Energy-adaptive Network Switching via Intra-device Scaling

Matthew Alexander Hodkin
EDS Lab
Lancaster University
Lancaster, UK
m.hodkin@lancaster.ac.uk

Louise Krug
Research & Network Strategy
BT
Ipswich, UK
louise.krug@bt.com

Peter Garraghan
EDS Lab
Lancaster University
Lancaster, UK
p.garraghan@lancaster.ac.uk

***Abstract*—We propose horizontal intra-device scaling for network switches. Our approach allows for a network device to dynamically scale energy use in response to changing network utilization at a finer-grain in comparison to existing monolithic approaches, and enables a reduction in cost and environmental impact via reduced network energy use outside peak operating periods. We demonstrate the feasibility of intra-device switch scaling by designing a network switch architecture comprising multiple, less powerful, network devices leveraging a Multiple Spanning Tree Protocol (MSTP) to operate in parallel in place of a singular powerful device. Our preliminary results demonstrate that our approach can reduce total network energy use by 66.3% in comparison to established approaches with minimal performance penalty, and outlines future work for further improvement for this new form of network switch architecture for reducing energy use within core network infrastructure.**

***Index Terms*—networking, energy, scalability, data center**

I. INTRODUCTION

The volume of internet traffic across core network infrastructure is growing at an ever increasing rate [1]–[4]. This increase is driven by the expansion in Internet-of-Things (IoT) devices [5], [6], the number of users and peripheral devices, [7] and transition to 5G wireless internet communication [8]–[10]. Increased network traffic has resulted in a net increase of overall energy consumption within the network due to greater computational power and number of switches underpinning core network infrastructure [1]. This growth in network energy consumption represents a significant economic and environmental concern for network users [11], [12].

Within current network architecture, information is transmitted in packets – independently addressed discreet portions of data – from servers to user devices being redirected by a chain of intermediate devices. These devices, known as network switches, are designed to operate at high performance and energy efficiency due to the need for consistent and stable access for industrial and financial applications as well as critical services such as emergency service communications. Current approaches to reduce network energy use consider whole-device techniques such as scheduling to deactivate specific sites and rerouting to reduce power use across the network during periods of low utilization. This generates issues when a device deactivated due to a hardware failure, causing an area to lose network coverage. The currently explored solutions at the individual switch level are sleep modes [13]–[17] and Dynamic Voltage and Frequency Scaling (DVFS) [18]. These solutions are difficult to implement as both solutions increase latency in the event of a surge in network utilization, with any DVFS algorithm recognizing the increase in utilization and increasing energy draw accordingly and sleep modes requiring wake time for the switch as a whole.

In this paper we propose energy-adaptive network switching by leveraging horizontal scaling of intra-device switches: a technique commonly leveraged within data centers [19]. Specifically, we use multiple less computationally powerful switches acting in parallel, instead of using a singular monolithic network switch. The switches acting in parallel utilize a Multiple Spanning Tree Protocol (MSTP) [20], [21] implementation over several Virtual Local Access Network's (VLAN) [22] to balance the load between switches at times of higher network utilization. This parallel setup allows for switches to be deactivated, and their workload redistributed among the remaining active machines whilst providing comparable service for reduced total energy consumption across a given time period as well as allowing for equivalent maximum computational power at times of increased workload for the switches, at a slightly increased maximum power draw.

In this paper we present our design approach and rationale, experiment setup, network layout, scalable topology and preliminary evaluation and results. Our specific contributions comprise the following:

- A new horizontal intra-device scale out topology for network switches to reduce total energy use of the core network infrastructure.
- An experimental evaluation of the proposed intra-device scaleable approach in comparison to a traditional monolithic network deployment.
- Discussion of further work to improve and augment intra-drive network switching.

II. BACKGROUND

Increase in Energy Usage: Total energy consumption of core network infrastructure is increasing [1]. This has lead an increase in scrutiny of switch energy consumption at an individual device level. One key issue is the constant power draw required to operate irrespective of network traffic load.

