Application of Genetic Algorithm to Maximise Clean Energy Usage for Data Centres

Raymond Carroll, Sasitharan Balasubramaniam, Dmitri Botvich,
and William Donnelly

TSSG, Waterford Instituite of Technology, Ireland
{rcarroll,sasib,dbotvich,wdonnelly}@tssg.org

Abstract. The communications industry is currently witnessing a continued increase in energy consumption, and this trend is predicted to increase even more in the coming years. This is largely driven by the popularity of the Internet, which continues to attract growing numbers of users who now rely on the Internet as part of their daily lives. A major factor behind this attraction is the multitude of services available on the Internet, ranging from web based services (e.g. facebook) to heavy power consuming services such as multimedia (e.g. youtube, IPTV). Therefore the data centres housing these services are seeing their energy consumption increase proportionally, now leading researchers to actively search for solutions to improve the energy efficiency of data centres. In this paper we propose a green data centre solution that makes data centres and services prioritise the usage of clean, renewable energy sources. The solution allows data centres to share information regarding renewable energy and cooling, in order to exploit variance between different countries energy and temperature profiles by moving services between data centres. We employ a genetic-algorithm to find the optimal placement of services on the data centres.

Keywords: Green Data Centres, Energy Efficiency, Genetic Algorithm.

1 Introduction

In recent years there has been a growing focus on the impact of the internet, and more specifically data-centres, on the environment, in terms of their increasing energy usage. Figure 1(a) shows that from 2000-2006 the energy usage for data centres in the US [1] more than doubled. It also depicts the predicted trends up to 2010, extrapolated based on both the historical data and also based on recent trends towards energy efficiency, where both show huge increases in energy usage.

While these trends do consider the impact of the move towards more energy efficient practices, they do not consider the impact that new technologies and computing models may have. For instance the growth in the usage of 'smart' phones in recent years has been exceptional. These phones are in essence resource limited computing platforms, where often times much of the processing is done in back-end service/application residing on the data-centre. Also, the recent move towards cloud

J. Suzuki and T. Nakano (Eds.): BIONETICS 2010, LNICST 87, pp. 565–580, 2012.

computing holds huge potential for increasing data centre usage. Cloud computing proposes to move all the majority of the application processing and data storage into the data centre, with 'thin' client devices running simple interfaces. These emerging trends suggest that data-centres could grow beyond what has been predicted, and continue this rate of growth into the future. At the same time, many countries are now actively pursuing more renewable sources of energy, through their own capital infrastructure projects or through grid feed-in tariff incentive schemes. This is illustrated in Figure 1(b), where we show the recent capacity increases in wind and solar energy within the EU states.

Based on these developments, our work attempts to address the problem of data-centre energy usage by allowing data-centre operators to determine a service placement strategy with the best renewable and cooling energy profile. This in turn reduces the overall carbon footprint of the data centre operator. To do this we employ a genetic algorithm (GA) based service placement approach, where the GA determines the most optimal service/data-centres pairings to maximize data-centre usage of renewable energy sources and minimise cooling energy.

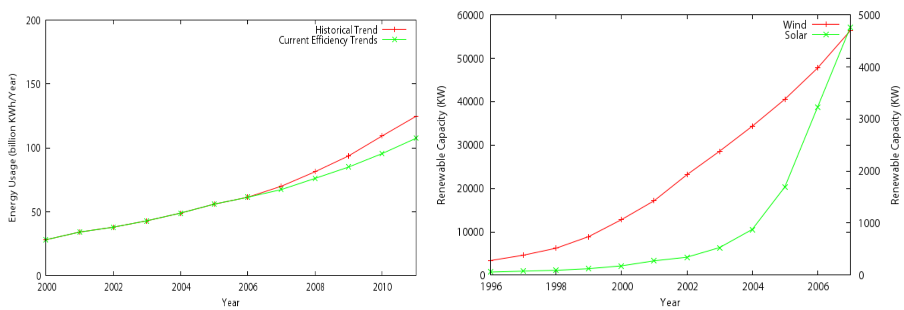

Fig. 1. (a) Data Centre Energy Trends (b) Wind and Solar Capacity increases of EU countries

2 Related Work

There is now an array of work being carried out in the area of energy efficiency for data centres. Many of these approaches focus on moving or scheduling workload in some way to achieve greater energy efficiency. Many of these have looked at consolidating workloads on a minimum number of servers in order to allow certain servers to be switched off or sleep to save power [2][3][4][5]. In [6] biological mechanisms are used to determine more efficient servers in a data centre where load is subsequently moved. Other works investigate how to make the data centres more efficient by reducing the load on the cooling systems through better workload placement and scheduling within and between data centres [7][8][9][10][11]. A similar approach to ours is taken in [12], where traffic load is moved between data centres based on electricity costs rather than renewable energy.

