Post-peak ICT: Graceful degradation for communication networks in an energy constrained future

Sofie Lambert,[1] Margot Deruyck,[1] Ward Van Heddeghem,[1] Bart Lannoo,[1] Wout Joseph,[1] Didier Colle,[1] Mario Pickavet,[1] Piet Demeester[1]

1) Ghent University-iMinds, Department of Information Technology, Ghent, Belgium

Abstract

In recent years, rising energy prices and increasing environmental concerns have boosted research in the so-called *green information and communications technology (ICT)* and *green networking* research tracks, aimed at improving the energy efficiency of communications while still offering maximal functionality. In this paper we explore a future scenario in which low power networking is no longer optional, but instead becomes a necessity due to fluctuating energy availability. The contribution of this work is twofold. First, we argue why a so-called *post-peak* future scenario – in which we can no longer rely on fossil fuels as our main resource for electricity production – is not unlikely, and what it might entail. Second, we explore the consequences of such a scenario for ICT: how well can current and future infrastructures cope with temporary energy limitations? As an illustration, we present a case study showing the impact of reduced energy availability on a wireless access network.

1 Introduction

In this paper we consider a *post-peak* future scenario, in which fossil fuels are no longer the main energy source for electricity production (fossil fuels are "past their peak"), but instead are increasingly replaced by alternative energy sources. We start our description below by motivating why a post-peak scenario may be imminent, and how this could result in temporary energy restrictions for ICT networks. Next, we outline a framework to evaluate the post-peak potential of technologies (Section 2), and propose new avenues of research to prepare ICT infrastructures for a post-peak situation (Section 3). In a specific case study for a wireless access network, we modified a wireless planning tool to make optimal use of energy in a post-peak situation (Section 4).

1.1 Why should we anticipate a post-peak future?

Present-day societies and economies run mostly on fossil fuels. Oil (petroleum), coal and natural gas made up 82% of the world's primary energy supply in 2012, and 68% of electricity was generated by burning fossil fuels, as shown in Fig. 1(a).

It is likely that in the coming years, societies will be forced to move from fossil fuels to alternative sources for their energy supply. We see three reasons why this may happen in the near future, in order of increasing importance: fossil fuel depletion, security of energy supply and the impact of fossil fuels on climate change. Some scientists argue that current trends in fossil fuel consumption will result in a peak of conventional oil production before 2030 (referred to as *peak oil*), leading to a global fuel shortage and steep increase in oil prices [1]. Whether this peak oil will indeed occur in the near future is a contested point in scientific literature, but even adversaries of the peak oil theory agree that political instability in oil-producing countries can result in severe oil price shocks, which have devastating effects for fossil fuel-based economies [2]. This brings us to the second argument: security of energy supply. Countries that switch to renewable energy sources reduce their vulnerability to the above oil price shocks [3]. Lastly and
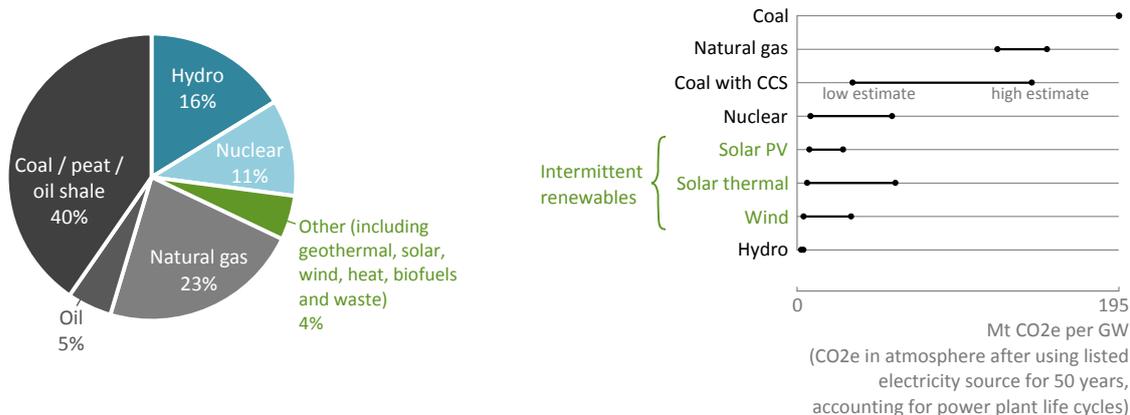
(a) Worldwide fuel shares of electricity generation in 2012 (source: Key world energy statistics by the International Energy Agency, published in 2014).

(b) Carbon footprint of various electricity sources (high and low estimates taken from [4]). CCS = carbon capture and storage.