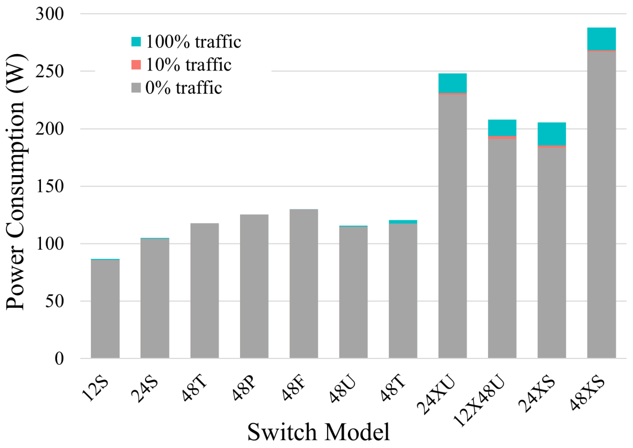

Fig. 1. Data from a collection of Cisco WS-C3850 series switches showing power usage while processing various levels of network traffic.

This characteristics is systemic to the majority of existing switching hardware currently available, as shown within Fig. 1 demonstrating that a collection of Cisco WS-C3850 series switches which use 89.3-99.7% of their total power consumption while idle. This is of significant concern given that networks must be provisioned to account for the maximum volume of traffic expected by the networks, with additionally mandated overhead, resulting in consistent under-utilization and a resultant increase in waste energy.

Hardware Redundancy: Switches currently used in core network infrastructure are used for emergency services communication. Therefore, the network must be capable of maintaining functionality despite hardware failures. Within current hardware, this requires a $2N$ redundancy for each switch. Within a scalable system, only the computer managing the switch would need a duplicate machine as a hot standby in case of failure. Any individual small-scalable-switch failures can be handled by activating one of the small-scalable-switches not currently in use, reducing their required redundancy to as low as $N+1$ and allowing the failed unit to be replaced without deactivation of the switch as a whole.

Implemented Improvements: Currently implemented improvements to networking systems focus on decreasing the total amount of energy the switch draws, making it subject to an indirect rebound effect [23], as noted in Jevons Paradox [24]. This leads to Internet Service Providers acquiring more hardware to increase the capabilities of the service they sell, with current focuses on download speeds, service reliability and consistency in the quality of the provided service.

Theoretical Improvements: While sleep states and DVFS have both been recommended, these two solutions are noncongruent. DVFS focuses on reducing the processing rate to match system demands. This allows computational work to be done constantly at a reduced rate. Sleep states rely on processing all available work as quickly as possible. This is done to allow the switch to re-enter its sleep state and reduce its rate of energy consumption. Both systems require hardware to be designed with these methodologies in mind and have the adverse effect of increasing latency.

III. RELATED WORK

There are several methodologies for attempting on device power scaling, with the current focuses being on DVFS and sleep modes as well as network scale adjustments. Reducing total energy consumption has been a focus in recent years, with each generation of network processor boasting reduced power consumption on its predecessor. Another area where this trend is apparent is in optical networks, where the focus is on reducing the energy consumption of the optical transmitter to reduce both its power consumption and the energy consumption of the required cooling [25].

DVFS and sleep modes have both been an area of focus for many years in attempting to reduce energy consumption of network switches using many algorithms. Sleep modes either require a buffer to hold packets while the switch wakes using wake-on-LAN functionality or must predict when the switch must be awoken using inbuilt models and a controlling computer [14], [26]. DVFS relies on reducing the energy consumption of certain sub-components within the network switch, reducing their processing speed to reduce energy consumption [18]. DVFS required network wide implementation of rate adaption as the surrounding switches must be aware of the reduced packet processing speed of their neighboring switches to ensure they transmit packets at the expected rate for the receiving switch [27].

Another area which has received attention is network scale adjustments, focusing on planning the utilization of the network as a whole, rather than individual switch level improvements. These adjustments take the form of whole site deactivation at times of low traffic and diverting traffic through fewer, more utilized sites while maintaining full network coverage [15]–[17], as required for emergency service communications. While this effectively reduces energy consumption of the network as a whole, it reduces network redundancy as the number of routes between locations is reduced. There is also a limit to the level of granularity in control provided by this implementation as sites which are required to maintain full network coverage must remain active regardless of the volume of network traffic they are processing. The other network wide adjustment has been the implementation of optical bypass [28] which is used to reduce the number of intermediate switches which must used to process a packet on its journey.