3 Sustainable Energy Prioritization Solution

In this section we introduce our proposed solution, outlining the service placement process and underlying genetic algorithm. Our goal is to place services on data centres in such a way as to maximise renewable energy usage and minimise cooling energy. The renewable energy consumption of a data-centre is measured through a value we call the renewable energy ratio (RER – Section 4), and this is used as a key metric in calculating the data centres which use more renewable energy. In terms of cooling, the ambient temperature of the country can have a significant effect on the efficiency of the cooling system. This efficiency is measured through a value termed the Coefficient of Performance (COP – Section 3.2) and so this value also plays a key role in calculating the data centres with the most energy efficient cooling. By using these values we can determine the best data-centres to place services on which will result in more renewable energy usage and lower cooling energy usage.

3.1 GA-Based Service Placement

The service placement process takes place in effectively three stages. Initially all data centres must co-ordinate and share information regarding their renewable and cooling energy usage levels, service usage details and data centre configuration information. This information forms the basis for our genetic algorithm to determine the fittest service configuration. In our solution the genetic algorithm is run periodically by a specific, pre-selected data centre (referred to as the GA-DC). However, for the purposes of redundancy each data centre is capable of running the algorithm and so the energy data is shared among all data centres.

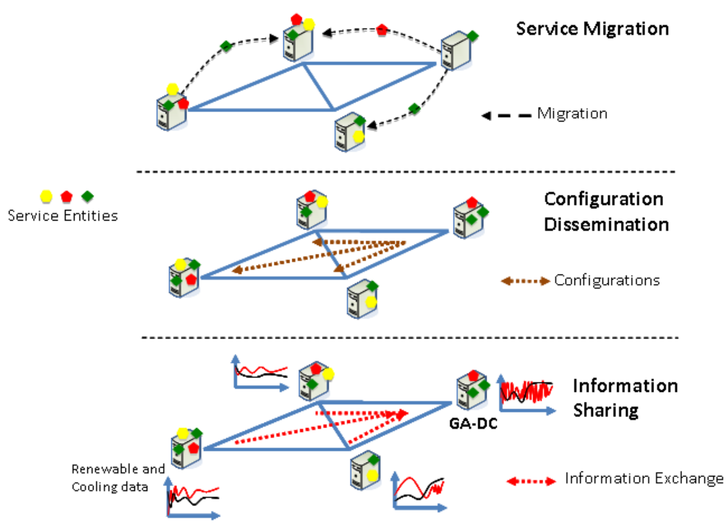

Fig. 2. Service Placement Process

As mentioned earlier the genetic algorithm will determine the most optimal service/data centre configuration, based on maximising the usage of renewable energy and minimising the consumption of energy in cooling the data centre. The genetic algorithm is described in detail in the following sections.

Once the optimal configuration has been found, this configuration is disseminated to all data centres. Each data centre then examines this configuration and implements its recommendations. To do this the data centre determines which of its currently hosted services have been selected to move and to what data centre. Once ascertained the data centres then migrates the selected services to their newly designated data centre. Each data centre is notified of the change in order to update their information registries and so that requests for the service at the originating data centre will be forwarded to its new location.

3.2 Genetic Algorithm

We begin the discussion of our genetic algorithm by looking at the fitness function we will use to determine the best solution. The fitness function is composed of two parts, the renewable energy consumed and the cooling energy consumed. These are conflicting optimization goals, since the data centres with the best renewable energy ratio are not necessarily the data centres with the best environmental conditions for efficient cooling. Below *(1)* we present our approach for maximising the renewable energy consumed and minimising the cooling energy used. Let the set of services be $S_i = \{s_1, s_2,...s_i...s_N\}$, where N is the total number of services; the set of Data Centres be $DC_j=\{DC_1, DC_2,....DC_j.....DC_M\}$, where M is the total number of Data Centres. Let RER_j be the Renewable Energy Ratio of data centre j, CE_j be the Cooling Energy of Data Centre j, sl_{ij} be the service load of service i on data centre j, and DCC_j is the capacity of data centre j. In *(1)* below we present our fitness function which attempts to maximise the renewable energy consumed and minimise the cooling energy used.

$$\max \sum_{j=0}^{j<K} \sum_{i=0}^{i<N} ((1-\alpha)(sl_i.RER_{DCj}) + a\frac{1}{CE_j}) \tag{1}$$

Subject to

$$\sum_{i=0} sl_{i\,j} < DCC_j \tag{2}$$

To determine the renewable energy consumed we use the renewable energy ratio (RER_j) of the data centre in question. The RER is the ratio of renewable energy production to total energy production in the data centres host country (see Section 4). This ratio gives us the best indication possible of what proportion of energy from renewable sources the data centre is consuming. However the quantity of renewable energy consumed is a factor of the load on the data centre also. As such we need to calculate the load exerted on the data centres by the service (sl_i).

