HOW AI REDUCES NETWORK ELECTRICITY BILLS AND CARBON FOOTPRINT

With energy costs on the rise, some good news at last

Context

The major sources of environmental impact for a telecom operator are their manufacturing footprint for network elements, direct electricity consumption, and the footprint associated with customers' devices. Notably, the industry represents 3-4% of global CO_2 emissions[1], and more than 70% of the energy spent by a mobile CSP derives from the radio access network, accounting for 17% of their total carbon footprint[2].

A typical operator with 15 million mobile subscribers spends 50-100M€ per year on electricity, based on 2021 prices (extrapolated from[3]). Rising energy costs will double this (or worse) in 2023, making it a growing source of OPEX.

Increasing network and bandwidth usage fights against more efficient network technologies, such as 5G and improved network and equipment design, getting us a 5-10% increase of electricity consumption year over year.

Optimizing the electricity consumption of networks is therefore a critical area of focus for CSPs and the global environment, making it a topic of choice to be addressed by the most modern techniques, including data analytics and AI.

This report will answer the following key questions:

- What are CSPs already doing to reduce their energy consumption?
- How do AI and Data drive the journey to estimate, measure, understand and optimize a CSP's network energy consumption?
- What is Project BOSE, Capgemini's solution for energy optimization?

1. https://www.bcg.com/publications/2021/building-sustainable-telecommunications-companies
2. Internal Capgemini benchmark based on CDP reports from 31 major CSPs
3. https://frankrayal.com/2022/05/10/how-much-mobile-operators-care-about-energy-consumption/

What is already being done

Before exploring the benefits of data & AI, we should review the major devices used to reduce electricity consumption today and those envisioned for tomorrow.

- **Improved hardware and software design:** NEPs (Network Equipment Providers) use better components and assemble them in a more energy-efficient way, due to the improvement in chip technologies, but also in better system design approaches. 5G is going to have a lower footprint than 4G for comparable bandwidth (but we still expect overall higher footprint due to increased bandwidth usage).
- **Energy-aware network design:** energy-aware design techniques are being applied to new networks, such as better design for cooling in data centers, with passive cooling instead of active cooling elements.
- **Static on/off mechanisms:** these enable CSPs to decide to switch off certain elements based on static scheduling, especially in urban areas when off-peak times do not require full capacity. Enough security margins are required to avoid impact on quality of service.
- **Energy-saving features:** RAN vendors provide the capacity to automatically reduce power and switch on and off capabilities based on static thresholds. This can provide still another layer of adaptation and energy saving, but the flexibility and intelligence of these automations is limited, and human expertise is needed to make them efficient.

These devices rely on technology and design progress, or on manual activity, and thus require human expertise. They cannot adapt finely in real-time to changing usage conditions, while maintaining a high quality of service.

How AI can help – the 4 levels of maturity

We want to go further, with existing networks, without adding to operations teams' workloads. More intelligent automation – based on better insight into the network – is necessary, and AI can deliver. To illustrate how data & AI can help, we will use a four-level scale:

The first level is about **estimating the global environmental footprint**, and the energy footprint in particular. Typically, physical models of network elements and networks are built, and then aggregated from bottom to top, to produce an estimate of energy consumption. This requires a deep understanding of the physical elements of the network – down to each circuit board – and potentially information from NEPs when available.

The precision of static estimation is limited, as traffic patterns are varying, and they impact the overall efficiency and consumption. Optimizing consumption is possible with "what if" scenarios in the design of a network, or in changes to be done, but cannot adapt to the behavior on a live, running network.

The second level is about **collecting real, detailed consumption data**, **from each network element and centralized real detailed consumption**. Such data can be obtained from the OSS, and are generally part of the KPIs aggregated in Performance Management Systems, alongside network usage and quality of service data.

They can then be transferred to a **sustainability data hub**, the data repository collecting all the organization's footprint and consumption data. The total footprint can then be calculated, day by day, hour by hour, based on actual consumption affected by varying network usage. A most accurate view is thus provided of the live network. The volume of data is quite high, but well within the capabilities of today's data platforms, especially those based on cloud technologies. Source data models are proprietary, dependent on equipment vendors. They thus need to be normalized in the sustainability data hub to be compared and reported in a homogeneous way.

The third level of maturity uses collected data to build **statistical models of the electricity consumption of a live network, and understand consumption factors and pattern** taking into account a maximum of factors affecting consumption, such as number of users, bandwidth, type of service (3G/4G/5G, voice/data), movement (and thus inter-cell handovers), network equipment model, end-user device models, local wave propagation factors, even humidity. Machine learning models can thus be built and validated, provided that enough data is available, in duration (are we covering a long enough time to have seen all configurations and variations), variety (do I get urban as well as rural areas? Different equipment models? Holidays times?...) and quality (is my data verified and correct? Do I get the same coverage for all networks, or am I missing some critical data elements?).

