Green data center networks: a holistic survey and design guidelines

Emna Baccour†, Sebti Foufou‡§, Ridha Hamila†, Aiman Erbad†
† College of Engineering, Qatar University
‡CS department, New York University, Abu Dhabi, UAE
§LE2i Lab, University of Burgundy Dijon, France
{ebaccour, hamila, aerbad}@qu.edu.qa, sfoufou@u-bourgogne.fr

***Abstract*—Data Center Networks (DCNs) are attracting immense interest from the industry, research and academia to keep pace with the increase of Internet services demands. One of the major concerns that draws the attention of researchers is the exponential growth of the energy consumption and carbon emission of the DCNs. Studies conducted to identify the causes of the increasing energy consumption have proved that the growing size of computing demand, the over-provisioning of the networking resources, the under-utilization of the infrastructure, the fault-tolerance, the high bandwidth exigence and the inefficient hardware and cooling structure are leading to considerable energy waste. Therefore, in recent years, new data center (DC) architectures are proposed where new hardware types and new technologies are implemented for the sake of energy efficiency. Other efforts are focusing on designing algorithms and strategies to enhance the utilization of the network resources. Replacing *brown power* by *renewable energy* was also one of the attractive ideas to minimize the energy costs. In this survey paper, we will present energy-related problems in data centers and review the state of the art of the research literature on energy efficient architectures, techniques, technologies, resource management, and thermal control and monitoring. Additionally, we present the challenges facing each approach and the strategies to build a green DC. This paper serves as a specification document that shows step by step how to minimize the energy consumption of different components of the system.**

***Index Terms*—Data Center Networks, energy efficiency, renewable energy, cooling management, network architectures, greening software.**

I. Introduction

In the few recent years, data center networks have witnessed an unprecedented growth [1]. This growing importance of data centers is caused by the emergence of IT operations worldwide that replaced the traditional business models. To keep pace with the unceasing demand in services, considerable research efforts are conducted to design a performant interconnection network with a cost-effective deployment and maintenance. Generally, the design goals of data center networks are: high scalability, simple installation, small number of wires, an efficient routing algorithm that must be itself scalable, low latency, high bandwidth, sufficient runtime performance by installing the highest performance equipment to be always ready for peak traffic times, and good thermal control by managing the cooling infrastructure and avoiding any over-heating problems. In addition, with data availability at stake, the robustness of the data center network becomes more critical than ever. Hence, storage machines and network devices are duplicated to protect the loss of clients data.

Based on these critical requirements, many efforts are trying to build high performance data centers that include all the QoS (Quality of Service) needs including PTNet [2] and LaCoDa [3]. However, implementing a huge number of devices for the sake of scalability, adding redundant machines to face failures, installing high capacity equipment to provide a 24/7 availability and investing in a power hungry cooling infrastructure contribute to make the data center networks one of the largest consumers of energy in the world. In fact, statistics in [4] predict that the world power consumption of data centers will rise from 1.1%-1.5% in 2011 to 8% in 2020 due to over-provisioning. This growth of electrical power consumption arouses many concerns to investigate the impact on the environment. Studies found that the IT sector contributes approximately in 2% of the total green-house gas emissions where 37% of this gas is caused by networking equipment [5]. Motivated by this high energy consumption and effects on the environment, many researchers, during the past years, have focused on the design of green data center networks. However, even if significant progress has been made, there is still a large opportunity to save energy since existing approaches generally handle a specific source of power waste and do not take a holistic approach to minimize the energy consumption of all parts of a DC infrastructure.

This survey presents a specification guide for business owners to build a green network or to implement new measures in order to minimize the energy consumption. The paper shows step by step the best practices to follow. For each step, an overview of the state-of-the-art research is presented. The different steps introduced in this paper are: (1) Renewable sources of energy, (2) Power-aware cooling, (3) Energy efficient equipment, (4) Energy efficient interconnections and (5) Power-aware routing algorithms and energy efficient software. We believe that our work differs from other efforts since it presents a guide to achieve the maximum of energy efficiency in DCNs and can be positioned to become a specification for data center industries and scientific research to build a green infrastructure.

