


Towards a More Energy Efficient Future Telecoms Network

Energy efficiency is increasingly important in telecommunications networks. This paper outlines an objective to improve efficiency by dynamically matching network power consumption to active traffic load, moving away from today's typically constant power consumption. We seek to work with vendors and wider industry to help develop these ideas and drive more energy efficient overall network solutions.

Introduction

Cost of energy and impact on the environment have become central to telecoms thinking. BT Group consumes nearly 1% of the UK's electricity¹. Much of this energy usage comes from legacy PSTN/TDM telecommunications technology which has seen little change for 30 years. In common with many operators these old networks will soon be switched off. This leaves the modern telecoms network based on FTTP, 4/5G and IP technologies. The future focus of telecoms energy consumption will be to ensure these modern networks make efficient use of energy and are able to scale to future demand whilst reducing overall power consumption.

Today most technologies make constant use of power independent of traffic load. When peak capacity growth demands more infrastructure, power consumption is consequentially increased at all times and not just at the peak load times. We believe that future telecoms technology must seek to improve on this, i.e. enable power consumption to track traffic load, enabling reduced consumption during periods of lower usage. Achieving this is expected to involve all layers of the network working in cooperation – software, hardware, orchestration and automation. The purpose of this paper is to foster collaboration across industry, including our vendors and other network operators to address this challenge.



Capacity Deployment

Providing sufficient capacity to ensure good customer experience during peak events, with sufficient resilient capacity to provide for protection of network failures, means the network must be sized to significantly exceed these extreme peaks.

The logistics associated with deploying network capacity often mean that new capacity is deployed many months in advance of the forecast demand. It often takes multiple years to upgrade capacity across 100s-1000s of sites. This means that new capacity, both switching equipment with spare ports waiting for new customers and inter-site transmission, are deployed long before the demand is in place to use them.

All of this equipment is today commissioned and using close to peak power usage from the initial date it is installed. This leads to an energy demand for our network which shows little time-of-day variation, and which is increasing over time as we build additional capacity for customers.

Traffic Profiles

In common with all telecommunications networks, we see a strong time-of-day variation in the traffic profiles of our networks. This varies by class of traffic, but the overall load on the backbone network shows a common trend. Figure 1 shows this backbone traffic in terabits per second (Tbps) for a representative week on BT's network. It shows clear daily variation from Monday through Sunday.

Traffic drops to a distinct trough overnight, around 20% of the peak. It rises during the day and reaches an evening peak driven predominantly by fixed broadband. Exceptional events such as live sport TV streaming drive irregular peaks which significantly exceed the normal peak. (The dotted line shows a recent irregular peak week which exceeds the normal peak by 25%).




Figure 1: Typical weekly traffic profile on BT's backbone network

¹ <https://newsroom.bt.com/how-bt-group-is-making-our-networks-more-energy-efficient>

Overall Opportunity

The opportunity for operators is to seek to provision a network which has a minimum level of power draw, and which varies consumption above this dynamically based on usage (Figure 2). It may meet exceptional demand with surge capacity which is only used during these events. It should also enable all capacity which is deployed but not yet needed to be held in low power mode until required. Network capacity operating in this mode should still be manageable and ideally will automatically transition to operational state where the network device has sufficient information to do so, such as when additional services are activated.




Figure 2: Target power consumption varies with demand and operational status

Considering the core traffic profile shown in Figure 1, the area under the traffic graph is around 55% of constant traffic at the weekly peak and 45% of constant traffic at the irregular peak. This indicates there could be an opportunity of up to approximately 50% reduction in power utilisation if this could be linked closely to traffic utilisation rather than peak capacity.

We see opportunity to develop power saving capabilities in two dimensions to help achieve this (Figure 3):

1. Vary power with traffic usage within individual active resources. For example, a switching ASIC or CPU may have a minimum power consumption and may vary the consumption above this with dynamic traffic load. Ideally the minimum is small relative to the variable component. A variant may be to turn down the capacity of an active resource, such as the capacity over a transmission link, to use lower power by reducing the capacity/speed during periods of lower demand.
2. Vary the number of active resources on a granular basis. For example, by being able to disable ports, ASICs, processors or transmission links, based on level of demand. This may be driven by a simple policy to disable capacity which has been installed but is not yet needed or may be driven by more complex time-of-day or load-related control. This may use automation to manage active resources in the network via a centralised control function. This is likely to be associated with a mechanism to routinely test the operation of components in disabled mode to verify operation.




