Energy-Aware Data Centre Management

Peoples, C., Parr, G., & McClean, S. (2011). Energy-Aware Data Centre Management. In *Unknown Host Publication* IEEE. <https://doi.org/10.1109/NCC.2011.5734700>

[Link to publication record in Ulster University Research Portal](#)

Published in:
Unknown Host Publication

Publication Status:
Published (in print/issue): 28/01/2011

DOI:
[10.1109/NCC.2011.5734700](https://doi.org/10.1109/NCC.2011.5734700)

Document Version
Publisher's PDF, also known as Version of record

General rights
Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

Energy-Aware Data Centre Management

C. Peoples, G. Parr and S. McClean

India-UK Advanced Technology Centre of Excellence in Next Generation Networks, Systems and Services

Faculty of Computing and Engineering, University of Ulster

Coleraine, Northern Ireland, BT52 1SA, United Kingdom

{c.peoples, g.parr, si.mcclean}@ulster.ac.uk

Abstract—Cloud computing is one way in which communications within and between data centres can be optimised by using resources which are physically close to the client, are exposed to lower electricity costs, contribute a smaller carbon footprint or have residual resources sufficient to fulfil Quality of Service requirements. Optimisation of activity involving data centres is a next generation network management objective due to continued growth in the number of plants and volume of operations within, factors which contribute to environmental concerns associated with energy consumption and carbon emissions from data centre facilities when renewable energy resources are not used. In this paper, we present an algorithmic mechanism developed to automate selection of a data centre in response to application requests, the Data Centre (DC) Energy-Efficient Context-Aware Broker (e-CAB). Through integration of the DCE-CAB in a case study scenario, operational improvement through reduction of carbon emission and balancing of other performance-related attributes including delay and financial cost is achieved, validating the DCE-CAB's positive impact.

Keywords—autonomy; context-awareness; data centre; quality of service; green networking; policy-based management.

I. INTRODUCTION

Cloud computing is one way in which communications within and between data centres can be optimised by using resources which are physically close to clients, are exposed to lower electricity costs, contribute a smaller carbon footprint or have residual resources sufficient to fulfil Quality of Service (QoS) requirements. Optimisation of data centre activity is a next generation network management objective due to continued growth in the number of plants and volume of operations within, factors which contribute to environmental concerns associated with carbon emissions from data centre facilities when non-renewable sources of energy are used to power them [1]. Cloud computing can benefit data centres without physical resources by enabling capacity leasing from solutions such as the IBM Cloud or Windows Azure platform. These can encourage companies without dedicated data centre resources to avoid their construction, restricting emissions associated with plant management. There are organisations however, with static data centre resources across the world, many driven to

expand capacity to keep up with growing demand. Recent growth rates, in parallel with impending environmental concerns, therefore drive recognition of opportunities to improve data centre operational efficiency [2] [3].

In this paper, a mechanism which automates selection of a data centre in response to application request and real-time network operation is presented, the *Data Centre (DC) Energy-Efficient Context-Aware Broker (e-CAB)*. When a real-time response is not prioritised by the application, carbon and financial costs associated with data centre communication and operation are included in the decision-making process. This takes into account environmental and operational sustainability concerns, a fact attracting increasing attention due to rising electricity costs [4] [5]. The overall objective of the DCE-CAB is to provision *energy-awareness* (step 1 in Fig. 1) in DC network operations and achieve *energy-efficiency* (step 2 in Fig. 1). In response to awareness of resource availability (e.g., bandwidth), network demand (e.g., transmission volume) and operational characteristics (e.g., bit error and packet loss rate), the DCE-CAB performs *dynamic management* to reach a scenario where energy-efficiency is achieved such that resource availability exceeds demand when task waiting time, carbon cost and/or electricity cost is optimised (objectives defined on a transmission-specific basis). In the DCE-CAB's development, our research objectives therefore

Figure 1. Energy-awareness driving efficiency of DC operations

include: (1) optimising carbon emissions during communications with DCs; (2) selecting DC facilities based on their carbon footprint; (3) selecting DC resources which achieve application QoS requirements; and (4) selecting DC facilities with lowest operating cost. We believe this to be a unique approach as it manages DC communications on the end-to-end path and does not optimise operational energy-efficiency within DCs only, a common strategy in related literature (e.g., [6]-[9]). The e-CAB optimises DC selection in response to application characteristics, DC and path from the device. Results in this paper show ability to achieve a balance between attributes which include electrical cost, latency, carbon emissions and number of path nodes.