The paper is organized as follows: Section II describes different problems that lead to energy waste in DCNs. Section III presents different steps to reduce the energy consumption of DCNs. Section IV provides a comprehensive discussion to help

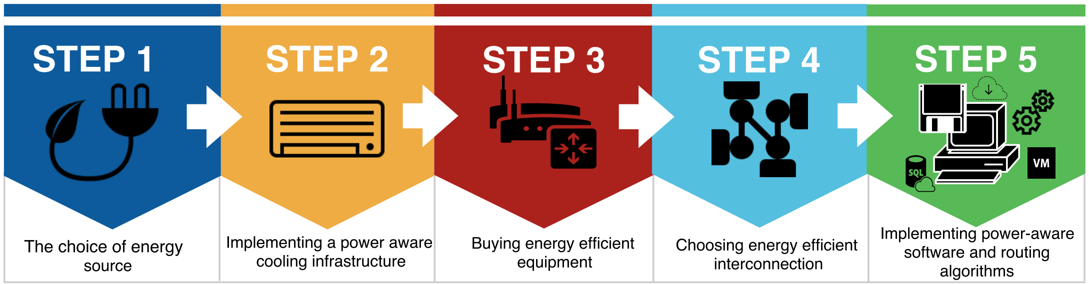

Fig. 1: Steps to green data center networks.

the business owners to adopt an effective strategy to green their networks adequately with their budget and QoS requirements. Section V concludes the survey paper.

II. DCN Problems Leading to Energy and Cost Inefficiency

Before conceiving a strategy to optimize the energy, we will understand, in this section, different obstacles that hold the network from being green.

A. Use of brown energy

Most of the energy consumed by DCs comes from coal burning, carbon-intensive sources, etc. Such energy produced from fossil-based-fuel is called *brown energy*. This type of energy costs the 24/7 available network billions of dollars annually. Unfortunately, the cost of energy is not the only problem for DC executors. In fact, consuming *brown energy* contributes in carbon emission and climate change as a single server can emit a carbon footprint similar to a sport vehicle. If continuing with the same rate of energy utilization, these billions of internet-connected devices could produce 3.5% of global emissions within ten years and 14% by 2040 [4]. This DC power consumption will double every four years [6].

B. Inefficient management of cooling infrastructure

A cost analysis described by *HP* and the Uptime Institute [7] showed that for most of data centers, 63% of the power is associated with cooling equipment. In fact, the compaction of the placement of computing hardware per unit area leads to a high cost related to the removal of heat dissipation. This problem is worsened by the adopted cooling approach. Indeed, the entire DC room is chilled by a constant temperature and the cooling machines are placed at fixed locations without taking into consideration the distribution of servers utilization and networking hardware that need special temperature.

C. Power-hungry equipment and architectures

The DC architecture is the interconnection between servers and network devices including switches and routers. This interconnection requires a deep study and an extreme consideration [8] [9] to guarantee the scalability, performance, and a good latency. Typical data centers use enterprise class equipment. These equipment are used to provide extremely stable and expandable environment. However, such hardware are very expensive and consume a large amount of energy [8].

D. Underutilized network

Servers utilization is one of the most important topics discussed for data center energy efficiency. In fact, studies conducted on DCNs traffics during different periods and times showed that the network utilization varies depending on the period of usage (holidays, working days, weekends, mornings, nights, etc.). Observations also showed that data centers operate only at 5% to 25% of their maximum capacities most of the time [10] [11]. This is explained by the fact that hundreds and even thousands of servers/switches in a data center, are powered on 24/7, waiting to receive or process data. In addition, data center operators plan the capacities of servers to be able to handle peak annual traffic such as Black Friday. For the rest of the time, the servers can stay unused (idle). This idle state is causing a huge waste of energy since a server consumes up to 70% of its peak power while being unused [10]. In addition, it has been noted that several interconnections implement a huge number of devices to guarantee scalability, fault-tolerance and high bandwidth, which contributes to the energy wasting.

III. Steps to Green Data Center Networks

In this section, we will describe different steps (summarized in Figure 1) that can be followed to solve the problem of power inefficiency. By considering these steps, a holistic strategy to deal with all sources of energy waste is applied.

A. Step 1: The choice of energy source

Choosing the right energy source is not only related to reducing the electricity bills but also to reducing the carbon footprint and waste landfills. In this context, industry and research gave a great interest to study the implementation of green sources of energy (solar, water or wind energy) that replace the *brown energy*.

1) Big companies and green sources of energy: Several big companies are motivated by the use of green sources of energy in their data centers including *Google* and *Microsoft*. For example, the goal of *Google* DCs was to achieve a *carbon free* energy. This means to empower all users with a cheap and clean source of energy. In this context, *Google* invested $2.5 billion in solar and wind projects to add green power to its grid, which made it the biggest purchaser of *renewable energy* in the world and contributed to reduce its energy use by 50%

compared to typical data centers. Google is now witnessing the results of its efforts by enlarging its use of *renewable energy* from 48% in 2015 to 61% in 2016 and 100% in 2017 [12]. *Microsoft* [13] owns also many data centers worldwide: USA, Europe and Asia. In order not to burden the world with more power consumption, *Microsoft* DCs are powered only by green resources of energy, since 2014. *Microsoft* is one of the biggest buyers of renewable energy in the world including buying 100 % of the energy produced by the Keechi project, purchasing 175 Megawatts of wind energy from the Pilot Hill Wind project and bringing 20 Megawatts of solar energy from Remington solar project [13].