The cooling energy (CE) of a data centre is calculated based on heat load (HL) to be removed from the data centre subject to the efficiency of the cooling system (COP)

in removing this heat (see *(3)*). The heat load is directly related to the energy being consumed by the computing equipment, which is then converted to heat. As such we calculate the heat load by determining the power being used in the data centre. This is shown in *(4)* , where P_{max} is the maximum power a single server consumes at peak load, P_{idle} is the power a single server consumes while idle (load = 0), *sl* is the load exerted on the server by a single service *i* and *ns* is the number of servers in the data centre. Since the idle power is consumed irrespective of server workload, the workload only impacts the power consumed above the idle power ($P_{max} - P_{idle}$).

$$CE = \frac{HL}{COP} \tag{3}$$

$$HL = \sum_{i=1}^{i<N} (sl_i).(ns.(P_{\max} - P_{idle})) + ns.P_{idle} \tag{4}$$

Before we discuss the COP in more detail, there are constraints on the GA which we must mention. The utilisation of the services assigned to a specific data centre cannot exceed the capacity of that data centre (2). In addition, a service must be assigned to only one data centre (especially important in crossover and mutation).

Coefficient of Performance (COP)

Critical to the calculation of the cooling energy fitness value is the Coefficient of Performance (COP) of each data centre. The COP value indicates the efficiency of the cooling system in removing the heat load from the data centre *(5)*. A high COP means the thermodynamic process is more efficient.

$$COP = \frac{HeatLoad}{CoolingEnergy} \tag{5}$$

$$COP = \frac{1}{T_H/T_C - 1} \quad \text{Th = Outside Temperature} \quad \text{Tc = Inside Temperature} \tag{6}$$

Under the principles of thermodynamics [13], the efficiency of a typical heat pump is highly dependant on both the inside (target) temperature and the environmental (outside) temperature to which the removed heat is rejected. Therefore the greater the outside temperature is (for a set inside temperature) the more inefficient the system.

In most cases, once the outside temperature drops below the indoor temperature air conditioning is typically not required. However in the case of data-centres, the primary heat load is not coming from heat transfer from the environment but rather the computing equipment, so cooling is still required. In line with best practices of data centre cooling, we design each data-centre with a free cooling system in addition to conventional cooling. Free cooling allows data-centres to utilize the outdoor environmental conditions to part, or even fully, cool the data centres when conditions allow. Typically this is when the outside temperature is below the indoor temperature, thus free cooling is also highly dependant on the weather conditions.

As a result we employ a COP model based on the assumption that both free-cooling and electric/mechanical-cooling are employed. Once the outdoor temperature is above the required cooling temperature, COP values based on standard electrical cooling are employed. However, when the outdoor temperature drops below the required indoor temperature we move to free cooling and adapt the COP model inline with the changeover. We do not subscribe to a specific free-cooling system, instead we generalize based on the assumption that free-cooling provides a significant improvement in efficiency of the cooling system.

For our non-free cooling COP model we adopt COP values from the ORNL [14] heat pump simulator. For free cooling we simply adapt the values of the ORNL COP such that when the outside temperature drops below the inside temperature, we adjust the COP relative to the original COP value (e.g. +40%). This aims to represent that, once the outside temperature is cooler than inside, the free-cooling system is in operation. However we do not assume that free-cooling COP is uniform, as the energy required by a free-cooling system can vary depending on the extent by which the outside temperature is cooler than inside. For instance air-pumps may need to pump less air to cool the server room the cooler the outside temperature gets. So, using this model we can determine the COP based on the known outside and inside temperature.

The next step in implementing our GA solution is to encode our problem into a chromosome representation. In essence each chromosome is required to represent a configuration of the entire set of services placed across the nine selected data centres. In our representation each gene represents a single service placed on a single data centre. Specifically, each gene contains two parts, the service in question and an ordered list of binary values indicating the data centre on which the service is placed. In this way we must first calculate the fitness of each individual gene, by examining the workload details of the service, and the details of the data centre on which it is placed. Once each gene's fitness has been calculated we can then sum these values to determine the overall chromosome fitness.

3.3 Genetic Algorithm Operation

The algorithm begins by randomly creating a population of chromosomes subject to the constraints outlined before (2). Once the initial population has been generated, we calculate the fitness of each chromosome. Then, based on the elitism approach, we select the two fittest chromosomes to be carried forward to the next generation. In order to populate the remainder of the new generation, we select two parent chromosomes and perform crossover in order to generate new offspring chromosomes. We employ roulette wheel selection to choose the parents and then perform single point crossover to create the child chromosome(s). Again care must be taken when performing crossover, that the resulting chromosome(s) do not cause any data-centre to exceed its maximum capacity. In our algorithm mutation is carried out by simply changing one of the binary values representing the data centre on which a service resides. In other words this results in the service being placed on a different data-centre than before, and again care must be taken to not break the constraints outlined.