The remainder of the paper continues as follows: in Section II, the research background is presented in more detail. This sets the scene for presentation of the DCe-CAB in Section III, including discussion of its core components and context data used in its optimisation algorithm. Operational characteristics of a test scenario are presented in Section IV, along with demonstration of the DCe-CAB's effectiveness using a selection of optimisation approaches. The conclusion and future work is presented in Section V.

II. RESEARCH BACKGROUND

Related literature relevant to this research applies energy monitoring and management of DC resources and operation. Bohra and Chaudhary (2010) propose VMeter which profiles power and other resources consumed by virtual machines hosted on physical nodes through monitoring CPU, cache, disk and DRAM, and has ability to predict real-time energy consumption [6]. This aims to improve resource utilisation according to power-aware policies. Hwang et al. (2010) define an energy utilisation model in order to optimise air and liquid cooling thermal management in DCs [7]. They review energy consumption incurred during operation of contrasting schemes in relation to other environment characteristics, such as chiller temperature and water flow rate. The overall objective in [7] is to apply DC cooling such that power overhead costs are optimised. Beloglazov and Buyya (2010) structure their approach in response to energy-efficiencies of operating virtualised cloud DCs [8]. They capitalise on energy-efficiencies from consolidating virtual machine operation in contrast to optimising utilisation across machines. Decisions are based on connections between virtual machines across the cloud and thermal state of each to optimise the energy-efficiency achieved. Tozer and Salim (2010) describe an approach to evaluate the effectiveness of air flow in the DC and perform cooling in both raised floor and non-raised floor buildings [9]. Air conditioning efficiency is reviewed using metrics which include negative pressure flow rate, bypass flow rate, recirculation flow rate, and balance of the Computer Room Air Conditioning Unit and server design flow rates. Their objective is to allow air management in the DC to be understood and enable efficiency to be improved.

As observed from these examples, a range of DC energy management schemes restrict optimisation to within the DC. There are opportunities however, to also take into account performance on end-to-end paths between client devices and DCs. This involves selecting the DC which operates most efficiently, and include performance trends on the paths leading to the DC in the selection process. It is with this objective that our research is pursued.

III. RESEARCH APPROACH: THE DCe-CAB

The DCe-CAB autonomically manages operational and communication energy-efficiency within and between data centres and client devices by applying dynamic responses to environment context which indicate performance achieved in terms of queuing delay, bandwidth availability, bit error and packet loss rate. Its primary objective is to distribute load from clients to DCs intelligently, with decisions made as a function of application QoS requirements and a drive to minimise the carbon and financial costs of DC operation. The DCe-CAB architecture includes an *orchestration agent*, and one or more *data centre* and *application agents* (Fig. 2). *Application agents* reside on all operational devices which communicate with data centres and capture QoS requirements and application characteristics (step 1a in Fig. 2). Attributes collected include acceptable request/response latency, transmission volume and required data rate, selected to influence decisions made by the orchestration agent which enable application QoS, while impacting on energy consumption, to be achieved. Context is also collected on the application side from operational devices sending traffic, including residual battery capacity, energy cost per packet and power source, attributes which indicate the importance of applying energy conserving techniques to a transmission and define device ability to achieve application QoS. *Data centre agents* reside in each DC under the orchestration agent's control, collecting context in parallel with the application agent (step 1b in Fig. 2) on the number of servers and racks in the DC, number of disks in each rack, available bandwidth within the data centre and job completion rate. Collected context attributes include those which define the DC's ability to achieve the QoS requirements associated with a transmission and its energy consumption characteristics. The *orchestration agent*

Figure 2. Flow of DCe-CAB control between client and data centre

is a repository where context collected from application and DC is evaluated. It is responsible for devices and DCs such as those shown in Fig. 3, their grouping determined on an organisation-specific basis. Transmission decisions originate from the orchestration agent using all context collected. Once evaluated, an intelligently-selected DC and path to the DC can be defined, over which communication from client devices (step 2) and data centres (step 3) will occur.

The e-CAB's overall objective is to provision Energy-Tolerant Networking (ETN). (The CAB was developed initially in response to reliability and sustainability challenges when communicating in Delay-Tolerant Networks (DTNs) [10] [11].) We develop solutions for ETN in a domain-specific approach, and this paper presents a version customised for the DC. The specification of context detail in the DCE-CAB architecture is important – a domain-specific range is required to allow energy-efficiency improvements, and simultaneously provision a solution with energy-awareness being in itself energy-efficient.