2) Using Green energy efficiently: Using *renewable energy* contributes to reduce the electricity bills and the carbon emission. However, this energy depends essentially on the weather, which is not always favorable. In this case, the *brown energy* is indispensable. To maximize the use of clean energy and ensure that the system is continually operating, an approach to schedule the use of *renewable* and *brown energy* should be adopted. GreenSlot [14] is a job scheduler designed for DCNs to predict the quantity of solar energy to be available in the next days. Then, it schedules the execution of incoming jobs depending on the availability of the *renewable energy* and the cost of *brown energy*. If the *brown energy* should be used, the GreenSlot selects the time slot when the electricity is cheap. Otherwise, it delays the jobs to ensure that the green energy will be used as long as the performance deadlines are not violated. Greenware [15] described in Figure 2 is also a system that tries to maximize the use of green energy while respecting a certain budget: First, the system computes the budget of energy based on workload rate in previous periods. Then, based on the time-varying brown/green energy cost and the availability of *renewable energy* in worldwide DCs, GreenWare decides how to dispatch requests to different networks, having different energy prices, while respecting the QoS requirements. The Net-zero network [16] is another

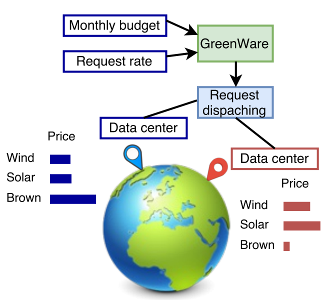

Fig. 2: GreenWare.

proposal to exploit green resources efficiently. This concept consists of producing an amount of *renewable energy* per day equal to the same amount of energy to be used. Net-zero consists of 4 modules: predicting the future electricity demand, provisioning resources for the incoming workloads to meet the Net-zero goal using optimization algorithms, deploying the scheduled planning, and verifying the efficiency of different modules to address any error in the next plan.

B. Step 2: Implementing a power-aware cooling infrastructure

As mentioned previously, more than 60% of energy is consumed by the cooling infrastructure and a huge budget is dedicated to the removal of heat to the outdoors. Many big companies are adopting free methods for cooling including *Google*. In fact, cooling its huge buildings requires 30% to 70% energy use overhead. Hence, *Google* decided to rely on free cooling mechanisms such as cooling with air or water rather than electrical chillers. For example, *Google* DC in Hamina, Finland has a 450 meter (1500 feet) long underground tunnel running into the Baltic sea used for the cooling [12]. Authors in [17] have given the cooling issue a high priority by conducting multiple researches and testbeds. One of these innovations is the intelligent sensors placement for smart cooling. This approach proposes to add sensors that monitor the temperature in different parts of the data center rooms. the control system continually manages the state of the cooling system depending on the workload and the active machines. It also provides the amount of cooling where and when it is needed, which reduces the energy consumption according to the workload and the infrastructure. Another reason leading to energy waste in DCNs is the unexpected changes of the workload and the pick traffic load that happens sometimes and leads to overheat. To address this problem, Authors in [18] derived a temperature control strategy for DCNs environment. This strategy combines an *air-flow control* that is responsible for long term cooling decisions and a *thermal-aware scheduler* that controls the short-term temperature fluctuations caused by unpredictable workloads. These two controllers receive the information concerning the temperature in each rack, and assume that workloads can be allocated to any non-busy server. In this way, based on the temperature data, high workloads are not assigned to the same racks and unexpected temperature changes cannot occur as seen in Figure 3.

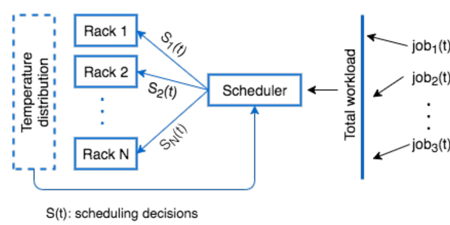

Fig. 3: Thermal-aware approach.

C. Step 3: Buying energy efficient equipment

To host a large number of servers accommodating a diversity of services without facing the problems of over-energy-consumption, energy efficient equipment should replace the energy hungry machines used in the typical data centers. The following types of technologies are tested in data centers and have proved to reduce the energy gain in the DCN: (1) Commodity network, (2) Wireless technology, (3) Optical technology. Hence, the next step to minimize the energy bills is to wisely choose the implemented hardware. In this section, we will describe some data center networks that use the aforementioned technologies.

1) The use of commodity network devices: The Three-tier data center is widely used in big data center companies [9]. Power-hungry enterprise-level switches are used in this kind of hierarchy because of their stability, reliability and security [8]. In response to this issue, several research efforts have been conducted to implement energy efficient devices (commodity network devices) instead of enterprise class hardware. Using commodity network switches contributes to reduce the cost, the energy consumption and the heat dissipation in the DC networks. Among these architectures, we can cite PTNet [19] and LaCoDa [3].