Figure 1: In order to reduce the carbon footprint of electricity generation, the share of fossil fuels must be reduced, and intermittent renewables (solar photovoltaic, solar thermal and wind energy) will need to contribute a bigger share in the energy mix.

most importantly, if policy makers take the climate change challenge seriously, renewables offer one of few virtually carbon-neutral alternatives to fossil fuels (see Fig. 1(b)).

Nuclear energy could make an increasing contribution to low-carbon energy supply, but a variety of barriers and risks exist: not only is it costly compared to alternatives, particularly when the high risks involved are factored in; a nuclear power renaissance would also increase the risk of nuclear terrorism and make efforts to control the spread of nuclear weapons much harder.

Due to their intermittent nature, renewables replacing fossil fuels will not be able to offer the same supply continuity as present-day energy provisioning systems. Some of the most mature renewable energy technologies that can be deployed at a large scale and at a relatively low cost, such as wind and solar energy, have varying outputs that depend on fluctuating weather conditions.

Since operation of the electricity grid requires energy production and energy consumption to be in balance at all times, low production periods must either be matched with low consumption, or the production deficit must be compensated with previously stored energy and/or energy imports. In an ideal smart-grid scenario, temporary low production can be matched on the demand side by postponing non-urgent energy consumption from homes and industries. The intermittency of renewable sources could also be (partly) compensated by constructing large energy storage facilities, using e.g. hydro pumped storage or batteries; or by interconnecting production capacities over a large geographical area, taking advantage of weather locality to reduce the chance of an overall low production.

However, bringing these upgrades to the power provisioning system in place will take time, as they require a widespread introduction of smart meters and controllable devices, and the construction of large energy storage facilities and long-distance high-capacity power lines. Meanwhile the consequences of the transition away from carbon-heavy and nuclear sources may already manifest themselves in the very near future: looming electricity shortages have been reported in news outlets of several developed nations in the past year (examples include Belgium, Germany, the United Kingdom and Japan). Moreover, even once future utility networks are properly dimensioned, there may still be temporary power shortages in rare periods of exceptional weather conditions.

To conclude this motivation, we remark that even though we can't be certain that we are headed for an energy intermittent future, the outlined trends indicate that the possibility is real, and we can only deal with the consequences properly if we are prepared. Even today, this knowledge could already be useful in other energy-constrained situations, such as disaster recovery or off-grid installations in developing countries, to

make optimal use of the energy available in emergency generators and back-up batteries.

1.2 Consequences for ICT

When there is a drop in energy production (lasting from a couple of hours to several days), and the energy gap can't be filled by burning abundant fossil fuels, governments may impose temporary restrictions on power consumption to avoid a collapse of the grid. A number of measures can be taken: cutting off (geographical) clusters of consumers[1], applying dynamic pricing schemes that reflect energy availability[2], or forcing large energy consumers to curb their consumption.

If governments decide to impose energy restrictions on large consumers, ICT service providers and telecom operators will also be targeted, as their extensive infrastructures use considerable amounts of power. For example, Telecom Italia is Italy's second biggest electricity consumer, and British Telecom consumed 0.76% of the U.K.'s national electricity consumption in 2011/12, making it the largest single electricity consumer in the country. Globally, the share of communication networks in total electricity consumption is around 1.7%; data centers contribute another 1.4% [5].

Evidently, pulling the plug on ICT infrastructure altogether is not a desirable option as it supports several critical applications. But some of the services on offer may be more dispensable than others (e.g. omnipresent broadband internet access vs. lifeline communications). Therefore we introduce the concept of *graceful degradation* under energy constraints: if the available energy for ICT is drastically reduced – falling back to 50%, 25% or a mere 10% of regular power levels – can we still offer a minimal functionality and connectivity? In order to answer this question, we need to determine what part of the functionality is truly indispensable, and how much power is needed to keep it running. This will be the focus in our evaluation of ICT technologies in the post-peak context.

1.3 Related work

This is a relatively new research domain as we are only just starting to see the signs of an impending post-peak future. The most notable existing work in the field is by B. Raghavan: his 2011 publication [6] was the first to consider the implications of a permanent energy crisis for ICT, listing a number of post-peak design principles and research questions. The main difference with our work is that Raghavan considered a scenario where energy demands would exceed supply overall and the current Internet architecture could no longer be used, whereas we focus on dealing with short term energy limitations, assuming the current network architecture is still in place and functioning normally most of the time.