3.4 GA Evaluation

In this section we perform some initial evaluations of the genetic algorithm itself, to ensure its correct operation. We vary the main GA parameters (population size, generations, crossover rate) as seen in the results presented in Figure 3.

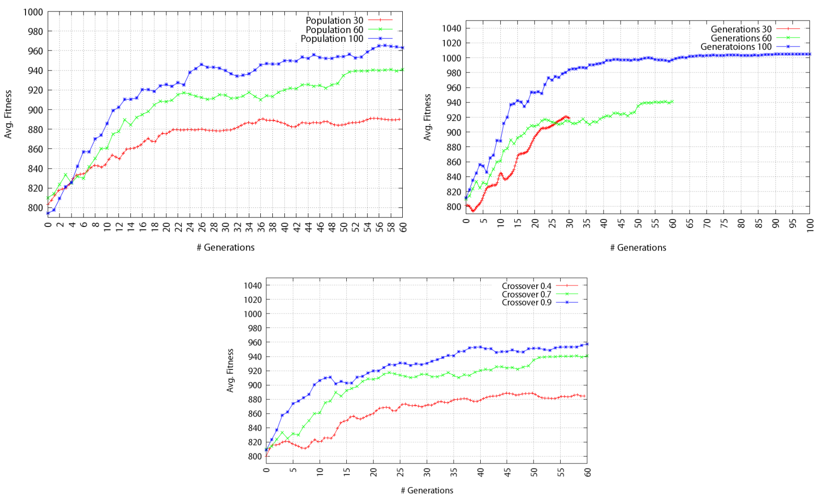

Fig. 3. (a) - Varying Generation Size (b) Varying Population Size (c) Varying Crossover Rate

We start by varying the number of generations for which the GA runs, as seen in Figure 3. As expected the greater the number of generations the better the overall result. The smaller generation size does obtain a reasonable fitness value quickly, but the 60 and 100 generation simulations are able to obtain higher values over time. Increasing the population size reduces the effects of randomness and gives a more diverse starting population. As expected, this leads to a stronger average population fitness, increasing in line with the population size increase. Finally, again as we would predict, higher crossover rates lead to more diversity in the populations and hence allow fitter, more optimal solutions to be found. At a low crossover rate we can see that the algorithm struggles to improve the population fitness since it is more difficult to breed new solutions from parents with higher fitness values.

4 Case Study

In order to properly present our solution we confine it to a fixed case study comprising a specific set of data centres and services, which will also be used later in the evaluation of our solution. We also specify the real temperature and energy data we have obtained, which are used in the calculation of the fitness function. Nine data centres were selected from major European cities, including Dublin, London, Lisbon, Madrid, Milan, Athens, Amsterdam, Berlin and Copenhagen. These were chosen in an attempt to give a significant variation in both climatic conditions and sources of renewable energy used.

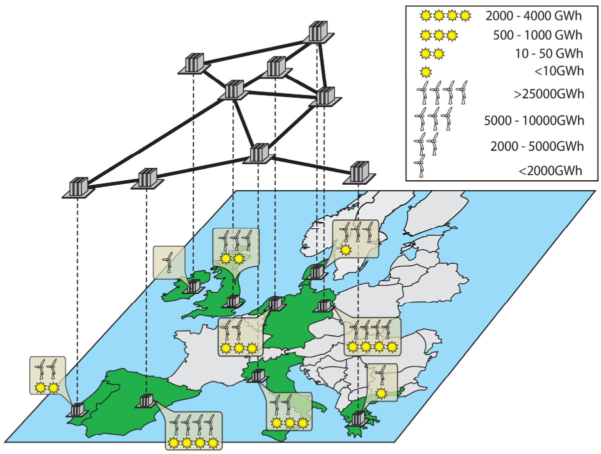

Fig. 4. Case Study Configuration

The data centres are connected in a network topology as depicted above, in line with the European Optical Network. The size of each data centre is determined by the number of servers it contains, which is relative to the population of the host country. Core to the approach taken in this paper is the use of real energy and weather data for the countries where the data centres reside. In line with this we have carried out a detailed search for data relating to the renewable energy production of each data centre host country, as well as temperature variations for a period covering January 2007 to December 2009. The energy production values, as described subsequently, are taken from [15].