2) Introduction of wireless technology: Energy saving in data centers can be realized by implementing wireless technology, since it has been proved that the energy consumption of 60 GHz transceivers is minor [20]. In fact, the 60 GHz transceiver consumes up to 0.3 Watts [21] while the power consumption of a switch port is equal to 12 Watts and the power consumption of a server port is equal to 4 Watts [22]. In this context, many efforts are focusing on building wireless DCs such as Cayley [21] and WFlatnet [23]. In addition to reducing the energy by using the wireless transceivers, wireless sensors technology [24] is also a way to help DC operators to monitor their network, detect the energy dissipation and act accordingly. In summary, wireless communications can revolutionize the data center construction. However, a completely wireless DCN cannot achieve alone all the QoS requirements. In fact, the maximum data rate of wireless links is 7GBps and can be reduced by interference, however, the speed of Ethernet links is increasing and becoming more performant.

3) The use of optical technology: Introducing optical DCNs can offer high bandwidth, high port density and low energy consumption compared to traditional electrical networks. Several studies reported that the optical networks can improve the energy efficiency of a data center by 15% to 20% per year [25] and proved that a complete optical network can contribute to gain 75% of energy consumption. In fact, it is estimated that the optical transmitter can reduce the energy consumption from 20 nJ/b consumed by an electrical router and 10 nJ/b consumed by an Ethernet switch to around 0.5 nJ/b used by optical technology [25]. Several optical networks are designed for DCs to minimize the power consumption including Helios [26] and WaveCube [27].

D. Step 4: Choosing energy efficient interconnection

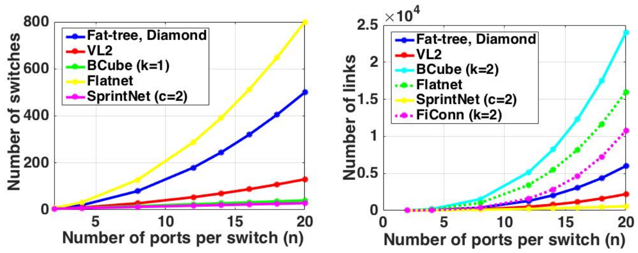

(a) (b)

Fig. 4: Comparison of number of network components.

Data center networks have different architectures and different structures. Each structure has its characteristics and its energy consumption depending on different parameters. Many efforts, such as [28], have tried to compare the power consumption in different DC topologies and concluded that the number of servers, switches and links used to build the network contributes to the power saving process. Figure 4 shows the number of implemented switches and links in different interconnection networks. For example, Fat-tree and Flatnet implement a larger number of switches compared to VL2, Diamond and SprintNet, which makes them power hungrier architectures. The description of different architectures can be found in the survey [9]. In the same context, other works aimed to implement the minimum of equipment by designing switchless networks, which interconnect servers directly without intermediate switches. This type of topology can save the energy consumed by switches, routers, racks and associated cooling machines. Among switchless DCs, we can cite the torus FleCube [29] and NovaCube [30].

Another fact that contributes to the power saving is the network interconnection, which determines the path length between communicating servers and impacts the number of operations processed by the network. A higher number of hops means that the packets pass by a higher number of intermediate switches and servers; meaning also a higher number of operations are accomplished by network devices (switches, servers and links) and a higher energy consumption. Figure 5(a) compares the number of operations accomplished by nodes in different networks (FlatNet, DCell, BCube and PTNet) having approximately the same size. We can see that Flatnet and DCell execute more operations because of their large APL (average path length).

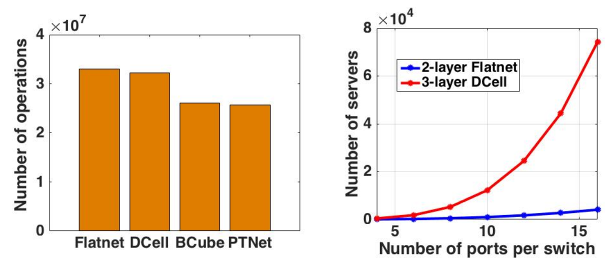

(a) (b)

Fig. 5: Interconnection efficiency.