2 The concept of graceful degradation

In a post-peak context, we want to know how the delivered service level of a network or device scales with the available power. Three generic power profiles are shown in Fig. 2. Without graceful degradation, the delivered service level will quickly drop to (almost) zero when the available power decreases, as indicated by the lower line in Fig. 2. These infrastructures have *low post-peak potential*, as they are unable to function under energy constraints. The upper two lines in Fig. 2 represent devices or networks that do allow graceful degradation: in an infrastructure with perfect power proportionality, the service level decreases at the same rate as the available power; infrastructures with *high post-peak potential* are able to offer a high service level even if the available power is decreased a lot.

As we will discuss in more detail in Section 3, a low post-peak potential is typically associated with dedicated devices, whereas a high post-peak potential is typically available when multiple resources can be

---

[1]For example, the Belgian government installed a "disconnection plan" (Dutch: afschakelplan) after the unforeseen shutdown of two nuclear reactors in 2014. The plan allows the electricity provider to temporarily cut off clusters of consumers when demands exceed supply, for example during peak hours (17h-20h) on a cold winter day.

[2]Throughout this paper, when we talk about energy restrictions or shortages, this does not necessarily mean there is no energy available at all, but rather, that scarcity can make it prohibitively expensive.
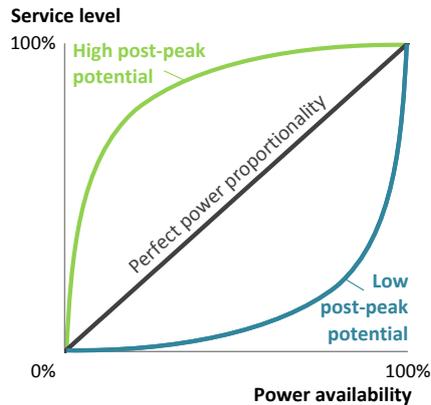
Figure 2: Depending on how their service level scales with power availability, technologies may have *high post-peak potential*, offering relatively good service for a fraction of their normal power, or *low post-peak potential*, losing service level rapidly even for a small power reduction. This relationship is influenced by the choice of service level metric, which can be based on criteria such as speed, reliability, user coverage,...

flexibly shared. The shape of the power profile of the device or network under study depends on the power proportionality that it exhibits, as well as on how its service level metrics are defined, as explained below.

Power proportionality — In Fig. 2, 100% power availability corresponds to what the device or network consumes in a normal energy situation to offer maximal service. This maximum power consumption will depend on the specific equipment under study, and whether it uses current (energy-hungry) or future (greener) technologies. The profile is also influenced by the impact on power consumption of a service reduction. Currently deployed networks typically exhibit a power consumption that is relatively constant despite strong diurnal variations in the traffic load or service, corresponding to the lower curve of Fig. 2. Reducing the maximum power consumption and improving the power proportionality (making power consumption scale with load, e.g. through load-adaptive network operation) have been important tracks in Green ICT research [7], and will become even more important when targeting graceful degradation for post-peak ICT.

Service level metrics — For a given device or network, the power profile will also depend strongly on the metrics used to define the service level. This can be a quantitative assessment, based on a weighted function of data throughput, bandwidth, uptime, error rate, user coverage, computation capacity, etc.; or it can be more qualitative, when certain applications or services are considered critical and thus part of the *minimal service*, while others are dispensable. For an existing example of graceful degradation, consider a smartphone with a battery that is running low. Certain applications, such as the camera functionality, may be temporarily disabled to ensure that the minimal service, texting and making phone calls, is guaranteed for as long as possible. If we define the service level of the smartphone as the number of calculations per second, the service under energy constraints may be quite poor, but if we define the service level as the ability to communicate, it is still relatively high.

3 Graceful degradation in ICT infrastructures

In this section we assess the potential for graceful degradation in various ICT infrastructures. We start by introducing some terminology concerning the areas under study; next we look into specific solutions for each of these areas.
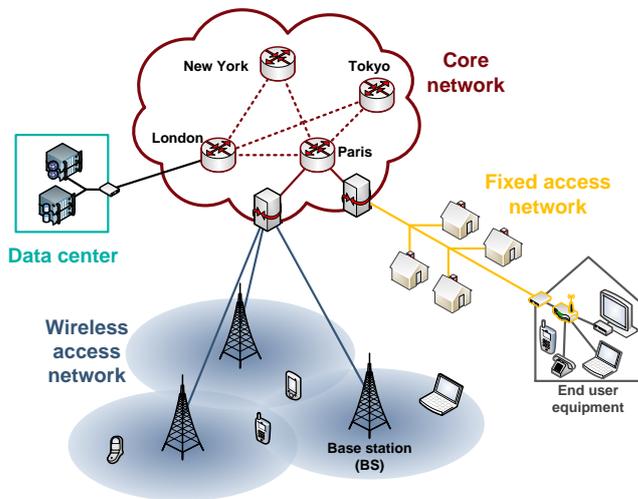
Examples of post-peak power reduction measures
Fixed access network
• Improve customer premises equipment (CPE) power proportionality
• Centralized control of CPE power states
Wireless access network
• Selectively switch off some BSs, based on different strategies:
– lowered capacity or reduced covered area
– support fewer wireless technologies
Core network
• Reduce level of resilience or overprovisioning
• Change of virtual topology
• Block specific traffic
Data center
• Reschedule user requests:
– to a later time
– to another location