Such models are providing actionable benefits in the pursuit of carbon neutrality and electricity bill reduction:

- Models enable the **identification of key factors** in emission reduction and increase, thus enabling the optimization of processes for better efficiency. They make it possible to answer questions such as: what is the impact of switching off some cells at night? How is the increased traffic affecting electricity

Figure 1: Maturity level for optimizing footprint and energy consumption

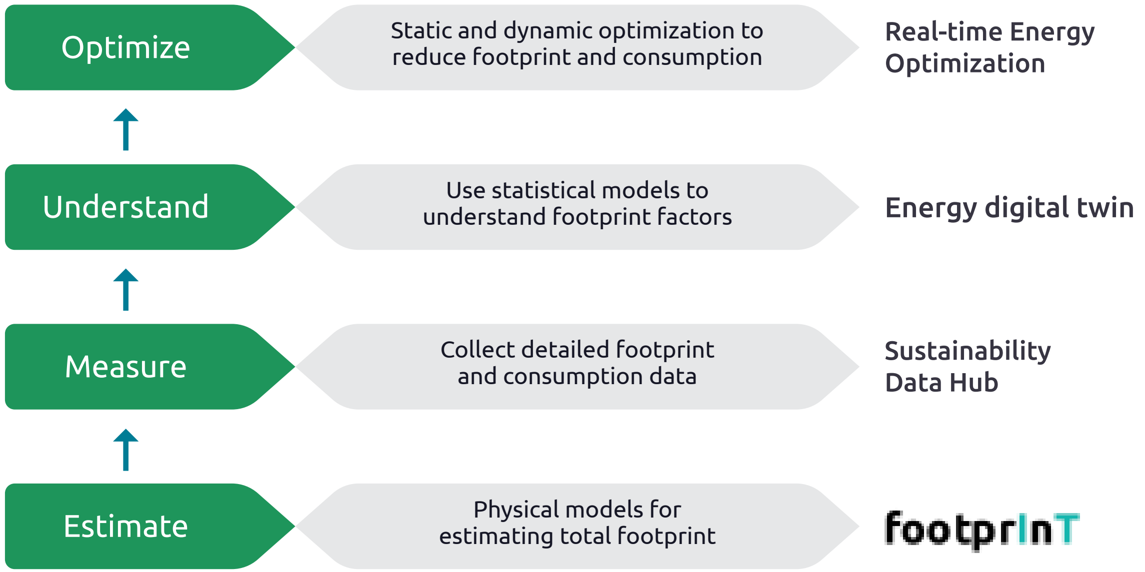

consumption? Which brands of network elements are proving to be more cost-effective? Should I replace older parts of the network?

- As CSPs are now implementing energy-saving features from NEPs, they get some level of dynamic adaptation to network usage, configured manually, usually based on general categories of towers with fixed thresholds. Statistical models help you **examine different configurations for those energy-savings features**, and compare the various possible configurations, based on types of cells (4G/5G, urban / rural, usage patterns…). You can then generate tailored configurations, specifically adapted to each cell, and not just based on rough categories.
- Finally, having a model enables the **identification of anomalies**: if a specific tower is deviating from the model, then it is performing under different, unidentified parameters. It is worth having a look, maybe to identify an error in the model, to find better parameters, or in reverse to uncover problems in specific elements – let's say a cell tower which is badly parameterized and wasting electricity.

Verizon has documented a similar approach, building an energy digital twin to understand and optimize energy consumption, saving more than $100M per year. Their energy digital twin combines a central repository of energy consumption readings from the network, data normalization and statistical models[4].

The fourth and final stage is **dynamic optimization to get maximum savings**: usage is estimated and anticipated per cell, for the whole network, based on machine learning models. Every day, every hour, every minute, future traffic is estimated, as well as the capacity required to provide adequate quality of service. Another set of machine learning algorithms is trained to derive optimal equipment parameters for each configuration of traffic. This enables CSPs to apply – on regular cycles, down to every few seconds or minutes – adequate configuration for each tower based on current and anticipated usage patterns. Energy consumption is thus optimized locally.

Verizon has documented a similar approach, building an energy digital twin to understand and optimize energy consumption, saving more than $100M per year. Their energy digital twin combines a central repository of energy consumption readings from the network, data normalization and statistical models.