Another problem that affects the power efficiency of the network is the high scalability of a data center. In fact, some architectures are rapidly scaling. It means when augmenting the size of a switch (enlarging the number of its ports), the size of the network increases dramatically. As an example, Figure 5(b) shows the scalability of DCell and Flatnet topologies. We can note that DCell is rapidly scaling. Supposing that we need only 10 thousand servers, 10-port switches must be implemented. Here, we obtain 12,210 servers which means we will get more than 2 thousands unneeded servers consuming extra energy. Hence, conceiving a gradually scalable topology is important to minimize the cost and the energy consumption of the network. Authors in [2] proposed a parameterizable topology called PTNet where the scalability level is determined by a parameter s. To summarize, when conceiving the data center interconnection, three parameters should be considered to minimize the energy: (1) the number of implemented devices, (2) the average path length and (3) the gradual scalability.

E. Step 5: Implementing power-aware software and routing algorithms

As mentioned in section II-D, the non-proportionality between the traffic load and energy consumption is a serious issue that faces the attempts to green the network. Many efforts stated that the traffic can be satisfied by only a subset of active components, and, the idle devices should be switched to a power-aware status. As an example, Figure 6 shows the increasing of servers load. The red nodes are the devices that receive or send data while the green nodes are the idle devices and can be set into an energy aware status. In this way, the energy proportionality can be achieved. Restricting the network to a subset of active devices and reducing the power consumption of the idle ones is called *dynamic power management*. Generally, four strategies are adopted for the dynamic power management: (1) Power-aware routing algorithms, (2) Dynamic voltage/frequency scaling, (3) Adaptive Link Rate and (4) Virtualization.

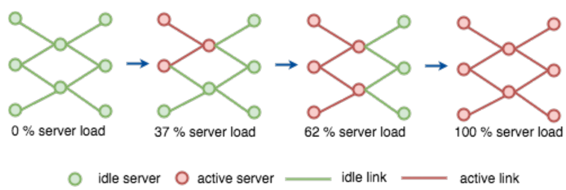

Fig. 6: Idle nodes.

1) Power-aware routing algorithms: The traditional routing algorithms (e.g. shortest path routing) contribute to worsen the power consumption issue inside a data center since they are not aware of the workload distribution and the utilization of different devices. Hence, designing a power-aware routing is an important step towards a green network. One of the most known power-aware routing algorithm is ElasticTree [31]. ElasticTree targets to find the minimum subset of the network which should be kept active and the set of network devices that are unused and can be shut down. The ElasticTree consists of three modules run in real-time: (1) The optimizer module: it defines the devices contributing in the traffic matrix; (2) The routing module: it calculates the shortest packets routes; (3) The power control module: it is responsible for adjusting the state of devices (on, off). Another power-aware algorithm is the vital nodes approach [10] which suggests not to calculate the best routing paths in real-time when receiving the traffic pattern. Instead, the network is abstracted to a graph and vital nodes in this graph are calculated once, using different methods (betweenness, closeness, degree, etc). At a given time t, when the traffic matrix is received, only the vital nodes are kept active. These approaches proved to reduce more than 20 % of energy consumption.

2) Dynamic voltage/frequency scaling (DVFS): This approach focuses on lowering the speed or the frequency of the networking devices to minimize their energy consumption. In this way, dynamic scaling allows to achieve the proportionality between power consumption and the supply voltage/frequency. The authors in [32] propose to reduce the energy consumption using the CPU DVFS and an energy aware task scheduling. The idea is to consolidate the tasks into a minimum number of processors and run these processors at their optimal frequencies. Thus, a large number of devices will stay idle and can be switched to a lower state. Two steps are proposed: (1) The energy aware task allocation: this step manages the tasks-to-processors allocation that achieves the minimum of energy consumption while respecting the QoS requirements; (2) Local task migration: When a task is completed in the allocated processor, two measures can be taken: either the frequency will be adjusted or tasks from other processors migrate to this available processor.

3) Adaptive Link Rate (ALR): Measurements in [33] showed that 10 Gbps Ethernet link consumes from 10 to 20 Watts more energy than 100 Mbps link. Also, experiments confirmed that the idle and fully utilized links consume approximately the same energy. Hence, these observations present an opportunity to save an important amount of energy by lowering the data rate during low traffic loads. The ALR approach intends to use different link data rates (10 Mbps, 100 Mbps, 1 Gbps, and 10 Gbps) by switching between them. Authors in [34] propose an ALR mechanism (presented in Figure 7). The idea is to attribute the responsibility of detecting the need to increase or decrease the data rate to the end of each link (ports). In case of change need, the end of the link sends an ALR REQUEST (goto high or goto low) to the link receiving partner. The link receiving partner can acknowledge the reception with either ALR ACK (agree) or ALR NACK (refuse) depending on the required QoS. Two thresholds are introduced; qLow and qHigh. If the port buffer queue length exceeds qHigh, the data rate should be upgraded. If the buffer decreases under qLow, the data rate can be lowered if the two sides agree.