Figure 3: Schematic of the network showing the areas under study (fixed and wireless access networks, core networks and data centers), along with examples of post-peak solutions for these areas.

3.1 Overview of ICT infrastructures

Fig. 3 gives an overview of a typical network infrastructure. The access network provides a physical connection to the end users through which they can connect to the Internet. In *fixed access networks*, a physical wire runs to the end user's premises, where the signal is decoded by a modem and distributed further in the local home or enterprise network. In *wireless access networks*, subscribers use standardized radio signals to connect to the nearest base station (BS), from where the signal is forwarded through a dedicated backhaul link to an aggregation point. Traffic from the access networks is aggregated and transmitted further through the backbone or *core network*, which provides high-capacity, low-latency connections across large geographical distances .

The focus in this paper is on devices that can be controlled centrally by the internet service provider (ISP). This means that, despite their significant contribution to ICT power consumption, end devices controlled by the consumer such as TVs, personal computers or mobile phones are outside the scope of this work. Modems, routers and wireless access points (APs), installed at the customer premises, may be considered part of the fixed access network, depending on whether they can be controlled remotely by the ISP. *Data centers* can also be controlled centrally, and form the last important contributors to network power consumption. These facilities house large numbers of computer and storage systems to host a wide range of applications, from websites over search engines to cloud computing.

3.2 Field-specific solutions and limitations

In the following, we explore the opportunities and limitations for graceful degradation in the four fields introduced above. The list of proposed measures is by no means exhaustive, but rather intended as a starting point for further research. While we discuss the solutions for these fields separately, post-peak strategy design should also keep a holistic overview of the interactions between different fields, and see whether some services can be substituted for others. For example, guaranteeing both wired and wireless connectivity may not be feasible under certain energy constraints, but as long as one of the communication channels remains available, this may be sufficient as a minimal service.

5

3.2.1 Fixed access networks

Since most of the power in fixed access networks is consumed by the customer premises equipment (CPE) [8], this is where we direct our first efforts to try and optimize the use of limited power resources. Currently deployed modems, wired and wireless routers typically show bad power proportionality with respect to data rates, corresponding to the low post-peak potential profile in Fig. 2.

However, if we define the service level as the number of customers being served, and consider the power consumption of the access network as a whole instead of on a single device level, switching off a selection of CPEs will scale down power consumption as the service level decreases, corresponding to a more power-proportional profile (the straight line in Fig. 2).

Depending on how the minimal service is defined, several strategies can be used to power off a well-chosen selection of access equipment. A *time division based strategy* could divide the access network into regions, and assign time slots for each region when connectivity would be available, powering off all access equipment in that region for the remaining time. This would give users periodic, deterministic access to the network. Alternatively, in dense urban networks a *wireless ad hoc strategy* could be applied, taking advantage of the relatively dense deployment of wireless APs. Here, a large fraction of the modems and APs could be powered down, while almost complete wireless coverage of the area (at reduced data rates) could be maintained by connecting users to neighboring hot spots. Note that for these strategies to work, the network provider should be able to power down modems and routers remotely.

In larger office environments, local area networks (LANs) can be structured so that *less critical parts* can be selectively powered off while critical connections are maintained.

3.2.2 Wireless access networks

In wireless access networks, BSs are the main power consumers. The power consumption of a single BS does not scale well with the traffic load, corresponding to low post-peak potential in Fig. 2. As a consequence, just as in fixed access networks, the only way to reduce power consumption in current networks is by switching off some of the equipment and looking at the service level on the network scale instead of the device scale. We propose two strategies for BS switch-off below, designed to get the maximum possible service out of the limited available energy during a post-peak shortage.