IV. DCE-CAB OPTIMISATION AND OBJECTIVES

The DCE-CAB evaluates context collected from applications, DCs and the environment through which traffic passes to optimise decisions made. A network scenario comprising client devices and DCs is represented in Fig. 3: a network $G=V(E)$ is composed of V $V:=V(G)$ nodes and E $E:=E(G)$ links. b is a sub-link on a path p between device and DC. b,a is a path p within network G , and DCI is a DC of G . Application traffic may be sent directly to the DC which will service its request or it may be sent via a DC to another. Valid paths in this topology therefore include: i,d (8(client),5,4(DC)) and y,c,a,k,p (13,2,3,5,1,6). Using this detail, high-level requirements to select a DC by the DCE-CAB allow costs associated with energy emissions, electricity costs and delay, each measured in relation to the number of nodes on the path to a DC, to be optimised subject to application bandwidth, bit error and packet loss rate QoS requirements. In this paper, we consider the effectiveness of approaches which use these operational characteristics to drive performance. Path p to the selected DC server from the client device will ideally be one which minimises the number of nodes n , carbon emissions EM , queuing delay Q and cost of electricity el from operating nodes along the path and at the data centre. These objectives are represented in Eq. (1), and EM may be replaced with n , Q or el :

$$\underset{p \in P^{*(b,c;G)}}{\text{minimise}} \quad EM \quad (1)$$

$P^{*(b,c;G)}$ is the set of all sub-links (b,c) in network G between client device and data centre, and Eq. (1) takes into account the need to minimise EM , n , Q , and el on the path external to the DC in network G . This however, is a constrained optimisation problem where minimisation of all performance-influencing attributes will not be possible; instead, one attribute will be prioritised and others controlled to meet, at worst, a threshold. While achieving

Figure 3. Paths between client devices and data centres

this objective, it is also necessary to ensure that the network:

- a. Has bandwidth B available in relation to application B_a QoS requirements on all links within network G :

$$\sum_{p \in P^{*(b,c;G)}} B \geq B_a \quad b,c \in V \quad (2)$$

where bandwidth B on all sub-links (b,c) between nodes in the network external to the DC in network G is greater than or equal to application bandwidth requirements B_a .

- b. Has the lowest, or one which the application a can cope with, bit error rate E and packet loss ratio R on all links within network G (R may replace ER):

$$\sum_{p \in P^{*(b,c;G)}} ER \leq ER_a \quad b,c \in V \quad (3)$$

where bit error rate ER and packet loss ratio R on all sub-links $P^{*(b,c;G)}$ between nodes in network G external to the DC is less than or equal to the acceptable error rate ER_a and packet loss ratio R_a for application a .

In summary, defining the communication process between client device and DC requires optimisation of carbon emission EM , electricity cost el , queuing delay Q and number of nodes n subject to application QoS requirements. The DCE-CAB therefore uses optimisation algorithms within its orchestration agent to achieve these high-level objectives.

A. DCE-CAB Optimisation: Case Study

The DCE-CAB's ability to improve communication efficiency between client devices and DCs is presented using a case study in Section IV.A. The network in Fig. 3 is considered in terms of emissions from each DC on the path (Table I), electricity costs on links (Table II) and queuing delay along the path (Table III), context characteristics collected by the DCE-CAB data centre, application and orchestration agents. In the case study, there are instances

where client devices are connected directly to a DC and a range are connected to multiple DCs, providing varying test conditions to explore DCE-CAB optimisation efficiency.

Optimisation decisions taken by the DCE-CAB are driven by priorities which include:

1. *Path selection as a function of average electricity $\bar{el}$ costs along the path for each node n on sub-link $b, c \in V$ in network G and maximum electricity $el_{\max}$ cost on the end-to-end path between device and DC on path p (Eq. (4)):*

$$\left(\frac{\bar{el}}{p \in P^*(n;G)} \cdot \frac{el_{\max}}{p \in P^*(n;G)} \right) \left(\frac{\bar{el}}{el_{\max}} \right)_{p \in P^*(DC;G)} el_{\max} \quad (4)$$

Eq. (4) selects a path and data centre DC for which electricity costs in network G are minimised, helping DC operators to achieve sustainability in the face of growing volumes of operations and increasing electricity costs. The relationship between average and maximum electricity costs is weighted by average and maximum values to ensure that the path with lowest cost overall is selected (a path may have the lowest maximum cost, but have a higher average cost than others). The value is also weighted by cost in the end DC on the basis that it is here where full processing of application traffic will occur and is therefore one of the most important costs on the path. The lowest cost path is selected as optimum when Eq. (4) is applied.