4. https://inform.tmforum.org/features-and-opinion/how-verizon-is-using-digital-twins-to-reduce-energy-costs/

Going further: optimization for entire networks

A local optimization policy will only consider cell-level or tower-level optimization. For still better performance, the centralized, global view of the network can be used to provide inter-cell optimization. For instance, specific elements in some towers can be switched off, transferring usage to adjacent towers. Applying such optimization without precise data and measurement would lead to uncertainties: am I really saving electricity? Am I not compromising quality of service? Am I leaving enough margin? Having precise data in real-time and estimating future usage allow you to switch off full radio heads with the confidence that you will be able to switch them back on in time to provide capacity when usage rises.

We need a high level of control and interaction with the radio network. An API-based integration with an EMS (Element Management System) or a SON controller can provide adequate control. Some adaptation is required to specific technologies and vendors. Such control is already possible with 4G networks, especially with recent releases providing advanced energy-saving features. 5G will provide another level of capabilities, and ORAN will provide still higher openness and standardization of interfaces.

A few operators have actually tested such dynamic energy management in mobile networks, with first results ranging from energy savings of 8% to 20% and more. We present in the next section project BOSE, a lab experiment in which Capgemini has obtained up to 18% savings.

	Data needs	Complexity	Precision
Optimize	Enables configuration parameter for RAN elements	Predictive capabilities on usage and QoS. Algorithmic optimization	High precision
Understand	Context: network usage, devices, quality of service, environment, topology	Algorithmic complexity: sophisticated machine learning models are required to reach high accuracy	High precision and understanding of key energy consumption factors
Measure	Electricity consumption	Higher volume of data, simple calculation	Accurate view of consumption if high-quality data is available
Estimate	Inventory and topology	Low, as only a static view of networks is needed	Limited, challenged by varying traffic patterns

Project Bose

Project Bose is a joint project between Capgemini and Intel. It uses a data-driven approach to 'enable sustainable 5G networks in a smarter way'. Cagpemini's NetAnticipate network AI framework is intimately connected with critical elements of a 5G network, using AI to understand and anticipate loads and usage patterns to configure them for minimal energy consumption, thus implementing the full layer 4 "Optimize" function. We obtained an average of 18% energy savings in a lab environment[5].

Further work is under way to continue driving innovation, increasing the scope and efficiency of Capgemini's project Bose Energy Optimization framework.

Project Bose overview

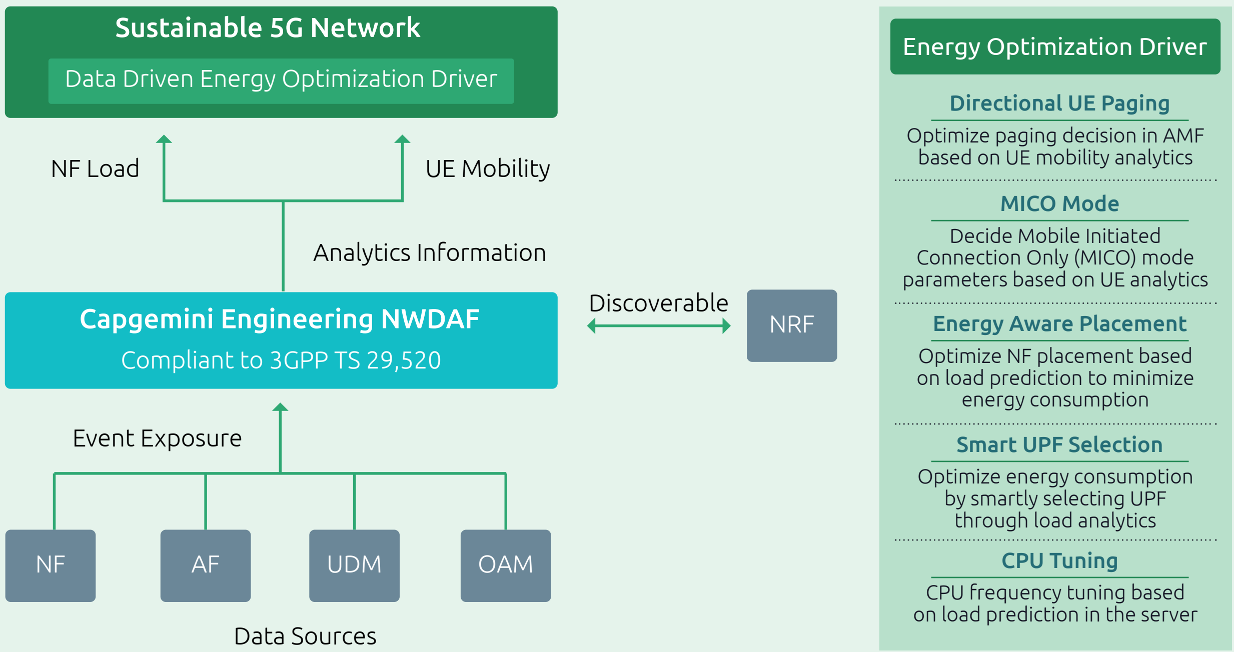

5. https://www.capgemini.com/insights/research-library/project-bose-enabling-a-sustainable-5g-network/