The first is to *cut capacity by reducing network density*. There is a trade-off between the capacity and the range of a single BS: for a given input power, decreasing the capacity (bit rate) results in longer ranges and thus a larger area covered by the BS. In a post-peak situation, the capacity per user can be limited so that the required network capacity would be reduced and, since the range of the BSs in a low-capacity network is bigger, the number of active BSs could be decreased. The achievable savings are calculated for a realistic case in Section 4. Depending on what fraction of the normal operating power is available, this strategy may result in reduced coverage. A wireless ad-hoc functionality for emergency communications, where information can be relayed through other end users' devices to reach the access network, may be worth investigating.

The second strategy is to *fall back to single-standard support*. In current mobile networks, overlapping coverage is offered for different standards. The most widespread technology is the second generation (2G) standard Global System for Mobile Communications (GSM). Third-generation (3G) Universal Mobile Telecommunications System (UMTS) and fourth-generation (4G) Long-Term Evolution (LTE) sites are often built at pre-existing 2G sites to increase peak data rates and the maximum number of user connections. During a post-peak temporary energy shortage, power consumption may for example be decreased by disabling the 3G and 4G BSs, falling back to the GSM network (GSM is currently still the most widely supported technology on handsets). This could result in a drastic reduction in power consumption, since BS consumption is similar across technologies (in the order of 1-2 kW/BS) [9]. For subscribers, this measure would result in a noticeable reduction in quality of experience, as network capacity and peak data rates would be reduced; but coverage would still be guaranteed.

3.2.3 Core networks

During a post-peak shortage, the traffic load in the core network will decrease significantly as a consequence of traffic reductions realized in access networks and data centers. Unfortunately, because of the current weak power proportionality in response to changes in the traffic load, the power consumption in the core network will remain almost constant, once again corresponding to a low post-peak potential profile in Fig. 2. In future deployments we might see new equipment and techniques that improve power proportionality, such as bit rate variable transponders and power aware routing schemes that increase the potential for sleeping interfaces [10]. In anticipation of these developments, we can think of four other approaches that could keep core networks running under post-peak power reductions.

The first is through *reduced resilience*. Typically, protection mechanisms are in place to almost instantly switch traffic over a secondary path in case of a link or node failure. Turning off backup equipment could reduce power consumption by a factor of two. As a result, recovery would not be instant (within 50 ms) anymore, but would take a couple of seconds when a recovery path is discovered and set-up automatically, or several minutes or hours when links need to be brought back online manually (in case of insufficient capacity). A second, similar approach is to *temporarily reduce the overcapacity* that is installed to handle peak-to-mean traffic variations and unexpected traffic spikes [11]. This could decrease power consumption in the core by another factor of two (or more). However, unexpected traffic spikes would not be handled flawlessly anymore and traffic bottlenecks would occur in anticipation of extra capacity being (manually) brought online. Third, we could *change the virtual network topology*. Several works have indicated that, with traffic demands currently being higher than the equipment line rates, fully-meshed virtual (IP) topologies are more energy efficient than ring topologies [12]. Further research is needed to investigate which topology (mesh, ring,...) would be most efficient in a post-peak scenario with reduced traffic demands, and how it would affect latency. The fourth and last approach we propose is *blocking service-specific traffic*, where non-critical services are temporarily blocked in the core, insofar as this is not already done at the network edge. This technique could also be used to reduce traffic demand to a level where a topology reconfiguration (introduced above) can bring additional savings.

3.2.4 Data centers

In [13], the authors discuss how data center management should be revised to maximize the use of off-grid renewables. These techniques can be adapted for use in a post-peak scenario, taking into account the more restrictive energy limitations. Below we list a number of ways to temporarily reduce the workload in data centers. Note that these measures can only reduce power consumption significantly if the data center manager is able to plan capacity by turning off selected groups of machines when the load is reduced, thus achieving high post-peak potential (cf. Fig. 2).

The first logical step is to delay system maintenance tasks such as system updates and backups. Incoming user requests can also be *rescheduled* to a later time or *migrated*, if possible, to another location where energy is more abundant. If workloads cannot be migrated but must be reduced, a *priority label* could be given to the most critical data and services, which would then be placed on machines that are kept on at all times to guarantee the minimal service. This priority label could be assigned manually based on service level agreements (SLAs) with customers, or automatically to the most frequently accessed content. Alternatively, all user requests could be handled with the same priority, but with less resources, keeping all services available, but at the cost of longer response times, including that of critical services.