IV. DESIGN

The objective of network switching via intra-device scaling is to directly align device energy usage to the volume of network traffic while maintaining network service quality. Networks are a combination of edge switches, which connect directly to devices, and core switches, which link edge switches together and can process more network traffic, as shown in Fig. 2. In this layout, a server and client connected to the same edge switch will communicate by sending packets via their shared edge switch without utilizing the core switch.

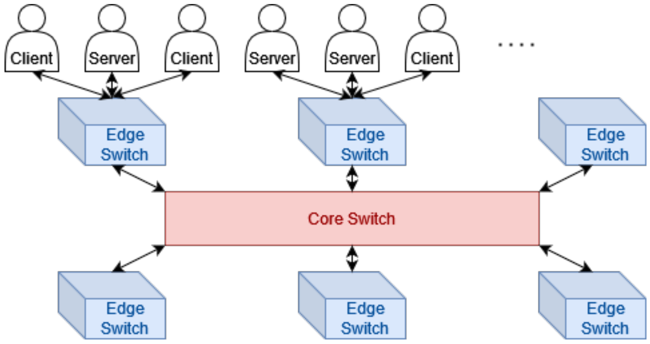

Fig. 2. A traditional monolithic switch setup.

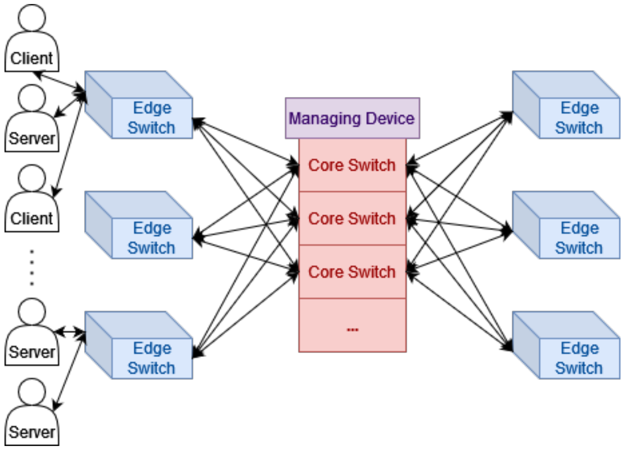

Fig. 3. The proposed intra-device scaling switch setup.

TABLE I
SUMMARY OF NOTATIONS IN DESIGN.

Notation	Description
n_t	Total number of switches
n_a	Total number of active switches
t	Volume of network traffic
$caps$	Capacity of individual switch
cap_a	Total capacity of all currently active switches acting in parallel
cap_t	Total capacity of all switches acting in parallel
i	Interval between iPerf calls
r	Time between recalculations of MSTP routes
u	Threshold for increasing the number of active switches
l	Threshold for increasing the number of active switches
$systime$	System clock time

Algorithm 1 Algorithm of proposed strategy.

```
$cap_t \leftarrow cap_s * n_t$
while system is active do
    if $systime \ \% \ i = 0$ then
        $cap_a \leftarrow cap_s * n_a$
        Calculate $u$ and $l$ using $cap_a$ and $cap_t$
        if $t > u$ then
            $n_a \leftarrow n_a + 1$
            $cap_a \leftarrow cap_a + cap_s$
        else if $t < l$ then
            $n_a \leftarrow n_a - 1$
            $cap_a \leftarrow cap_a - cap_s$
        end if
    end if
end while
```

If a server and client are connected on different edge switches, then the server would send packets to their local edge switch, which would forward them to the core switch. The core switch would then forward the packets to the clients local edge switch, which would relay the packets to the client machine.

Intra-device Scaling entails replacing the large singular core switch with a number of less powerful switches, as shown in Fig. 3. These switches can process less network traffic individually, and can work in parallel, with an equivalent maximum volume of network traffic to that of a traditional monolithic design. We manage the activation and deactivation of these switches via a managing device. This device monitors traffic flow and determines the number of switches which must be active as a result. The deactivation of switches, when they are not required to process the volume of network traffic, leads to a reduction in total energy usage for the system as a whole and enables energy scalability.