The following paths from four client devices in Fig. 3 are

TABLE I. CARBON EMISSIONS FROM DATA CENTRES

Data Centre ID	Data Centre Carbon Emissions (KG Carbon Dioxide)
DC1	0.5
DC2	0.8
DC3	0.75
DC4	1
DC5	0.8
DC6	0.65

TABLE II. ELECTRICITY COST AT DATA CENTRES

Data Centre ID	Electricity Cost (pence per unit)
DC1	13.3
DC2	15
DC3	11
DC4	12.2
DC5	10
DC6	14

TABLE III. QUEUING DELAY ON PATHS BETWEEN DCs

Data Centre ID	Queuing Delay (seconds)
DC1	0.5
DC2	1
DC3	0.2
DC4	0.05
DC5	1
DC6	0.02

selected when Eq. (4) is applied:

path f (DC5) for device 7; **path j (DC5)** for device 9;
path z, c, a (DC2, DC3, DC5) for device 10;
path q, k (DC1, DC5) for device 12.

For device 7, Eq. (4) selects the shortest path with lowest electricity cost (10.0 pence per unit) in relation to costs of 15.0, 11.0, and 12.2 pence per unit on other one-hop paths. Similarly, for device 9 the lowest cost path (10.0 pence) is selected over those with higher costs (13.3 and 12.2 pence). One hop paths are not selected from devices 10 and 12 due to higher operating costs on them and prioritisation of cost over node number when Eq. (4) optimisation is applied. In the case of device 12, path q, k is selected, with the final choice of DC5. Average path cost is 11.65 pence which can be compared with the average cost of 13.3 (path q), 12.2 (path o), and 14.0 pence (path r) on other one-hop paths. These results therefore validate the DCE-CAB's ability to execute effective selection in relation to electricity cost at the end DC and at points along the path.

2. *Path selection as a function of average carbon emission $\overline{EM}$, queuing delay $\bar{Q}$ and electricity cost $\bar{el}$ between device and DC, weighted by number of nodes n (Eq. (5)):*

$$n \left(\left(\frac{\overline{EM}}{p \in P^*(n;G)} \right) \left(\frac{\bar{Q}}{p \in P^*(n;G)} \right) \left(\frac{\bar{el}}{p \in P^*(n;G)} \right) \right) \quad (5)$$

The relationship between average carbon emission, delay and electricity cost is weighted by the number of path nodes to optimise the opportunity that the least costly selections overall are made when these network characteristics are prioritised. The attributes are contextualised in relation to number of nodes on the path – Q may be higher, for example, when the path is longer (in terms of having more nodes), but with a lower average carbon and financial cost it should not be classified as unsuitable. The least cost path and DC is selected as optimum when Eq. (5) is applied.

Optimum paths according to Eq. (5) between devices and DCs are determined by the orchestration agent to be:

path f (DC5) for device 7; **path j (DC5)** for device 9;
path z, c, a (DC2, DC3, DC5) for device 10;
path q, k (DC1, DC5) for device 12.

Using this approach, it is possible to observe balance between network characteristics. From device 7, the one-hop route selected is exposed to lowest electricity cost (10.0 pence), maximum delay (1.0 second), and average emissions (0.8 KG), achieving a balance between attributes. From device 10, a two-hop path is selected when Eq. (5) is applied, and balance between context attributes is achieved through selection of the DC with average carbon emissions (0.8 KG), lowest electricity cost (10.0 pence) and highest delay (1.0 second). The results of this optimisation therefore highlight that while path length is important, balancing a range of attributes is required to optimise performance.

3. Path selection as a function of average carbon emission $\overline{EM}$, queuing delay $\overline{Q}$ and electricity cost $\overline{el}$ between device and DC, weighted by number of nodes n and average carbon emissions on path p within network G (Eq. (6)):

$$n \left(\left(\overline{EM}_{p \in P^*(n;G)} \right) \left(\overline{Q}_{p \in P^*(n;G)} \right) \left(\overline{el}_{p \in P^*(n;G)} \right) \right) \left(\overline{EM}_{p \in P^*(n;G)} \right) \quad (6)$$

Eq. (6) is comparable to Eq. (5) in that it strikes a balance between a range of attributes. In addition, it weights the calculation by average carbon emissions and node number to prioritise these attributes in the path selection made. The least cost path is selected as the optimum route.

When Eq. (6) is applied, the following paths are selected by the DCe-CAB orchestration agent as optimum:

path f (DC5) for device 7; **path l (DC1)** for device 9;
path t (DC6) for device 10; **path q (DC1)** for device 12.