4 Case study: wireless access network with reduced power

4.1 Realistic scenario for a suburban area

As a concrete case study, we evaluate the post-peak potential of a wireless access network by simulating an LTE (4G) network in a suburban area of Ghent, Belgium. Fig. 4 shows the selected area together with the
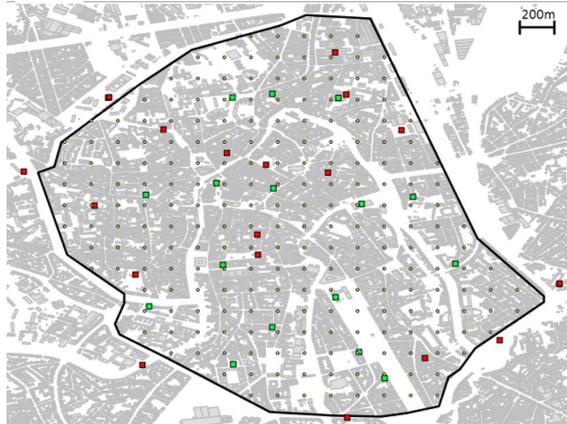
Figure 4: The selected suburban area of 6.85 km^2 in Ghent, Belgium. Users are located on the yellow dots; squares indicate base station (BS) sites. As an illustration of the sleep mode principle, the active and inactive BSs for a simulation with 70% available power and random user distribution are colored green and red, respectively.

Table 1: Properties of the considered LTE (4G) radio technology.

Property	Value
Carrier frequency	2600 MHz
Channel bandwidth	5 MHz
Minimum power/active BS (antenna input power 1 dBm)	1204 W
Maximum power/active BS (antenna input power 43 dBm)	1672 W
BS power in sleep state	
- scenario zero sleep power	0 W
- scenario 45% sleep power	752 W
(Coding rates) Modulation schemes	(1/3, 1/2, 2/3) QPSK
	(1/2, 2/3, 4/5) 16-QAM
	(2/3) 64-QAM
Bit rate/BS [1/3 QPSK - 2/3 64-QAM]	2.8 Mb/s - 16.9 Mb/s
Range/BS [1/3 QPSK - 2/3 64-QAM] (NLOS[*], @43 dBm)	1090 m - 194 m

[*] NLOS: non-line-of-sight.

locations of the 35 simulated BS sites (based on real BS positions), and Table 1 lists the LTE characteristics.

We assume 224 candidate users who want to be simultaneously active, each requesting 64 kb/s, which is more than enough to make a phone call. The location of the users within the selected area is known (a realistic mix of indoor and outdoor users is assumed), but the order in which they connect to the network varies randomly across simulation runs. A heuristic deployment algorithm, based on the one described in [14], chooses which BSs are active and which are in sleep state. The objective of the algorithm is to maximize the number of active users, i.e. users connected to an active BS, by switching BSs on and off in response to the simulated user locations and load, and setting the power level for each BS based on its optimal reach. The original algorithm in [14] was designed to achieve a predefined QoS and coverage with minimal power consumption. In the post-peak version, we introduce an additional constraint: the overall power consumption (summed over all BSs) cannot exceed a pre-defined power budget. This may result in some loss of QoS and coverage.

We consider two scenarios for the power consumption of BSs in sleep mode: perfect sleep, where a sleeping BS consumes no power; and a more realistic sleep power, which would typically be around 45% of the maximum power [15] (see Table 1 for values).
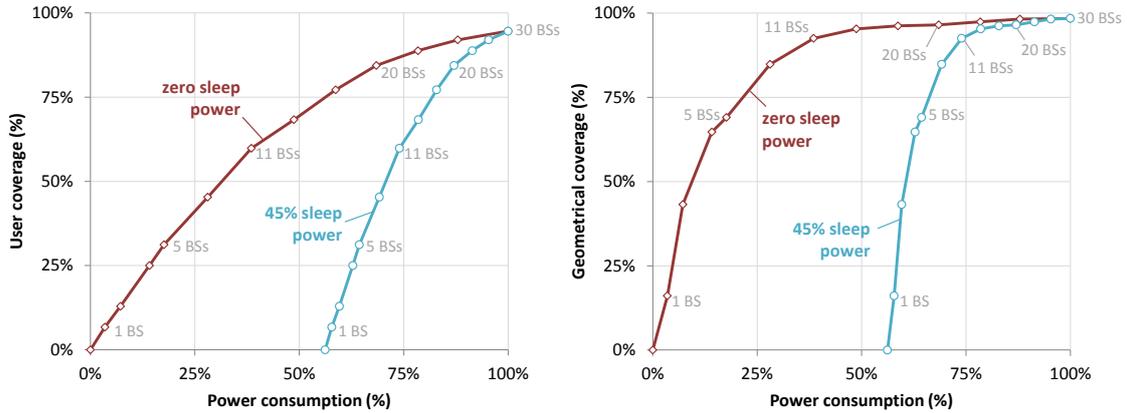
Figure 5: Simulation results of the wireless case study for LTE (4G) technology. When base stations (BSs) consume negligible power in the inactive state (zero sleep power, dark red lines), both user and geometrical coverage have high post-peak potential. When BSs require a significant fraction of their maximum power in sleep state (45% sleep power, light blue lines), most of the post-peak potential is lost.