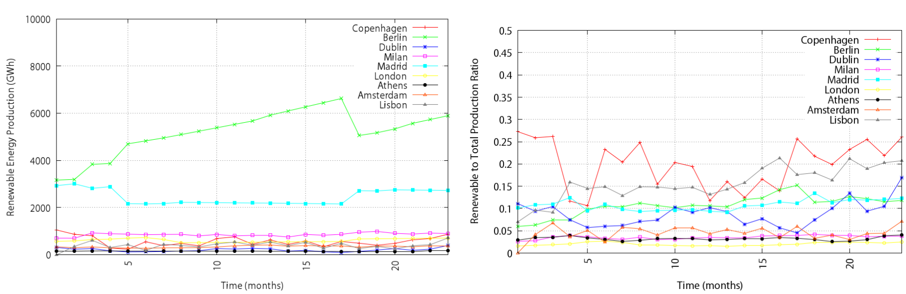

Fig. 5. (a) Renewable Energy Production, (b) Renewable Energy/Total Production Ratio

In Figure 5(a) we show the total renewable production for each of the selected data centres countries over the aforementioned period. Instinctively the larger countries will have greater volumes of renewable energy production (e.g. Germany and Spain), with the country's policy for sustainable growth also affecting capacity values. Other factors, such as weather conditions and capacity increases account for variations in the values from month to month. In Figure 5(b) we present the renewable energy values as a fraction of the overall energy production of the data-centres home country. This gives a clearer representation of the countries with the most energy production, and hence those that are more desirable candidates for services to migrate to. In terms of the temperature variations of each country/data centre, we used data from the European Climate Assessment & Dataset (ECA&D) [16] project, recording real temperature data from across Europe. Also, in order to calculate the cost impact on data centres we also used the real energy unit price as reported by the European Commission (Eurostat [17]).

5 Simulation and Results

In the following section we perform a case study simulation of the potential renewable energy gains possible for a small sized data centre operator, based on the genetic algorithm and scenario outlined in the previous sections. The operator runs nine data centres distributed as seen in Figure 4. Within each of these data centres there are a varying number of servers (8-200) and services (16-400), proportionate to the population of the country. In terms of the server specifications, we stipulate a standardized server across all data centres with a maximum power draw (P_{max}) of 400w and an idle power draw (P_{idle}) of 150w. To represent the workload exerted on the server, each service is randomly assigned a value that denotes how much of the servers processing capability it is using. This value effectively represents each services utilisation at a given time. In this work we keep the request rate uniform (i.e. we do not alter the service workload values) in order to allow clear comparisons in the evaluation of our solution. The request does vary between data centres however, proportionate to the population of the host country.

The simulation runs over 23 simulated months where the evaluation of data centre/service configurations by the GA takes place each month. In our simulations we compare our proposed approach using the genetic algorithm to the scenario where services remain statically on their allocated data centre. In the static case services are allocated relative to the size of the data centre and remain there throughout the course of the simulation. In Table 1 we show the parameters used for our simulations.

Table 1. Simulation Parameters

Parameter	***Value***
Population Size	100
Mutation Rate	.7
Crossover Rate	.1
# Generations	60
Free-Cooling Efficiency	+40%
α	0.5

In Figure 6(a) we present the overall quantity of renewable energy used when employing both the genetic algorithm approach and the static approach. As you can see the GA based solution out performs the static solution. In this case the GA utilises, on average, 15.9% more renewable energy than static services which accounts for approximately 1566MWh of electricity. The overall energy usage (of IT equipment) remains constant for both solutions, indicating that the GA did not increase the renewable quantity simply by increasing the total energy utilisation.