Within the intra-device scaling layout, the total capacity of the network core cap_t is found by multiplying the network traffic capacity of an individual switch cap_s and the total number of switches that can act in parallel n_t. The number of active switches n_a is determined by comparing the volume of network traffic t to the threshold for increasing the number of active switches u and the threshold for decreasing the number of active switches l. The capacity of all active switches in the core cap_a is found by multiplying the network traffic capacity of an individual switch cap_s and the number of active switches that can act in parallel n_a.

The strategy for linking energy use to volume of network traffic is to take the workload normally processed by a single, more powerful switch and distributing it between several smaller switches. These smaller switches will act in parallel with the number of switches active proportional to the volume of network traffic at that time, as shown in Algorithm 1 using the variables in Table I.

The energy utilized is tied to the volume of network traffic t by increasing or decreasing the number of active switches n_a dependent on t. The setup is controlled by a managing device capable of activating and deactivating switches while the flow of traffic and balancing between all n_a switches is carried out by an MSTP implementation refreshing every r seconds. This is conducted to reduce switch under-utilization or over-utilization, leading to a reduction in the waste energy by providing different routes for network traffic to take. The strategy proposed operates with a discreet number of energy levels, that being the number of switches active and working in parallel and an additional energy level when all switches inactive. This means that with a layout of n_t switches, there are $n_t + 1$ discreet energy levels. In this topology, a greater number of switches, each with a lower total energy use, would allow for more discreet energy levels and therefore a closer correlation between the network utilization and the energy utilization, allowing for greater efficiency of the system as a whole while under use.

V. EXPERIMENT

A. Setup

Within our experiments, we evaluated the energy and performance of both the traditional monolithic architecture (Fig. 2) and a horizontal intra-device switching architecture (Fig. 3). The monolithic architecture captures the model of a typical network, with several edge switches feeding edge devices connected to each other through a core switch, which can process a higher volume of traffic. However, this core switch is not directly connected to any edge devices. Intra-device switching instead leverages multiple smaller switches whose activation and deactivation is controlled by a managing machine, which makes determinations on resource provisioning, based on the volume of network traffic.

The network is setup with the edge composed of a series of TP-Link JetStream 8 Port Switch TL-SG2210MP's [29]. In our experimental setup, 7 of these switches are used, each connected to 3 machines acting as server/client machines with a 1Gbps connection rate. In the monolithic setup, the switch at the core of the testbed network is a single TP-Link JetStream 24 Port Switch TL-SG3428 [30] connected to each edge switch using 3 ports acting in unison utilizing Link Aggregation [31] to ensure both setups exhibit a connection rate of 3Gbps between the switch at the core of the testbed network and edge switches. In the intra-device approach, switches at the core of the testbed network are 3 TP-Link JetStream 8 Port Switch TL-SG3210 V3's [32], each connected to an edge switch using a single 1Gbps connection, totaling 3Gbps between all 3 switches when all are operational.

These smaller switches are designed to act in parallel using a combination of MSTP and VLAN to allow for switches to be disconnected and rejoin the setup as required, as well as to ensure that traffic is balanced across the switches to avoid any switches being over-utilized while others remain idle. The time between recalculations of MSTP routes in the network r is set to 20 seconds. VLAN assignment was split evenly across each server/client machine connected to any given edge switch to ensure load balancing between all switches at the core of the testbed network when acting in parallel under the scalale setup. Switch activation and deactivation is controlled by a PowerBOX 4KG [33], which is an Ethernet controlled 4 UK socket power supply, allowing switches to be automatically activated and deactivated as the system requires.

All machines send and receive data using the iPerf command, with i set to 15 seconds , allowing for control of the volume of network traffic as well as access to the required measurement metrics of Jitter and Number of Packets lost. Both network setups receive two different profiles of network traffic, one provided by BT based on real world utilization data, as shown by the line in Fig. 4(a), and one as a sine curve of utilization over the experiment run, as shown by the line in Fig. 5(a). We then measure the (1) energy consumed by the device (W), (2) Network Jitter (ms), and (3) packet loss under these traffic loads.