Conclusion

High electricity consumption, associated cost and footprint are becoming major problems for CSPs. Newer network elements may lower consumption, thanks to manufacturers improving technology and benefiting from lower-consumption chips. But manufacturing and swapping has a footprint also, and it is worth considering optimizing within a layer of intelligence added on top of existing technologies, minimizing additional hardware. We explored in this paper how data analytics and artificial intelligence can help in a multi-level approach, adding software solutions on top of existing networks, to reduce the consumption of existing network elements.

Such AI-based optimizations can be applied to existing 4G and 5G networks, and are adaptable to future 5G and 6G technologies, which will provide new capabilities and openness to further facilitate optimization. This will become increasingly necessary, as with growing network usage those new, more energy-efficient technologies will still require additional effort to lower electricity cost and footprint as much as possible.

There's no better time to start than now. Let's talk!

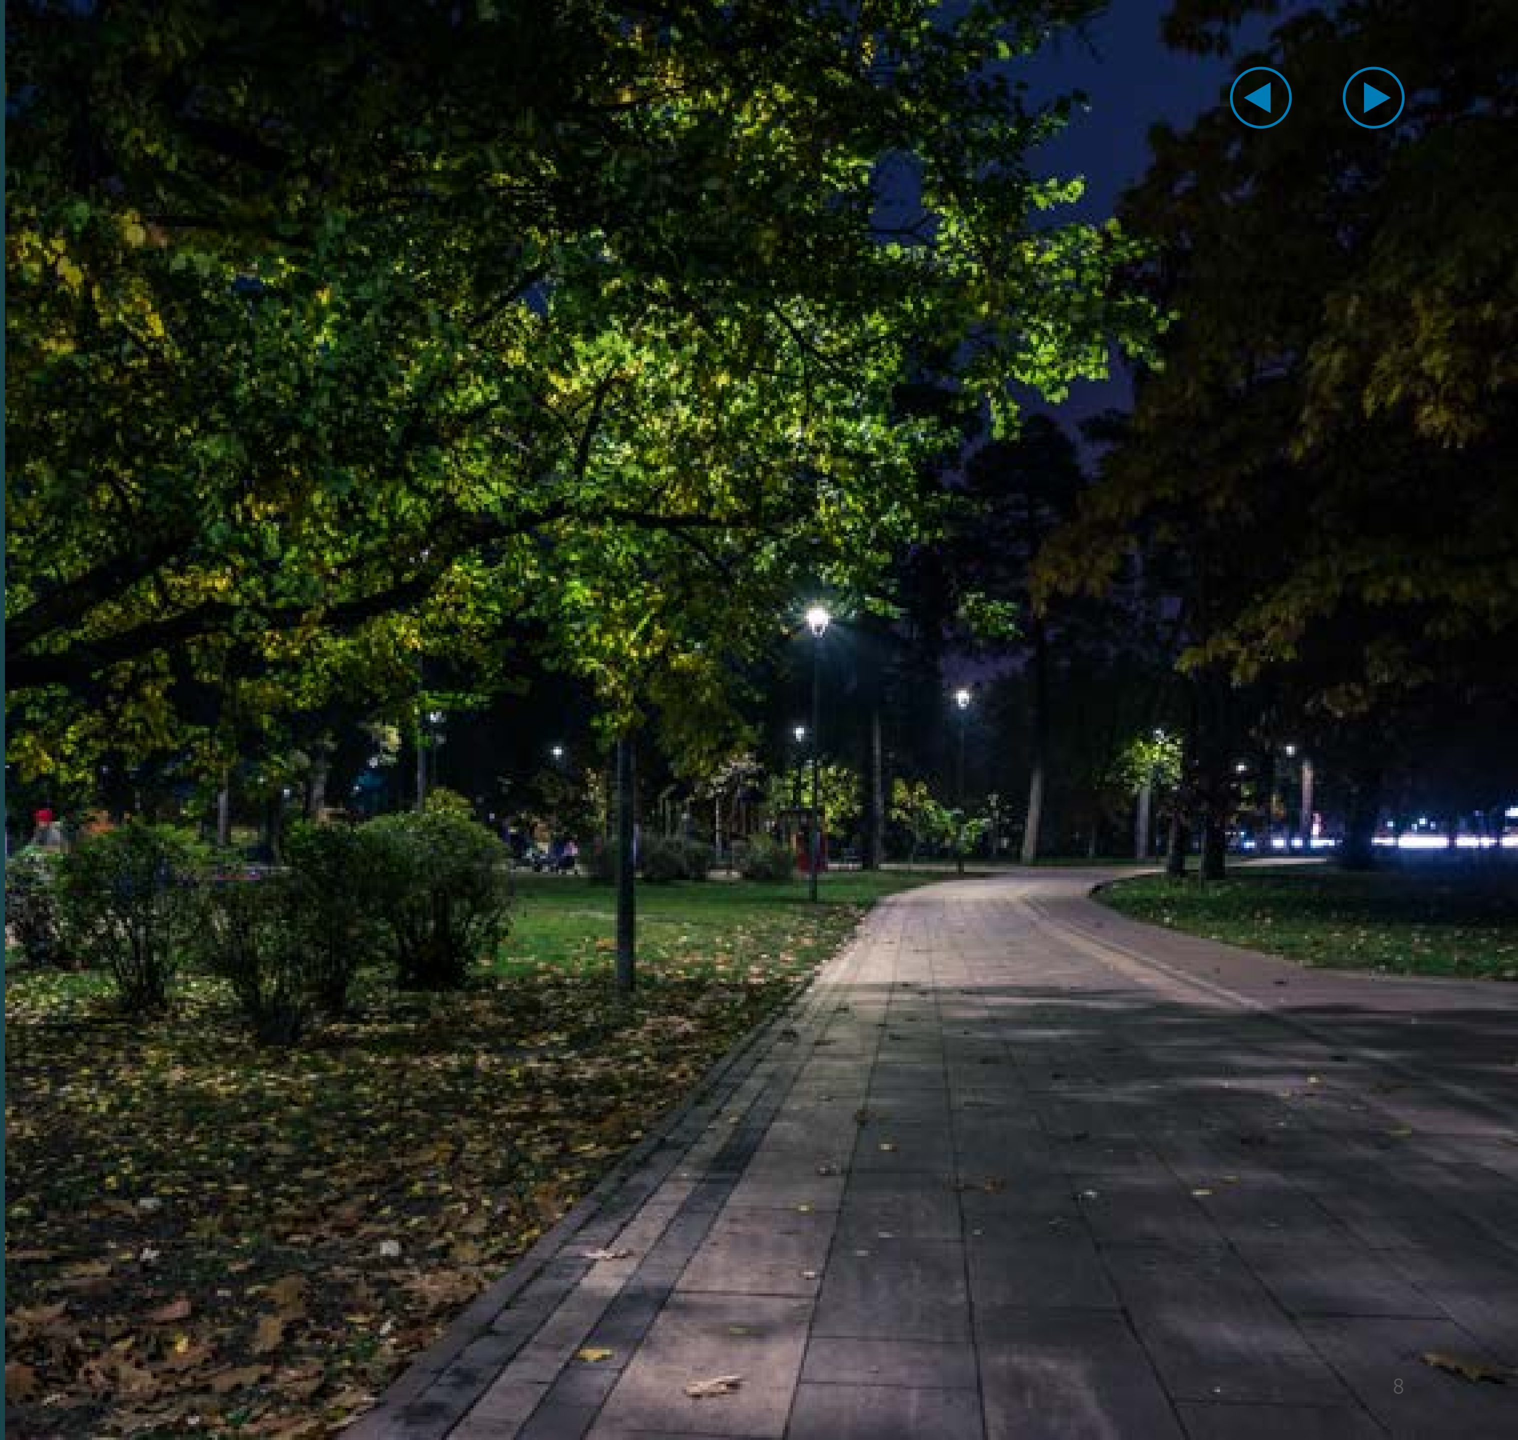

About Capgemini

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 350,000 team members in more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2021 global revenues of €18 billion.

Get the Future You Want | www.capgemini.com

Get in touch

Yannick Martel
Data Analytics lead for telco industry
yannick.martel@capgemini.com

Caroline Vateau
Director
Sustainable digital
caroline.vateau@capgemini.com

Subhankar Pal
Global Innovation Leader for
Intelligent Networks
subhankar.a.pal@capgemini.com

Industry Marketing_22022023_MK