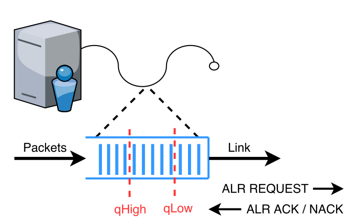

Fig. 7: The adaptive link rate approach.

4) Optimizing energy using Virtualization: Several years ago, virtualization was used as an environment to test servers. However, recently, researchers have given the network virtualization great attention since it has been an adapted technique in data centers. The virtualization technology is based on creating multiple virtual machine (VM) instances on a single physical server [35]. Since the under-utilization of the network and the non-proportionality between the traffic load and the energy consumption are the major causes of energy waste, virtualization can be an important technique to consolidate the traffic load, improve the resource utilization and achieve energy efficiency. The general idea is to transfer virtual machines (migration) and map them to a set of servers (allocation), then, assign to them the resources and services requested by the clients. In this way, the workload is consolidated in a set of servers (consolidation) to maximize their utilization and the idle machines can be shut down, which leads to a great energy saving. Figure 8 presents the virtualization concept. An optimal VM management (migration, allocation and consolidation) is needed to accomplish the best performance and power saving.

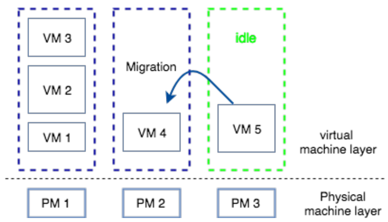

Fig. 8: Virtualization concept.

Several efforts are conducted to optimally solve many key challenges: (1) Where to migrate the VMs to maximize the utilization of devices while minimizing the migration time overhead; (2) How to monitor the network to respect the Service Level Agreement (SLA) fixed by the client while saving the maximum of energy; (3) How to ensure the independence of the virtual machines from their physical machine. Among these efforts, we can cite [36] and [37].

IV. Summary and Discussion

To effectively and comprehensively address the energy problem in huge data centers, we should adopt a holistic approach that takes into consideration all the problems causing energy inefficiency in a DC environment. It means, to green a data center, one or a combination of the following approaches should be considered:

- Clean energy sources: this can be accomplished by reducing the energy cost and impact on the environment using *renewable sources of energy*. This approach faces several challenges including the location of the data center, the randomness of the weather, the space for energy farms, the availability of a budget to build these farms and the possibility of buying clean energy.
- Clean cooling and heat disposal: by reducing the cost of cooling which is burdening the electricity bills of DCNs. In addition to the location of the DC, this approach depends on the space in the building to implement equipment that use the outside sources to cool the inside infrastructure. Then, a budget should be planned for the new infrastructure or the heat sensors.
- Efficient equipment and interconnection: by implementing energy efficient components, servers and switches, energy consumption can be reduced. However, this approach depends on the applications and operations hosted in the DC since some of them need high-performance machines. These machines are not necessarily energy friendly. Also, some DCs are being built with containerized architecture, which constrains the implementation of several energy-efficient topologies such as wireless/optical networks.
- Power-aware algorithms: accomplished by reducing the under-utilization problem and aligning the energy consumption with the traffic load. This approach depends on the QoS and reliability requirements of the applications. In fact, unlike the hardware solutions that work on fine time granularities (milliseconds to seconds), software solutions operate at coarser granularities (seconds to hours) [38]. Therefore, the energy efficiency is inversely proportional with the optimal performance. In other words, power-aware algorithms add time overhead to calculate idle nodes or migrate VMs, which cannot always be implemented in real-time. Hence, data center owners should evaluate their business needs and study if their applications truly require 24/7 accessibility, and then they can sacrifice wisely in terms of performance in order to be green.

Implementing all these approaches is every business owner's dream, however, this is not evident. Business owners need to evaluate their needs and outline their (1) aims, objectives and constraints, (2) budgets, (3) scope and time granularities, (4) customer's QoS requirements, (5) available equipment and possibility to enhance, change or add new hardware, (6) tolerance of devices for the frequency/voltage scaling or to be switched to on/off status (6) global and local resource management, (7) plans and possible actions and (8) monitors to observe the progress and achievements.

After highlighting these points, two strategies can be adopted:

- Incremental strategies: In this strategy, the enterprise already has its infrastructure and cannot invest a big budget to restructure its DC. So, it only adds some simple measures to achieve moderate green goals. These measures can include switching off idle devices, optimizing the DC temperature, voltage/frequency scaling, etc. The aforementioned approaches are easy to implement and have negligible costs.
- Strategic strategies: In this strategy, the enterprise conducts an audit of its infrastructure, buildings, location and weather and then develops a greening plan addressing several aspects. For example, new networking devices and technologies can be implemented, clean cooling systems can be built and new interconnections can be adopted. In addition, deeper measures can be used such as carbon filters, refurbishing and reusing old computers, etc.