4.2 Simulation results

The heuristic deployment algorithm is applied to the area under study with varying constraints on the overall power consumption. This results in the power profiles shown in Fig. 5, where each simulated point corresponds to the median of 40 simulation runs. The horizontal axis in Fig. 5 shows the network power consumption as a percentage of the maximum power consumption (100% = power for a network optimized for maximum coverage). The vertical axis shows the service level, which can be defined using two different metrics. The *user coverage* indicates the percentage of users that can get the minimum bit rate they are requesting from the active BSs (in this case, 64 kb/s). The *geometrical coverage* indicates the percentage of the (outdoor) area that can be reached by signals from the active BSs at the lowest coding rate and modulation scheme (1/3 QPSK), offering a bit rate of 2.8 Mb/s per BS (see Table 1). User coverage is a more strict metric, as a geographical location may be within range of a BS while a user in that location may not be covered as such, due to an uncovered indoor position or capacity limitations of the BS.

When BSs consume negligible power in the inactive state (zero sleep power, dark red lines in Fig. 5), the post-peak potential for user coverage is limited: when 25% of the power is available, user coverage for LTE is only around 30-40%. The power profile for geometrical coverage looks more promising, since about 75% coverage can be guaranteed with only 25% of the power. This is because once most of the area is covered, the energy cost of adding geometrical coverage increases as filling a small coverage gap will still require a complete additional BS to be switched on, and the minimum power per BS is 1204 W (cf. Table 1).

The geometrical coverage could be considered the fraction of users that can be served when only outdoor coverage is guaranteed and extremely low data rates are allowed, offering basic communications such as text messaging. Depending on how the minimal service is defined – should it include voice communications and possibly even low-data rate exchange; what about indoor coverage? – one could opt for higher geometrical coverage (avoiding large outdoor coverage gaps) or better user coverage in densely populated areas (leaving large coverage gaps in remote areas, but keeping indoor coverage and higher rates available in urban areas).

When BSs require a significant fraction of their active power in sleep state (45% sleep power, light blue lines in Fig. 5), most of the post-peak potential is lost: 56% of the maximum power is still needed even when *all* BSs are in sleep mode (=minimum network power consumption) and thus no service at all is offered. Note that theoretically only 45% of the power would be needed if the network were perfectly dimensioned and all BSs were switched on at maximum load. However, activating 30 out of 35 BSs sufficed to reach the maximum attainable coverage for maximum load; hence the 100% power consumption already

corresponds to the consumption of 30 active and 5 sleeping BSs. In any case, the post-peak potential of a deployment with relatively high sleep power is clearly insufficient, as no service at all can be offered when energy availability is temporarily below 50%.

5   Conclusion

5.1   Conclusions of this work

We studied the effect of post-peak energy shortages on various ICT infrastructures: fixed and wireless access networks, core networks, and data centers. We assessed whether these infrastructures could still offer a basic functionality when the available power was significantly lowered compared to normal operating conditions.

We distinguished between low and high post-peak potential infrastructures and devices (Fig. 2). Based on the ICT fields we studied, we conclude that the post-peak potential on a single device level is typically low: most devices are unable to offer half of the normal service (defined as number of calculations, data rate,...) when the available power is halved. Therefore, the post-peak potential of dedicated devices to which users are statically assigned, is very limited. In networked environments on the other hand, when users share resources and *flexible* resource assignment is possible, the post-peak potential on a wider scale can be much higher. This is for example the case in large-scale data centers or wireless access networks, where some of the capacity can be temporarily disabled while still offering a minimal service to all users.

We also considered the concrete case of a suburban wireless access network under varying energy constraints. Simulation results showed that the post-peak potential can only be realized if low power modes are truly low-power and inactive BSs consume negligible power. This is in contrast with present-day sleep modes for BSs which typically consume around 45% of the active power.

5.2   Moving forward

This is a relatively new research domain, and though we touched upon a number of strategies already, handling a post-peak situation will no doubt require further research into a broad range of applications and domains. When investigating post-peak features, there are a number of pitfalls that need to be kept in mind.

A first issue is that indirect effects of a post-peak situation on user behavior are hard to predict. For example, offering a slower service could lead to a reduced load if users give up non-urgent activities in response to the reduced quality of experience, but it could also increase the load if users' connection times are increased due to the slower service. Indirect effects on user behavior could also be used to enhance post-peak savings, for example by suppressing TV signals for the indirect effect of having less television sets switched on. Interdisciplinary research is needed to predict these kinds of effects.