Figure 3: Two dimensions of power saving

Together we expect individual resource optimisation and power aware orchestration which control active devices to enable significantly more variable power consumption for the whole network.

A further opportunity may be to vary the amount of active capacity provided for restoration purposes (for use on failure of active capacity), such that the amount of restoration capacity which is maintained in an active state is varied based on the current load and operational state of the network. By doing so it is expected that trade-offs can be made between the speed of restoration on failure and the amount of capacity which must be pre-provisioned and active to provide protection capacity.

It is conceivable that operators will be willing to balance speed of protection switching against amount of active protection capacity, which is consuming power but not carrying any traffic. This could lead to new protection regimes which do not maintain all spare capacity in a permanently active state in order to save power, at the expense of extending time to recover beyond the traditional telecoms 50ms target.

The remainder of this paper considers different classes of network equipment where we see varying levels of opportunity based on the existing maturity of these concepts.

Core Software Platforms

Much of the future telecoms network will be based on commercial off-the-shelf server technology running telco specific software applications. We see many classes of this already with varying levels of applications from mobile packet cores, IMS and customer specific NFV.

We see power saving capabilities of these platforms broadly categorised in the two dimensions identified above. Category 1 focuses on CPU power management (use of c-states, p-states, uncore frequency scaling) as well as considering power policies of other server components, for example fan speeds. Category 2 involves powering off entire servers and we predict that this will save an increased amount of power as servers can consume up to 60% of their peak power when idle.

The main downside to these techniques is the latency in returning to an active state. Category 1 solutions will generally have a smaller exit latency, for example c-states are on the order of microseconds, meaning these solutions could be considered 'less risky' when compared to category 2, where the exit latency is the delay in powering on and provisioning servers, on the order of minutes. Category 2 will likely require more in-depth investigations and trialling, especially with the complex nature of orchestration required. Work is required to determine the optimal use of both, and we encourage vendor solutions to address this agnostic to the running application and considering different virtualisation layers and server hardware.

Where the application is predominantly performing control plane functions, slightly more aggressive policies may be appropriate, when compared to user plane dedicated servers, given the application may not be realtime and may not have a direct impact on customers. It is important to understand the level of resource required to meet surges in control plane activity and ensure that this is not compromised by power saving measures.

Where the application is predominantly a user plane or packet forwarding application, the use of power management solutions are complicated due to the use of poll mode drivers such as DPDK for fast packet processing. Workloads that make use of this typically appear as 100% utilised at the OS level meaning category 1 low power states will never be entered. True utilisation telemetry should be used to ensure solutions optimally use category 1 power management and category 2 capacity management, whilst also being mindful of the impact on customer experience.

Software platforms offer the most opportunity for rapid innovation and are also the least efficient in terms of power consumed per unit of capacity processed today. We see active development in this domain and encourage vendors to continue to develop both the underlying optimisation capability and orchestration with us. We also encourage vendors to develop software compatible with a range of CPU architectures e.g. Intel, AMD and ARM for optimal performance per watt in software platforms.

Core Routing Platforms

The backbone of the telecoms network will continue to deploy hardware switching to provide the IP layer. We see three levels of power management to be important:

- i. Ability to deploy hardware as part of a capacity planning build ahead process, but not enable power consumption on a granular level until there is demand. This means being able to place optics, ports, cards and internal capacity such as ASIC, SerDes etc into low power or off mode until they are required. Equipment should be able to be deployed with a live management plane and powered up on a granular basis as customers are connected or core links commissioned.
- ii. Ability for the forwarding hardware to vary power consumption with traffic load. Ideally ASIC based functions would provide similar capability to CPU based forwarding planes where the power consumption of the ASIC varies with level of traffic without any intervention. Historically we have seen very little difference in power consumption of device with zero traffic and full traffic. There is scope for significant improvements in this domain.
- iii. Ability to dynamically power up and down components with traffic load on a granular basis to provide the equivalent of scale out and scale back dynamics for physical devices. This may be internal to the device, reducing the number of internal switching components which are active with traffic load, or may extend network wide to disabling ports and physical links during periods of low traffic to allow the network to reroute onto the remaining links. Some component optimisation may be controlled by internal automation within the device, but there may be more opportunity for network wide optimisation with external SDN control of a number of devices.