V. Conclusion

The design of green data center networks has gathered the attention of industry and academia because of its environmental and financial impact. In this context, extensive research has been conducted to minimize the power consumption and carbon emission of the network infrastructure. This survey presents a specification guide to build an energy efficient data center network. In fact, the paper describes significant insights about various problems and aspects causing the huge power waste when building a traditional data center network. Then, different steps to build a green network are proposed starting from choosing the right source of energy, implementing and managing the cooling infrastructure, choosing the network interconnection and the power-aware hardware and technologies arriving to the software part where the routing algorithm, the virtual machines and power-aware schemes should be well studied. Finally, we summarized and introduced the required considerations before choosing the right greening strategy.

ACKNOWLEDGMENT

This publication was made possible by NPRP grant 8-627-2-260 and NPRP grant 6-718-2-298 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors.

REFERENCES

[1] J. Zhang, F. R. Yu, S. Wang, T. Huang, Z. Liu, and Y. Liu, "Load balancing in data center networks: A survey," *IEEE Communications Surveys Tutorials*, vol. 20, no. 3, pp. 2324–2352, 2018.

[2] E. Baccour, S. Foufou, and R. Hamila, "Ptnet: A parameterizable data center network," in *2016 IEEE Wireless Communications and Networking Conference*, April 2016, pp. 1–6.

[3] Z. Chkirbene, S. Foufou, R. Hamila, Z. Tari, and A. Y. Zomaya, "Lacoda: Layered connected topology for massive data centers," *Journal of Network and Computer Applications*, vol. 83, pp. 169 – 180, 2017.

[4] (2017, December) Tsunami of data could consume one fifth of global electricity by 2025. [Online]. Available: https://www.climatechangenews.com/2017/12/11/tsunami-data-consume-one-fifth-global-electricity-2025/

[5] K. Bilal, S. U. R. Malik, O. Khalid, A. Hameed, E. Alvarez, V. Wijaysekara, R. Irfan, S. Shrestha, D. Dwivedy, M. Ali, U. S. Khan, A. Abbas, N. Jalil, and S. U. Khan, "A taxonomy and survey on green data center networks," *FGCS*, vol. 36, pp. 189 – 208, 2014.

[6] (2017, December) Why energy is a big and rapidly growing problem for data centers. [Online]. Available: https://www.forbes.com/sites/forbestechcouncil/2017/12/15/why-energy-is-a-big-and-rapidly-growing-problem-for-data-centers

[7] L. Stapleton. (2016) Getting smart about data center cooling. [Online]. Available: http://www.hpl.hp.com/news/2006/oct-dec/power.html

[8] K. Bilal, S. U. Khan, L. Zhang, H. Li, K. Hayat, S. A. Madani, N. Min-Allah, L. Wang, D. Chen, M. I. Iqbal, C.-Z. Xu, and A. Y. Zomaya, "Quantitative comparisons of the state-of-the-art data center architectures." *Concurrency and Computation: Practice and Experience*, vol. 25, pp. 1771–1783, 2013.

[9] T. Chen, X. Gao, and G. Chen, "The features, hardware, and architectures of data center networks: A survey," *Journal of Parallel and Distributed Computing*, vol. 96, pp. 45 – 74, 2016.

[10] E. Baccour, S. Foufou, R. Hamila, and Z. Tari, "Achieving energy efficiency in data centers with a performance-guaranteed power aware routing," *Computer Communications*, vol. 109, pp. 131 – 145, 2017.

[11] M. Dayarathna, Y. Wen, and R. Fan, "Data center energy consumption modeling: A survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 1, pp. 732–794, Firstquarter 2016.

[12] "Google environmental report," Tech. Rep., 2018. [Online]. Available: https://storage.googleapis.com/gweb-sustainability.appspot.com/pdf/Google_2018-Environmental-Report.pdf

[13] (2019) Microsoft. [Online]. Available: https://www.microsoft.com/about/csr/environment/

[14] I. Goiri, R. Beauchea, K. Le, T. D. Nguyen, M. E. Haque, J. Guitart, J. Torres, and R. Bianchini, "Greenslot: Scheduling energy consumption in green datacenters," in *International Conference for High Performance Computing, Networking, Storage and Analysis*, Nov 2011, pp. 1–11.

[15] Y. Zhang, Y. Wang, and X. Wang, *GreenWare: Greening Cloud-Scale Data Centers to Maximize the Use of Renewable Energy*. Springer Berlin Heidelberg, 2011, pp. 143–164.

[16] M. Arlitt, C. Bash, S. Blagodurov, Y. Chen, T. Christian, D. Gmach, C. Hyser, N. Kumari, Z. Liu, M. Marwah, A. McReynolds, C. Patel, A. Shah, Z. Wang, and R. Zhou, "Towards the design and operation of net-zero energy data centers," in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2012, pp. 552–561.