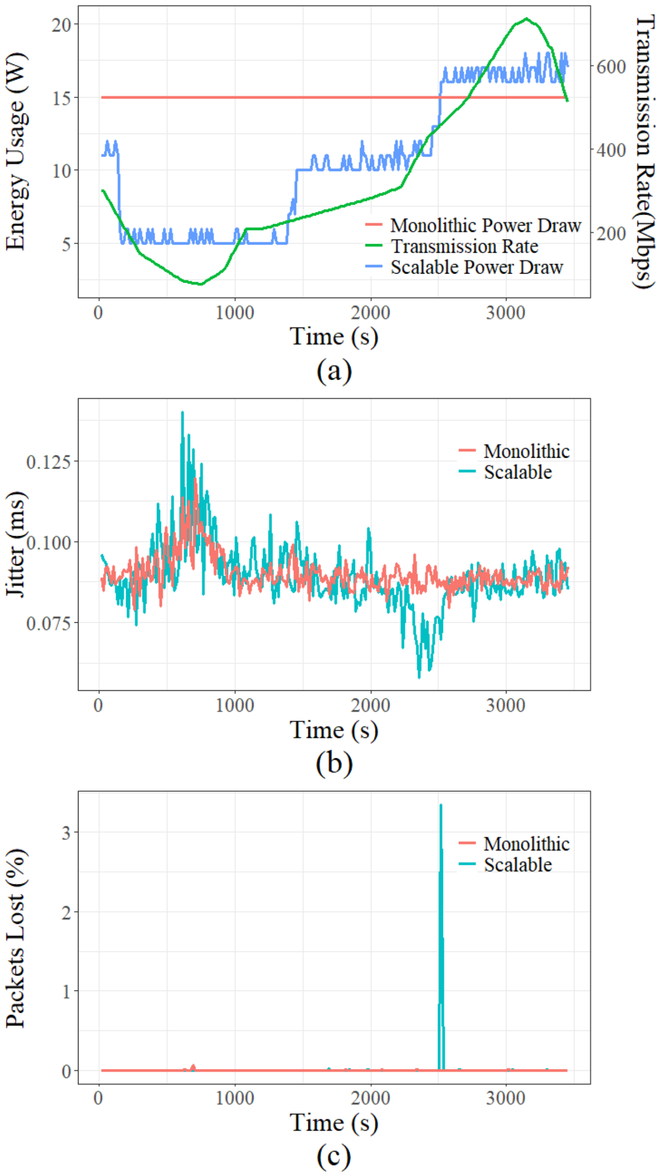

(a)

(b)

(c)

Fig. 4. (a) Average energy consumption of both setups against the network traffic passing through the setup using BT traffic profile, (b) Average Network Jitter of both setups over the hour run using BT traffic profile, (c) Average Packet loss of both setups over the hour run using BT traffic profile.

B. Preliminary Results

Experiment results are depicted in Fig. 4 & Fig. 5, with all results averaged across 10 runs. In Fig. 4(a) and Fig. 5(a), noting the Monolithic Power Draw, we observed that the energy consumption of the monolithic setup is consistently 14-15W, regardless of switch utilization, following the trends seen in Fig. 1 in the CISCO WS-C3850 series switches. In Fig. 4(a) and Fig. 5(a), noting the Scalable Power Draw, it is observable that the scalable setup operates at 3 levels of energy consumption, one for each switch activated, and

