. Katz, "Greening the switch," December 2008.

[5] L. Chiaraviglio, D. Ciullo, M. Meo, and M. Marsan, "Energy-efficient management of UMTS access networks," in *Teletraffic Congress, 2009. ITC 21 2009. 21st International*, Sept. 2009, pp. 1–8.

[6] A. A. Kist and A. Aldraho, "Dynamic topologies for sustainable and energy efficient traffic routing," *Computer Networks*, vol. 55, no. 9, pp. 2271–2288, 2011.

[7] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," *Communications, 2009. ICC '09. IEEE International Conference on*, pp. 1 –6, June 2009.

[8] W. Liu, H. Sirisena, and K. Pawlikowski, "Building dependable next generation networks (ngns): A new bluegreen design philosophy," in *Australasian Telecommunication Networks and Applications Conference (ATNAC), 2011*, nov. 2011, pp. 1 –6.

[9] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," in *NSDI'08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*. Berkeley, CA, USA: USENIX Association, 2008, pp. 323–336.

[10] J. Baliga, R. Ayre, W. V. Sorin, K. Hinton, and R. S. Tucker, "Energy consumption in access networks," in *Optical Society of America*, 2008, optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference.

[11] A. A. Kist, “Modelling of staged routing for reduced carbon footprints of large server clusters,” *Int. J. Communication Networks and Distributed Systems*, vol. 6, no. 1, pp. 28–44, 2011.

[12] M. Noakes, J. Cain, J. Nieto, and E. Althouse, “An adaptive link assignment algorithm for dynamically changing topologies,” *Communications, IEEE Transactions on*, vol. 41, no. 5, pp. 694–706, May 1993.

[13] J. B. Cain, S. L. Adams, M. D. Noakes, T. Kryst, and E. L. Althouse, “A ”near-optimum” multiple path routing algorithm for space-based sdi networks,” in *IEEE Military Communications Conference (MILCOM) 1987*, vol. 2, Oct. 1987, pp. 0578–0584.

[14] C. Minkenberg, R. P. Luijten, F. Abel, W. Denzel, and M. Gusat, “Current issues in packet switch design,” *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 1, pp. 119–124, 2003.

[15] I. Keslassy, S.-T. Chuang, K. Yu, D. Miller, M. Horowitz, O. Solgaard, and N. McKeown, “Scaling internet routers using optics,” in *SIGCOMM ’03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM, 2003, pp. 189–200.

[16] A. Wassal and M. Hasan, “Low-power system-level design of vlsi packet switching fabrics,” *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 20, no. 6, pp. 723 –738, Jun 2001.

[17] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, “Power awareness in network design and routing,” *Proceedings of IEEE INFOCOM 2008*, April 2008, phonix, USA.

[18] L. Chiaraviglio, M. Mellia, and F. Neri, “Energy-aware backbone networks: A case study,” in *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, June 2009, pp. 1–5.

[19] L. Chiaraviglio, D. Ciullo, E. Leonardi, and M. Mellia, “How much can the internet be greened?” in *GLOBECOM Workshops, 2009 IEEE*, 30 2009-Dec. 4 2009, pp. 1–6.

[20] N. Vasic and D. Kostic, “Energy-aware Traffic Engineering,” EPFL, Tech. Rep., 2008.

[21] Y. Wu, L. Chiaraviglio, M. Mellia, and F. Neri, “Power-aware routing and wavelength assignment in optical networks,” in *Optical Communication, 2009. ECOC ’09. 35th European Conference on*, Sept. 2009, pp. 1–2.

[22] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, “Saving energy in ip-over-wdm networks by switching off line cards in low-demand scenarios,” *Optical Network Design and Modeling (ONDM), 2010 14th Conference on*, pp. 1 –6, feb. 2010.

[23] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, “Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines,” *Comput. Netw.*, vol. 54, pp. 1245–1265, June 2010.

[24] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot, “Feasibility of ip restoration in a tier 1 backbone,” *Network, IEEE*, vol. 18, no. 2, pp. 13 – 19, 2004.

[25] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, “An approach to alleviate link overload as observed on an ip backbone,” in *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, 2003.

[26] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, Y. Ganjali, and C. Diot, “Characterization of failures in an operational ip backbone network,” *IEEE/ACM Trans. Netw.*, vol. 16, pp. 749–762, August 2008.

[27] Y. Wang, H. Wang, A. Mahimkar, R. Alimi, Y. Zhang, L. Qiu, and Y. R. Yang, “R3: resilient routing reconfiguration,” *SIGCOMM Comput. Commun. Rev.*, vol. 40, pp. 291–302, August 2010.

[28] M. Gupta and S. Singh, “Greening of the internet,” *the ACM SIGCOMM*, 2003.

[29] C. Gunaratne, K. Christensen, and B. Nordman, “Managing energy consumption costs in desktop pcs and lan switches with proxying, split tcp connections, and scaling of link speed,” *Int. J. Netw. Manag.*, vol. 15, no. 5, pp. 297–310, 2005.

[30] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows*. Prentice Hall, 1993.

[31] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang, “Experience in measuring backbone traffic variability: models, metrics, measurements and meaning,” in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, ser. IMW ’02. New York, NY, USA: ACM, 2002, pp. 91–92.

[32] A. Gunnar, M. Johansson, and T. Telkamp, “Traffic matrix estimation on a large ip backbone: a comparison on real data,” in *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, ser. IMC ’04. New York, NY, USA: ACM, 2004, pp. 149–160.