A second important issue is that systems will become more vulnerable at the time of energy shortages. While solutions should be chosen carefully to minimize vulnerability, some loss of reliability during a post-peak energy shortage will be inevitable: we should keep in mind that a minimal service level is still preferable to no service at all.

Last but not least, ICT is not only an important power consumer, it can also play a key role in controlling the power consumption of other loads. In a post-peak future, a *smart grid* that gathers information about suppliers and consumers and helps control them, will be an important instrument to balance electricity supplies and demands. This should be studied in a dedicated research track, considering incentives that shape consumer behavior and opportunities for automated power control. Post-peak measures for ICT specifically will need to guarantee the exchange of smart grid control signals. Further research is needed to assess how smart grid control can be set up independently from existing communications infrastructures, or how it can be guaranteed as part of the minimal service.

Acknowledgments

References

[1] S. Sorrell, J. Speirs, R. Bentley, A. Brandt, and R. Miller, "Global oil depletion: A review of the evidence," *Energy Policy*, vol. 38, no. 9, pp. 5290–5295, Sep. 2010.

[2] M. Radetzki, "Peak Oil and other threatening peaks—Chimeras without substance," *Energy Policy*, vol. 38, no. 11, pp. 6566–6569, Nov. 2010.

[3] C. Lutz, U. Lehr, and K. S. Wiebe, "Economic effects of peak oil," *Energy Policy*, vol. 48, pp. 829–834, Sep. 2012.

[4] N. P. Myhrvold and K. Caldeira, "Greenhouse gases, climate change and the transition from coal to low-carbon electricity," *Environmental Research Letters*, vol. 7, no. 1, p. 014019, Mar. 2012.

[5] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Trends in worldwide ICT electricity consumption from 2007 to 2012," *Computer Communications*, vol. 50, pp. 64–76, Sep. 2014.

[6] B. Raghavan and J. Ma, "Networking in the long emergency," *Proc. ACM SIGCOMM Workshop on Green Networking*, 2011.

[7] C. Lange, D. Kosiankowski, R. Weidmann, and A. Gladisch, "Energy consumption of telecommunication networks and related improvement options," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 17, no. 2, pp. 285–295, Mar. 2011.

[8] J. Baliga, R. Ayre, K. Hinton, and R. S. Tucker, "Energy consumption in wired and wireless access networks," *IEEE Communications Magazine*, vol. 49, no. 6, pp. 70–77, Jun. 2011.

[9] M. Deruyck, E. Tanghe, W. Joseph, and L. Martens, "Modelling and optimization of power consumption in wireless access networks," *Computer Communications*, vol. 34, no. 17, pp. 2036–2046, Nov. 2011.

[10] F. Idzikowski, E. Bonetto, L. Chiaraviglio, A. Cianfrani, A. Coiro, R. Duque, F. Jiménez, E. Le Rouzic, F. Musumeci, W. Van Heddeghem, J. López Vizcaíno, and Y. Ye, "TREND in energy-aware adaptive routing solutions," *Communications Magazine, IEEE*, vol. 51, no. 11, pp. 94–104, Nov. 2013.

[11] D. Kilper, K. Guan, K. Hinton, and R. Ayre, "Energy challenges in current and future optical transmission networks," *Proceedings of the IEEE*, vol. 100, no. 5, pp. 1168–1187, 2012.

[12] W. Van Heddeghem, F. Idzikowski, F. Musumeci, A. Pattavina, B. Lannoo, D. Colle, and M. Pickavet, "A power consumption sensitivity analysis of circuit-switched versus packet-switched backbone networks," *Computer Networks*, vol. 78, no. 0, pp. 42 – 56, 2015, special issue: Green Communications.

[13] C. Stewart and K. Shen, "Some joules are more precious than others: Managing renewable energy in the datacenter," *Proceedings of the Workshop on Power Aware Computing and Systems*, 2009.

[14] M. Deruyck, W. Joseph, E. Tanghe, and L. Martens, "Reducing the power consumption in LTE-Advanced wireless access networks by a capacity based deployment tool," *Radio Science*, vol. 49, no. 9, pp. 777–787, 2014.

[15] M. Gonzalez, D. Ferling, W. Wajda, A. Erdem, and P. Maugars, "Concepts for energy efficient LTE transceiver systems in macro base stations," in *Future Network Mobile Summit (FutureNetw), 2011*, Jun. 2011, pp. 1–8.