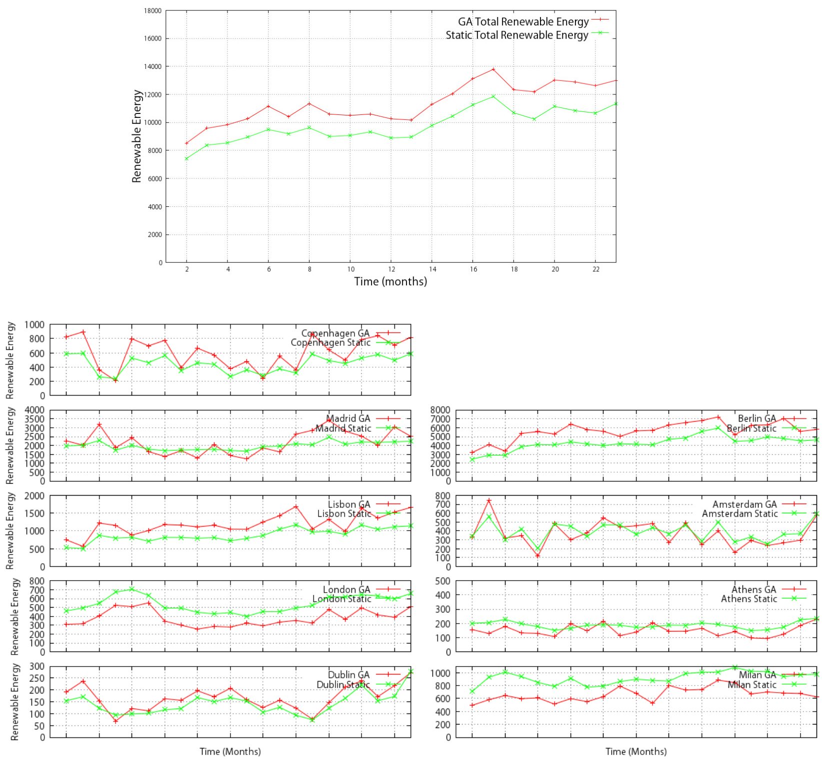

Fig. 6. (a) Total Renewable Energy (b) Renewable Energy per Data Centre

In Figure 6(b) we break this renewable energy usage down according to the individual data-centre usage. As expected, the change in the level of renewable energy used varies from data-centre to data-centre, depending on it's renewable ratio. Many data centres (e.g. Lisbon, Copenhagen, Berlin) increase their renewable usage while others (e.g. London, Milan, Athens) perform worse, using more fossil-fuel based energy. Increases in renewable energy are as a result of more favourable conditions (i.e. higher renewable ratios) in that country and vice versa. We can see a strong correlation here with the utilisation as indicated in Figure 8, as data centres that increase utilisation also increase renewable enregy utilisation, while those with lower utilisation decrease renewable usage. Correlation can also be seen with the renewable ratios in Figure 5(b), as those with generally higher renewable ratios gain renewable while those with lower ratios again perform worse.