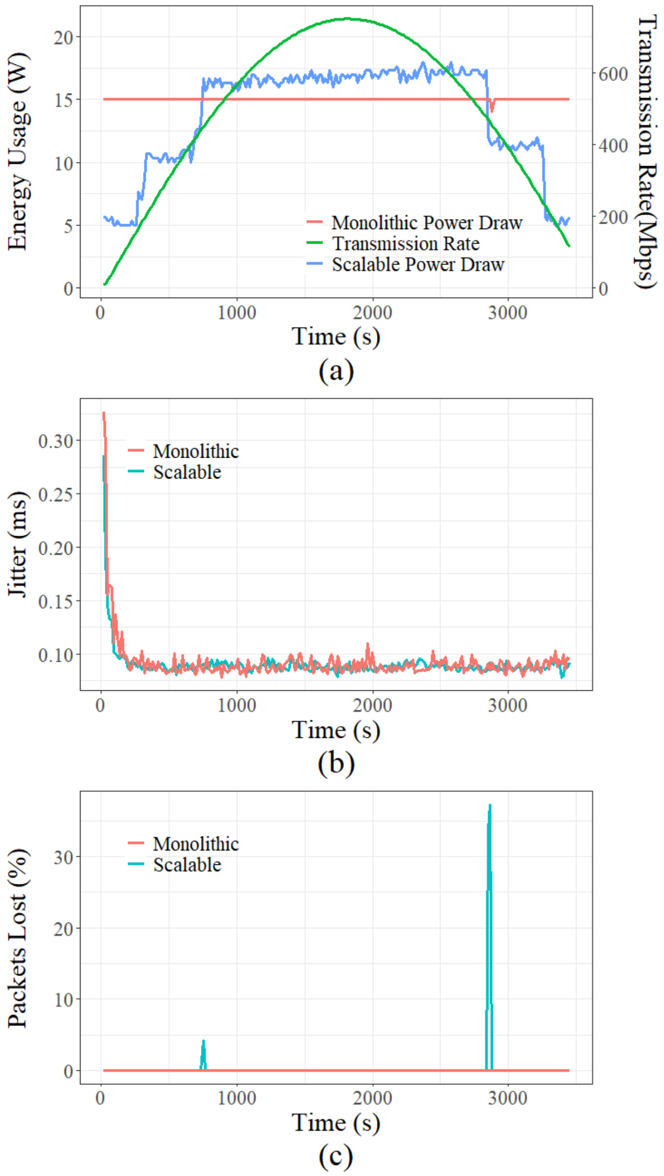

Fig. 5. (a) Average energy consumption of both setups against the network traffic passing through the setup using Sine traffic profile, (b) Average Network Jitter of both setups over the hour run using Sine traffic profile, (c) Average Packet loss of both setups over the hour run using Sine traffic profile.

TABLE II
RESULTS OF BOTH EXPERIMENT NETWORK SETUPS USING BT TRAFFIC PROFILE (3 S.F.).

BT Traffic Profile	**Scalable Setup**	**Monolithic Setup**
Total Energy Usage (J)	35500	51800
Average Energy Usage (W)	10.3	15
Average Jitter (ms)	0.0900	0.0901
Range of Jitter (ms)	0.0820	0.0406
Maximum packet loss (%)	3.35	0.0741

TABLE III
RESULTS OF BOTH EXPERIMENT NETWORK SETUPS USING SINE TRAFFIC PROFILE (3 S.F.).

Sine Traffic Profile	**Scalable Setup**	**Monolithic Setup**
Total Energy Usage (J)	47500	51700
Average Energy Usage (W)	13.8	15
Average Jitter (ms)	0.0903	0.0922
Range of Jitter (ms)	0.208	0.248
Maximum packet loss (%)	37.2	0.0147

more closely follows the network traffic and device utilization. This shows that the experimental setup is functioning as intended in tying energy consumption more closely to the packet routing work done by the switches. We observed that intra-device scaling enabled a minimum and maximum power draw of 5W and 17W, respectively, and an overall 35500J in comparison to 51800J (a reduction to 66.3%) in comparison to monolithic. Fig. 4(b) and Fig. 5(b) show the jitter across the same experimental runs, with the difference in average jitter in the BT and Sine traffic profiles being 0.0001ms and 0.0018ms, respectively. However, as seen in Fig. 4(c) and Fig. 5(c), we can see the packet loss rates can spike at the points when a new switch is added or removed from the setup. This is a result of desynchronisation between the MSTP implementation recalculating the layout of the network and the managing computer activating or deactivating switches. In the scalable setup there is a 2 minute delay between a new switch being activated and the switch being available to the network to process packets. During this initialization period the new switch draws 2W of power.