We are seeing positive early developments in this domain but encourage vendors to develop this further and consider how cross device orchestration can enable more power saving opportunities.

Optical Platforms

Optical WDM networks comprise optical fibre links and terminal devices, which provide connectivity with required network capacities. Coherent optical transponders/line optics include complex technologies and require higher power consumption from lasers/receivers and DSP ASICs.

We see scope for three levels of energy optimisation. These apply to both dedicated optical equipment and pluggable optical modules used in routing equipment.

- i. Ability to shut down individual client and line ports in optical transponder cards based on their usage. There are usually multiple client ports and more than one line port per card. Ability to shut down individual ports based on their usage would allow network operators to reduce the power consumption as well as adapt to changes in traffic demands.
- ii. Ability to turn down the speed of individual optical wavelengths based on demand. Modern coherent optics often offer scaling of the wavelength channel speed such as between 100 and 400G based on trading off rate and reach. It may be possible to turn down the wavelength capacity over an optical link until there is more traffic demand to use the full rate. The potential benefit would be to reduce power consumption, for example, by using less DSP resources. It is also worth exploring other power savings within coherent optics, e.g., through reduced dispersion compensation DSP load for different optical routes/reach.
- iii. Ability to turn on and off individual wavelengths based on network capacity demand. We see opportunity to scale resources by disabling some wavelength capacity at periods of low traffic or before the capacity is required. It may be viable to perform this on an occasional basis based on peak capacity demand or on a regular basis based on time-of-day traffic levels. Understanding the level of power savings and whether some or all components of a coherent transponder and optical pluggable (e.g., laser, DSP etc) can be disabled and reenabled on a regular basis is a significant area of interest to understand what level of power efficiency may be achieved long term.

There have been continuing efforts in new technology developments to reduce power consumption of optical components as well as DSPs, which mainly improve static power consumption. We have yet to see significant developments to enable dynamic power savings in this domain and encourage vendors to review what options may be viable to provide dynamic power to match the customer load on the network.

Radio Access Networks

The power consumption of RAN is a significant contributor to an operator's overall energy usage. Operators have been exploring innovative approaches to minimise this energy footprint. We are already implementing power-saving modes in our radio networks. These modes are designed to reduce power consumption during periods of reduced demand, typically at nighttime.

During these periods, RAN equipment is able to disable specific radio carriers or cells, effectively reducing network capacity. However, this reduction does not impact connectivity for customers. Mobile devices automatically select the remaining carriers for connection, ensuring uninterrupted service while lowering energy consumption.

We see cutting-edge technologies to further enhance energy efficiency. For instance, the utilisation of the latest ultra-lightweight Massive MIMO radio technology supplied by various vendors to improve 5G network performance provides up to a 40% reduction in energy consumption compared to previous generations.

Additionally, collaborations with technology providers are already resulting in development of advanced equipment and software features, such as ability to put components into 'deep sleep' mode with the potential to save up to 70% of power per radio during low-traffic hours.

From an operator's perspective, RAN vendors must prioritise key factors to enhance energy efficiency. These include maximising data transfer efficiency and minimising power consumption during idle periods through innovative hardware and software solutions. Additionally, the integration of AI/ML technologies can enable rapid, real-time decision-making, considering site configurations and traffic conditions, leading to optimisation of RAN energy consumption.

Radio access is one of the most mature domains in seeking to apply dynamic control to power usage, in part due to the high volume of power consumed. We encourage vendors to continue to innovate and support us in minimising energy usage without compromising network quality.