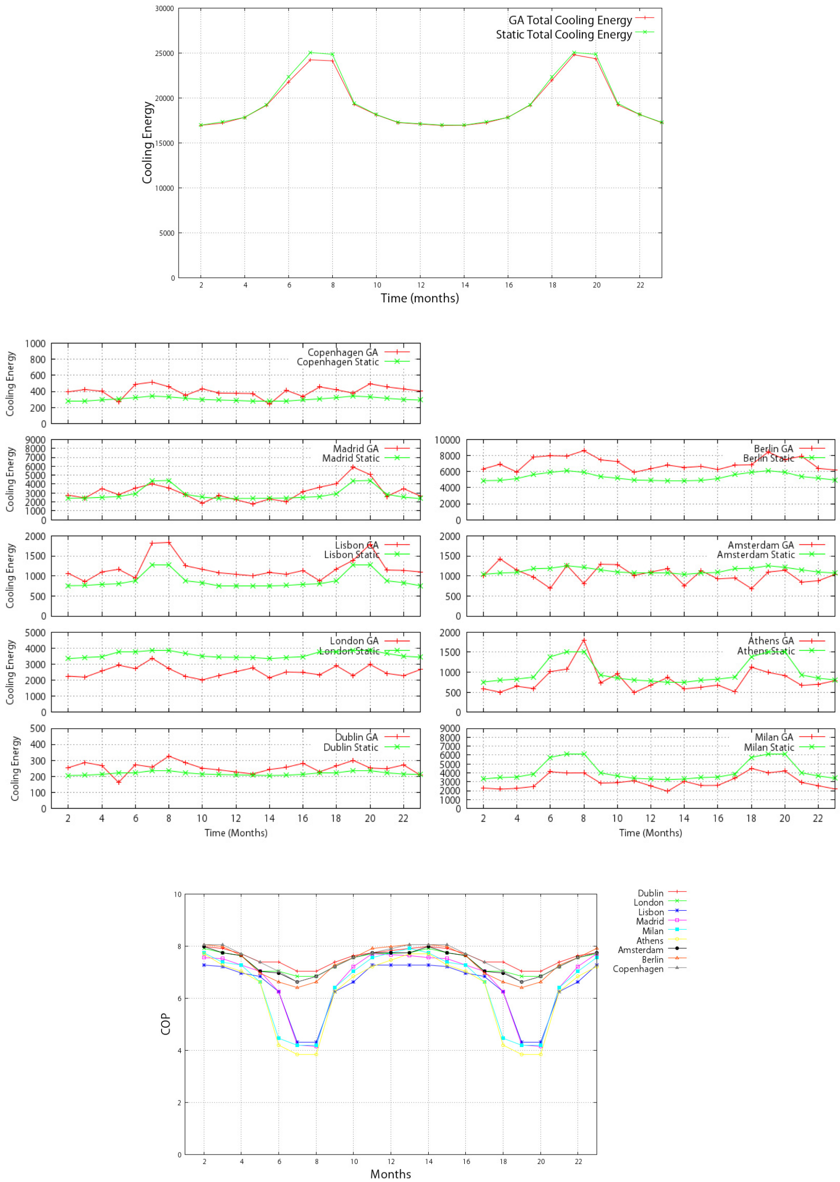

Fig. 7. (a) Total Cooling Energy (b) Cooling Energy per DC (c) COP per DC

In Figure 7(a) we present the total cooling energy used by both the genetic algorithm and static simulations. In the fitness function we aim to minimize the cooling energy expended, where cooling energy is directly proportionate to the COP of the data centre and the load exerted on the data centre. Savings in cooling energy are made by moving load to countries with lower temperatures and hence better COP values. If there is no significant variance in the temperatures of the countries then there is little opportunity to make significant savings.

Fig. 8. Utilisation per DC

In Figure 7(c) we can see that, during the winter periods the COP values for each data centre are very close and so the savings observed are very small. During the summer periods the COP values begin to diverge and so the GA can find service placements that can provide energy savings. In general however the savings observed for cooling energy are small (at best 3% for months 6-8), and this is due to the geographic proximity of the data-centres. In Figure 7(b) cooling energy values for each individual data centre are shown. Again there is a strong correlation here with the utilisation values in Figure 8, where DCs with higher utilisation will see higher cooling energy. The cooling energy values may appear somewhat counter-intuitive at first, given that the data-centres with the best COPs generally show increased cooling energy values for the GA (e.g. Dublin, Copenhagen, Berlin). However, given that the data centres with the best COP values are targeted for service placement, this will lead to increased utilisation and hence increased cooling energy usage. Since these have the most efficient cooling conditions, the cooling cost is lower than on those data centres with higher COPs for the same load. In other words, by removing load from lower-efficiency data centres and placing it on more efficient data centres we reduce the cooling energy consumed. There are some exceptions though, such as London and Amsterdam, who generally have good COP values but their cooling values do not necessarily reflect this. However we must consider the effects of the renewable aspect of the fitness function. Dublin, Copenhagen and Berlin also have good or very good renewable energy ratios, which make them more attractive for placement (i.e. higher fitness) while London and Amsterdam have poor or very poor renewable values. This offsets the effect of a positive COP value.

Figure 8 presents the utilisation experienced by each data centre over the course of the simulation. The utilisation values presented here are relative to the overall capacity of the entire data-centre group (i.e. all 9 data centres). As expected we can see that many of the data centres in the GA approach decrease capacity compared to the static while others increase. The data centres that consistently increase (Copenhagen, Lisbon, Dublin, Berlin) can be seen to correspond to those data centres that perform well in terms of renewable energy and also cooling energy. It must be noted that utilisation is also influenced by the capacity of the data centres. For instance Copenhagen is generally the best performer in terms of renewable energy and one of the top performers for cooling, yet Berlins utilisation increase is significantly larger. This is simply because Berlin is a considerably larger data centre and can handle a much larger utilisation increase. In terms of reduced utilisation, we can see that London and Milan show significant reductions with Athens, Amsterdam and Madrid show varying levels of reduction. For London and Milan, both have very poor renewable energy ratios (specifically London).

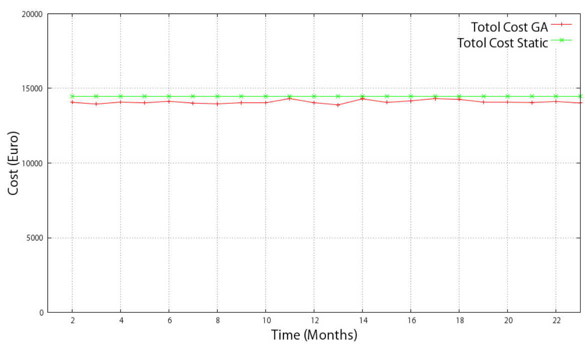

Fig. 9. Total Cost

Finally in Figure 9 we show the cost values incurred in both simulations. Since cost is not part of the fitness function this is presented only to evaluate the cost impact of our solution. Here it can be seen that the cost of our solution is lower than the cost of the static approach. However the cost decrease is very small (approx. 2.5% average), but the aim here is not to considerably reduce the cost, merely to ensure that our proposed approach does not come at a financial burden to the data centre operator. This is important in promoting the proposed solution to data centres operators, as increased costs will negatively impact the renewable energy benefits of the system.