VI. CONCLUSIONS

In this paper we have proposed a new form of energy-adaptive network switching via leveraging intra-device scaling based on network workload. Our preliminary results indicate that our approach is is capable of reducing total energy use to 66.3%, against a network traffic flow modeled on real world data, in comparison to conventional approaches, whilst maintaining similar performance levels. Our work indicates that the horizontal intra-device scale out topology is effective in reducing total energy consumption of a network, without increasing Jitter within the network, with an increase in the number of packets lost.

A. Future Work

While initial experiment results are promising and highlight the potential value of our approach, there are various aspects to further study and refine to improve the approach. Specifically, the current approach resulted in an increased packet loss when a switch enters and leaves the proposed setup. This is, as stated, an issue in timing disconnect between the recalculation of the network shape in the MSTP implementation and the computer managing switch activation. This issue causes packet loss as a switch is only removed from the MSTP network view when a recalculation is performed after a switch has already been deactivated, leading to the loss of any packets intended to be routed via that switch. To solve this issue we intend to use a P4 [34] implementation. This implementation would be managed directly from the controlling computer and would be used to remove a switch logically from the network, allowing it to clear its buffer of packets currently being processed, before deactivating the switch and allowing for the removal of the MSTP and LAN implementation currently in use.

Another possible improvement is to develop a model which can predict future traffic demands to allow for additional small scalable switches to be activated and added to the setup before the increase in network traffic occurs in order to avoid a significant increase in latency, jitter and packet loss. A prediction model would also allow for a higher degree of certainty that a small scalable switch can be removed from the system. An alternative to this would be the development of a new packet, either propagated down a specific path between server and user or throughout the network to warn of an immediate increase in traffic through the informed switches, allowing them to increase their capacity as required and allowing them to react to sudden increases in demand. Under a redesigned switch hardware architecture, sleep modes could be integrated allowing for near-immediate response to increases in traffic. This would be achieved by setting the next switch required in sleep mode rather than fully deactivated.

ACKNOWLEDGMENT

The authors would like to thank BT for their funding and the network traffic data supplied for this project. This work was supported by the EPSRC (EP/V007092/1).

REFERENCES

[1] J. Morley, K. Widdicks, and M. Hazas, "Digitalisation, energy and data demand: The impact of internet traffic on overall and peak electricity consumption," *Energy Research & Social Science*, vol. 38, pp. 128–137, 2018.

[2] J. Baliga, K. Hinton, and R. S. Tucker, "Energy consumption of the internet," *COIN-ACOFT 2007-Joint International Conference on the Optical Internet and the 32nd Australian Conference on Optical Fibre Technology*, pp. 1–3, 2007.

[3] G. Fettweis and E. Zimmermann, "ICT energy consumption-trends and challenges," *Proceedings of the 11th international symposium on wireless personal multimedia communications*, vol. 2, no. 4, 2008.

[4] R. S. Tucker, J. Baliga, R. Ayre, K. Hinton, and W. V. Sorin, "Energy consumption in IP networks," *2008 34th European Conference on Optical Communication*, 2008.

[5] S. Al-Sarawi, M. Anbar, R. Abdullah, and A. B. Al Hawari, "Internet of things market analysis forecasts, 2020–2030," *2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*, 2020.

[6] Y. L. Cheng, M. H. Lim, and K. H. Hui, "Impact of internet of things paradigm towards energy consumption prediction: A systematic literature review," *Sustainable Cities and Society*, vol. 78, p. 103624, 2022.

[7] Cisco, "Cisco annual internet report (2018–2023) white paper." [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html

[8] E. Hajlaoui, A. Zaier, A. Khlifi, J. Ghodhbane, M. B. Hamed, and L. Sbita, "4G and 5G technologies: A comparative study," *2020 5th International Conference on Advanced Technologies for Signal and Image Processing (ATSIP)*, 2020.

[9] X. Cheng, Y. Hu, and L. Varga, "5G network deployment and the associated energy consumption in the UK: A complex systems' exploration," *Technological Forecasting and Social Change*, vol. 180, p. 121672, 2022.

[10] L. Williams, B. K. Sovacool, and T. J. Foxon, "The energy use implications of 5G: Reviewing whole network operational energy, embodied energy, and indirect effects," *Renewable and Sustainable Energy Reviews*, vol. 157, p. 112033, 2022.