This evaluation improves upon the effectiveness of Eq. (5) for energy-efficiency objectives by considering the range of context detail available and providing decisions according primarily to number of path nodes and emissions from each. Paths selected in this instance are one hop, limiting the carbon emission, queuing delay and electricity cost to which transmissions are exposed. For devices 12 and 9, Eq. (6) results in carbon emissions reduction by 50% from the maximum possible, and average carbon emissions are achieved in the communication with device 7 (20% lower than the maximum, 5% higher than the minimum). For device 10, carbon emissions on both one-hop paths are the same; the algorithm therefore makes the selection as a function of queuing delays and electrical operating costs, selecting the lowest of both. The results of Eq. (6) subsequently validate ability of the optimisation approaches to select paths on the basis of carbon emissions while balancing financial cost and queuing delay (see paths selected by Eq. (4), Eq. (5) and Eq. (6) for device 9 in Table IV) to provide energy-aware management of DC operation.

V. CONCLUSION AND FUTURE WORK

In this paper, we demonstrate optimisation approaches which improve efficiency of operations within and between client devices and DCs. Algorithms presented focus on context attributes which include queuing delay, carbon emission, electrical cost and number of nodes on the network path. The optimum solution uses the full range of context available, weighting the relationship between attributes by number of path nodes and carbon emissions from the network to result in selection of shortest hop paths; a balance is achieved between the other context attributes. These results therefore validate the optimisation approaches which select paths on the basis of carbon emission and

TABLE IV. ONE-HOP PATH CHARACTERISTICS FROM DEVICE 9

Client	DC	el	Q	EM	Eq. which selects path
9	1	13.3	0.5	0.5	Eq. (6)
9	5	10.0	1.0	0.8	Eq. (4), Eq. (5)
9	4	12.2	0.05	1.0	-

provide energy-aware DC management.

As part of future work, we propose to extend the optimisation algorithms presented in this paper to include decisions in relation to specific application QoS requirements such as those presented in Eq. (2) and Eq. (3).

ACKNOWLEDGMENT

This research is supported by the India-UK Advanced Technology Centre of Excellence in Next Generation Networks, Systems and Services, funded by the UK EPSRC Digital Economy Programme and the Government of India Department of Science and Technology (DST). We also acknowledge Professor A. Jhunjunwala and Dr. D. Jaliha from the Indian Institute of Technology, Madras for their contributions to this paper.

REFERENCES

- [1] European Commission Renewable Energies Unit, "Code of Conduct on Data Centres Energy Efficiency," Version 1.0, Oct. 2008, pp. 1-20.
- [2] R. Schmidt, "Information Technology Energy Usage and our Planet," in *Proc 11th Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems*, May 2008, pp. 1255-1275.
- [3] Department of Energy & Climate Change, "Impacts of Climate Change in the UK," Available [Dec. 2010] at: <http://www.decc.gov.uk/>.
- [4] Department of Energy & Climate Change, "Quarterly Energy Prices," 2008-2010; Available [Dec. 2010] at: www.decc.gov.uk/.
- [5] British Telecommunications Plc., "A Realist's Guide to Green Data Centres," White Paper, 2007, pp. 1-8.
- [6] A. E. H. Bohra and V. Chaudhary, "VMeter: Power Modelling for Virtualized Clouds," in *Proc. IEEE Int. Symp. on Parallel & Distributed Processing, Workshops and Phd Forum*, Apr. 2010, pp. 1-8.
- [7] D. C. Hwang, V. P. Manno, M. Hodes and G. J. Chan, "Energy Savings Achievable Through Liquid Cooling: A Rack Level Case Study," in *Proc. IEEE Intersociety Conf. on Thermal and Thermomechanical Phenomena in Electronic Systems*, Jun. 2010, pp. 1-9.
- [8] A. Beloglazov and R. Buyya, "Energy Efficient Resource Management in Virtualised Cloud Data Centers," in *Proc. 10th IEEE/ACM Int. Conf. on Cluster, Cloud and Grid Computing*, May 2010, pp. 826-831.
- [9] R. Tozer and M. Salim, "Data Center Air Management Metrics – Practical Approach," in *Proc. IEEE Intersociety Conf. on Thermal and Thermomechanical Phenomena in Elec. Systems*, Jun. 2010, pp. 1-8.
- [10] C. Peoples, G. Parr, B. Scotney and A. Moore, "A Context-Aware Policy-Based Framework for Self-Management in Delay-Tolerant Networks: A Case Study for Deep Space Exploration," *IEEE Communications Magazine*, Jul. 2010, pp. 102-109.
- [11] C. Peoples, G. Parr, B. Scotney and A. Moore, "Operational Performance of the Context-Aware Broker (CAB): A Communication and Management System for Delay-Tolerant Networks (DTNs)," in *Proc. of 2nd IEEE Int. Conf. on Advances in Satellite and Space Communications*, Jun. 2010, pp. 128-133.