Fixed Access Networks

While individually each piece of equipment and component in the access network typically use less power than those in the core, the volume of equipment required for the access network means that the total power consumption of the fixed access network equipment is significantly higher than the core network.

As with the IP core equipment the components in the access fibre network typically consume similar power whether loaded or unloaded. This means the access network has a large potential for savings if power consumption can be reduced in line with actual traffic rather than the capacity of links. Given the relatively low numbers of users served by each piece of equipment in the access network compared to the core, operators may be willing to accept more risk as a result of power reduction techniques in the access network relative to the core.

We see 3 opportunities to reduce the power consumption:

- Improving the raw peak rate efficiency of the equipment (Watts per Gb/s). We see DSP and compute developments where the move to smaller nodes sizes has reduced the Watts per Gb/s over recent years, but this has not always been applied to access devices. By including power efficiency higher up in the design criteria we see opportunity to improve Watts per Gb/s of access equipment.
- Scaling the power consumption with port utilisation. Typical vendors produce a limited number of products to minimise the hardware design costs. Operators then use these products with very high maximum throughputs in a range of different deployments in many cases with less than 100% ports or port capacity enabled. This includes not all ports active, ports running at lower rates e.g. 2.5Gb/s PON in a 10Gb/s capable port, total end user services rates below capacity or backhaul capacity below the access capacity. In these cases, power reducing modes can be used that are semi-static, i.e. that are adjusted over a period of years, possibly during planned engineering downtime. These may be much simpler to implement than power saving features that respond dynamically to traffic.
- Scaling the power consumption with actual traffic. There may also be opportunity for saving power by responding more actively to daily periods of low traffic demand outside of the busy hours.

Network Wide – Environment and Cooling

A further consideration is the impact of the physical environment on power consumption of telecoms devices. Most modern equipment has variable cooling fans which optimise the power consumed by cooling to ensure the device runs within operational limits. We see opportunity for system wide optimisation of power and cooling including both the accommodation and electronics. This optimisation should be aware of how equipment operating temperatures will affect reliability. We would like to work with vendors for greater access to data to support this wider optimisation. As networks become more dynamic to vary their power consumption with traffic, we may also expect to carry out these optimisations in real-time.

Introduction of alternative cooling solutions such as liquid/water cooling and alternative form factor devices may also play an important role in future telecommunication networks. We are open to alternative suggestions – for example a consideration about how we can use “free” cooling and understanding the relationship between space occupied and the cooling requirements. We can see that whilst our old equipment is power hungry, it had significantly reduced cooling requirements compared to modern systems as the heat density was lower.

Another possibility could be high temperature operation, removing the need for any external cooling systems. This is also of interest as it is not unlikely that we could exceed 45 degrees external temperature in the medium term, perhaps during the lifetime of the next generation equipment.

Network Wide – Power Specification, Instrumentation and Standardisation

We see scope for significant additional focus from vendors on publishing power consumption expectations and providing telemetry to measure real time power usage. Historically vendor datasheets have provided maximum consumption for power planning purposes, but far less information related to typical consumption at a range of traffic load and ambient temperatures. As the industry becomes more focussed on power consumption, this level of measurement and specification will become increasingly important to operators.

Greater standardisation of how power consumption is tested and published is also encouraged, together with common approaches to providing real time power consumption telemetry. We also encourage standardisation of interfaces for controlling the power state of network components to allow SDN controllers to monitor and provide closed loop control of the active power status of the network.

Conclusions

We have outlined the importance of power efficiency for future telecommunications networks and identified that there is considerable opportunity to improve on existing primarily static power consumption driven by number of installed devices. We believe the future telecoms network should offer power consumption which is driven by the number of active customers and level of active traffic and so should show considerable variation by time of day and customers connected.

We expect future energy efficiency to come from a combination of energy efficient components which can be turned up and down and a level of SDN orchestration and control which is able to manage which components are live at any time. There will be important trade-offs for operators to consider in order to optimise energy efficiency whilst ensuring a level of active resilient capacity is maintained for different services across different times of day.