[11] Z. Allam, S. E. Bibri, and S. A. Sharpe, "The rising impacts of the COVID-19 pandemic and the Russia–Ukraine war: energy transition, climate justice, global inequality, and supply chain disruption," *Resources*, vol. 11, no. 11, p. 99, 2022.

[12] B. Francis-Devine, P. Bolton, M. Keep, and D. Harari, "Rising cost of living in the UK," *Research Briefing, House of Commons Library, October*, 2022.

[13] M. Gupta and S. Singh, "Greening of the internet," *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 19–26, 2003.

[14] D. Ali, F. R. Raja, and M. A. Saleem, "Energy efficient resource utilization : architecture for enterprise network : towards reliability with sleepalert," 2020.

[15] Y. Wei, X. Zhang, L. Xie, and S. Leng, "Energy-aware traffic engineering in hybrid SDN/IP backbone networks," *Journal of Communications and Networks*, vol. 18, no. 4, pp. 559–566, 2016.

[16] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "Elastictree: saving energy in data center networks," *NSDI*, 2010.

[17] L. Zhou, L. N. Bhuyan, and K. K. Ramakrishnan, "Dream: Distributed energy-aware traffic management for data center networks," *e-Energy*, 2019.

[18] S. Nedevschi, S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing network energy consumption via sleeping and rate-adaptation," *Symposium on Networked Systems Design and Implementation*, 2008.

[19] L. A. Barroso, J. Clidaras, and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, 2013.

[20] A. de Sousa and G. Soares, "Improving load balance and minimizing service disruption on ethernet networks with IEEE 802.1 s mstp," *Workshop on IP QoS and traffic control*, 2007.

[21] N. Finn, "Multiple spanning trees in 802.1 Q," *IEEE*, vol. 802, pp. 1–11, 1996.

[22] V. Rajaravivarma, "Virtual local area network technology and applications," *Proceedings The Twenty-Ninth Southeastern Symposium on System Theory*, 1997.

[23] S. Sorrell, "Jevons' paradox revisited: The evidence for backfire from improved energy efficiency," *Energy policy*, vol. 37, no. 4, pp. 1456–1469, 2009.

[24] W. S. Jevons, "The coal question; an inquiry concerning the progress of the nation and the probable exhaustion of our coal-mines," *Macmillan*, 1866.

[25] Nokia IP Networks, "Improving the environmental sustainability of IP networks," 2023. [Online]. Available: https://onestore.nokia.com/asset/210560?_ga=2.84354141.552477406.1690896198-333856616.1690896198. [Accessed: Sept. 23, 2023]

[26] J. Chen and B. Sikdar, "A mechanism for load proportional energy use in wireless local area networks," *2013 IEEE Global Communications Conference (GLOBECOM)*, pp. 2466–2471, 2013.

[27] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 rate adaptation: a practical approach," *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, pp. 126–134, 2004.

[28] J. M. Simmons, "Optical network design and planning," *Springer*, 2014.

[29] TP-Link, "Jetstream 8 port switch TL-SG2210MP," 2020. [Online]. Available: https://www.tp-link.com/uk/business-networking/omada-sdn-switch/tl-sg2210mp/. [Accessed: Oct. 8, 2023]

[30] TP-Link, "Jetstream 24 port switch TL-SG3428," 2021. [Online]. Available: https://www.tp-link.com/uk/business-networking/omada-sdn-switch/tl-sg3428/. [Accessed: Oct. 8, 2023]

[31] A. Hendel, "Link aggregation trunking," *Sun Microsystems, IEEE*, 1997. [Online]. Available: http://www.ieee802.org/3/trunk study/tutorial/ahtrunk.pdf

[32] TP-Link, "Jetstream 8 port switch TL-SG3210 V3," 2020. [Online]. Available: https://www.tp-link.com/uk/business-networking/managed-switch/tl-sg3210/. [Accessed: Oct. 8, 2023]

[33] Netio, "PowerBOX 4KG," 2021. [Online]. Available: https://www.netio-products.com/en/device/powerbox-4kx. [Accessed: Oct. 8, 2023]

[34] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese *et al.*, "P4: Programming protocol-independent packet processors," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 87–95, 2014.