[17] X. Wang, X. Wang, G. Xing, J. Chen, C. Lin, and Y. Chen, "Intelligent sensor placement for hot server detection in data centers," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 8, pp. 1577–1588, Aug 2013.

[18] N. Vasic, T. Scherer, and W. Schott, "Thermal-aware workload scheduling for energy efficient data centers," in *Proceedings of the 7th International Conference on Autonomic Computing*. ACM, 2010, pp. 169–174.

[19] E. Baccour, S. Foufou, R. Hamila, Z. Tari, and A. Zomaya, "Ptnet: An efficient and green data center network," *Journal of Parallel and Distributed Computing*, vol. 107, pp. 3 – 18, 2017.

[20] E. Baccour, S. Foufou, R. Hamila, and M. Hamdi, "A survey of wireless data center networks," in *2015 49th Annual Conference on Information Sciences and Systems (CISS)*, 2015, pp. 1–6.

[21] J.-Y. Shin, E. G. Sirer, H. Weatherspoon, and D. Kirovski, "On the feasibility of completely wireless datacenters," in *Proceedings of the Eighth ACM/IEEE Symposium on Architectures for Networking and Communications Systems*. ACM, 2012, pp. 3–14.

[22] L. Popa, S. Ratnasamy, G. Iannaccone, A. Krishnamurthy, and I. Stoica, "A cost comparison of datacenter network architectures," in *Proceedings of the 6th International Conference*. ACM, 2010, pp. 16:1–16:12.

[23] E. Baccour, S. Foufou, R. Hamila, and M. Hamdi, "wflatnet: Introducing wireless in flatnet data center network," in *2015 IEEE Globecom Workshops (GC Wkshps)*, 2015, pp. 1–6.

[24] T. Akiyama, M. Matsuoka, K. Matsuda, Y. Sakemi, and H. Kojima, "Secure and long-lived wireless sensor network for data center monitoring," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, July 2018, pp. 559–564.

[25] R.S.Tucker, "Green optical communications part ii: Energy limitations in networks," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 17, no. 2, pp. 261–274, 2011.

[26] N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, "Helios: a hybrid electrical/optical switch architecture for modular data centers," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, 2010.

[27] K. Chen, X. Wen, X. Ma, Y. Chen, Y. Xia, C. Hu, Q. Dong, and Y. Liu, "Toward a scalable, fault-tolerant, high-performance optical data center architecture," *IEEE/ACM Transactions on Networking*, vol. 25, no. 4, pp. 2281–2294, Aug 2017.

[28] L. Gyarmati and T. A. Trinh, "How can architecture help to reduce energy consumption in data center networking?" in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. ACM, 2010, pp. 183–186.

[29] D. Li, Y. Shen, and K. Li, "Flecube: A flexibly-connected architecture of data center networks on multi-port servers," *Computer Communications*, vol. 77, pp. 62 – 71, 2016.

[30] T. Wang, Z. Su, Y. Xia, B. Qin, and M. Hamdi, "Novacube: A low latency torus-based network architecture for data centers," in *2014 IEEE Global Communications Conference*, 2014, pp. 2252–2257.

[31] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "Elastictree: Saving energy in data center networks," in *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, 2010, pp. 17–17.

[32] S. Wang, Z. Qian, J. Yuan, and I. You, "A dvfs based energy-efficient tasks scheduling in a data center," *IEEE Access*, vol. 5, 2017.

[33] H. Yan, C. Gueguen, B. Cousin, J. P. Vuichard, and G. Mardon, *Green Networking and Communications : ICT for Sustainability*. Shafiullah Khan, Jaime Lloret Mauri Eds. Publisher: CRC Press, USA, 2013.

[34] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption of ethernet with adaptive link rate (alr)," *IEEE Transactions on Computers*, vol. 57, no. 4, pp. 448–461, 2008.

[35] A. Varasteh and M. Goudarzi, "Server consolidation techniques in virtualized data centers: A survey," *IEEE Systems Journal*, vol. 11, no. 2, pp. 772–783, June 2017.

[36] T. M. Nam, N. H. Thanh, H. T. Hieu, N. T. Manh, N. V. Huynh, and H. D. Tuan, "Joint network embedding and server consolidation for energyefficient dynamic data center virtualization," *Computer Networks*, vol. 125, pp. 76 – 89, 2017.

[37] A. Marotta, S. Avallone, and A. Kassler, "A joint power efficient server and network consolidation approach for virtualized data centers," *Computer Networks*, vol. 130, pp. 65 – 80, 2018.

[38] W.-c. Feng, *The Green Computing Book: Tackling Energy Efficiency at Large Scale*, 1st ed. CRC Press, Inc., 2014.