We summarise the level of maturity of each class of network equipment in Figure 4 below. We believe the industry is at the early stages of maturity (left) for most domains with radio access showing most progress, but no domain yet reaching the full maturity (right) we expect to be important for the long term.

We would like to work with industry, our vendors and other operators to explore these concepts further.

	Always On	Granular Control Static [planning]	Granular Control by time of day	Variable with Load for active components	Dynamic Scaling of active / inactive components	SDN controlled Optimisation incl Resilience
Software Platforms		●		○ Emerging	○	
Core Routing		●				
Optical	●					
RAN			●	○ Emerging		
Fixed Access	●					

Figure 4: Maturity of dynamic power saving per network domain

About this paper

Written by members of BT's Tech Fellowship cross company community of technical leaders as part of an initiative to foster collaboration across industry on improving energy efficiency.

Authors: Muhammad A. Goraya, Abigail Griffiths, Ian Horsley, Louise Krug, Adrian Smith, Yu-Rong Zhou.

Contact: adrian.ca.smith@bt.com (editor)

Annex

We expect power efficiency and dynamic power consumption with traffic load to become increasingly important to our network and become a key part of future technology selection.

We have included below example questions vendors can expect to see in our future procurement processes.

Energy Use

- What options are available for disabling ports/cards/components to enable power to be saved when no customers are connected to them?
- Can ports/cards/components be enabled and disabled dynamically under management system or local device control.
- What is the energy per bit at optimal usage?
- What is the power consumption at peak usage and when idle (or a graph of traffic power)?
- How does this vary with external temperature?
- What is the capability for reporting energy use and how granular is this by port/card/component?
- What are the options available for dynamically reducing energy use and what trade-offs do these options make?
- Can energy efficiency be improved by providing network wide control of active resources and how do your resources participate in such an ecosystem?
- What is the expected MTBF (reliability) of the device and how does that change if the device or portions of the device are depowered daily?

Circular Economy

- Do you have an estimate of the energy used to create the device?
- What method is used to estimate this?
- Do you take back devices for re-use or recycling at end of life?
- What percentage of returned devices are reused?
- When recycling what percentage of components can be reused?
- When recycling what percentage of materials do you recover?
- What materials on the EU critical risk list are used in the device?
- What percentage of these materials do you recover?

About the authors

BT's Tech Fellowship is a career development programme run within BT's networks business unit which brings together those who want to further their development as technical specialists, keep abreast of the latest technology and make a significant impact in an ever-changing industry. The Fellowship offers dedicated time away from the day job and helps build collaboration across the BT technical community, providing new opportunities and access to developing technologies. This article has been written by one of the collaboration squads established as part of the Fellowship comprising the authors below.




Muhammad A. Goraya is a Network Design Specialist & Accomplished Engineer in the RAN Design team. Prior to BT, he held senior technology consultant positions with

Ericsson, Nokia, Huawei, and a number of operators. He has published widely on the latest telecom trends, particularly future networks and the potential of virtual & cloud-based solutions in the context of Open-RAN.




Abigail Griffiths is a technology graduate working in BT's Open Networks Research team. She holds a Master degree in Physics from the University of Southampton. Since

joining BT her research has focused on cloud computing and improving the energy efficiency through optimisation of telco cloud platforms.



Ian Horsley is a BT Distinguished Engineer responsible for access network research. His team focusses on the future of PON technologies, reducing power consumption, speeding up and reducing the cost

of rollout and management of FTTP technologies. He is active in industry bodies such as the ITU-T SG15 and FSN.



Louise Krug is the technical lead for BT's Sustainability Research. She has been at the forefront of developing BT's major cross-platform energy model which has driven major changes leading directly to energy and carbon reductions.



Adrian Smith is BT's Network Infrastructure Architecture Director and a BT distinguished engineer. He has spent the last 30 years helping build BT's IP core networks and is today focussed on the evolution of BT's network platforms.



Dr Yu Rong Zhou is a principal researcher and an accomplished engineer at BT working on optical networks technologies and network synchronisation playing a key role in leading edge technology trials and demonstration, she has published widely in technical journals and conferences.