6 Discussion and Future Work

Here we will briefly discuss some issues and notions that we feel warrant highlighting. Firstly, we abstract the energy grid of each country to be a 'black box', where we assume each energy source is fed into the grid and its source becomes indiscernible. In other words, the energy produced from renewable sources is not partitioned or reserved for specific usage, but is available in the grid for common usage in direct proportion to the rate at which it was produced. In electrical grids, detailed forecasting determines the required electrical capacity required at any given time. Consumption beyond this capacity requires additional generators to be brought online (typically using dispatchable, fossil fuel based energy sources such as natural gas) to cater for this additional consumption. In this regard, it may be suggested that causing a data centre to consume more energy in fact only utilises more non-renewable energy. However, when looking at the quantity of additional consumption our system places on the data centre, it equates to approximately 2×10^{-5} percent of the country's electricity production. This minute change we believe would be covered by the grids forecasting model. If the system were to be adopted in a large scale this may present a more significant issue, however, given the disparity in data-centre locations it might still not be directly discernable. That is to say that not all data-centre operators will have data-centres in the same countries, so the load increase will not always affect the same countries. However, in future work we aim to factor more detailed information from the grid, so that processing load only moves to data centres when its energy consumption drops below the forecasted consumption.

Another issue that we discuss here is the additional energy and latency costs that may be incurred by moving quantities of services between data centres. Moving services could cause higher loads on networking equipment along the migration path, hence increasing their energy consumption. Also, moving services further from the source of requests could potentially increase delays times and hence reduce end user Quality of Experience (QoE). In future work we plan to expand our simulations to evaluate these effects on the underlying network infrastructure. However, we feel that these effects could be limited by placing a distance limit on migrating services or indeed integrating a distance metric directly into the move metric objective function itself. In this way we could ensure that we always attempt to minimise the effects of migration on energy and QoS.

7 Conclusion

Due to the increasing popularity of the Internet, the communication systems of the future are predicted to consume large quantities of energy. In particular, the data centres that house various types of Internet services are poised to be the most significant consumer of energy. While improving energy efficiency is one objective of modern society, another key objective is to move towards green, renewable energy sources to reduce our carbon footprint. In this paper we have proposed a green data centre solution that uses a GA-based service placement approach based on targeting countries with the highest production of renewable energy and the best conditions for cooling. To validate our proposed solution, we carried out some demonstrative simulations by gathering data regarding the renewable energy production and temperature profiles of each country, and implementing a genetic algorithm that aims to maximise the quantity of renewable energy consumed and minimise cooling energy expended. From our simulations we have demonstrated that by employing this technique it is feasible to make significant improvements in the proportion of renewable energy utilised in data centre operation, hence reducing the quantity of fossil fuels burned and ultimately carbon emissions. We also demonstrated that cooling energy can be reduced in circumstances where there is significant variance in the country's temperature profiles. At the same time, we showed that this improved renewable energy utilisation did not come at an increased monetary cost for the operator.

References

1. ENERGY STAR. Report to Congress on Server and Data Center Energy Efficiency Public Law, 109-431. Environmental Protection Agency, Washington, D.C. (2007), `http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf`
2. Bradley, D.J., Harper, R.E., Hunter, S.W.: Workload-based Power Management for Parallel Computer Systems. IBM Journal of Research and Development 47, 703–718 (2003)
3. Das, R., Kephart, J.O., Lefurgy, C., Tesauro, G., Levine, D.W., Chan, H.: Autonomic multi-agent management of power and performance in data centers. In: The Seventh International Conference of Autonomic Agents and Multiagent Systems (May 2008)
4. Meisner, D., Gold, B.T., Wenisch, T.F.: PowerNap: Eliminating Server Idle power. In: ACM SPLOS (2009)
5. Rusu, C., et al.: Energy-Effcient Real-Time Heterogeneous Server Clusters. In: Proceedings of RTAS (April 2006)
6. Barbagallo, D., Di Nitto, E., Dubois, D.J., Mirandola, R.: A Bio-Inspired Algorithm for Energy Optimization in a Self-Organizing Data Center. In: Weyns, D., Malek, S., de Lemos, R., Andersson, J. (eds.) SOAR 2009. LNCS, vol. 6090, pp. 127–151. Springer, Heidelberg (2010)
7. Moore, J., Chase, J., Ranganathan, P., Sharma, R.: Making scheduling "cool": temperature-aware workload placement in data centers. In: Proceedings of the USENIX Annual Technical Conference 2005 on USENIX Annual Technical Conference, Anaheim, CA, April 10-15, p. 5 (2005)
8. Patel, C.D., Sharma, R.K., Bash, C.E., Graupner, S.: Energy Aware Grid: Global Workload Placement based on Energy Efficiency. In: IMECE 2003-41443, 2003 International Mechanical Engineering Congress and Exposition, Washington, DC (2003)
9. Bash, C., Forman, G.: Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the data center. In: Proceedings of the USENIX Annual Technical Conference, pp. 363–368 (2007)
10. Wang, L., et al.: Towards Thermal Aware Workload Scheduling in a Data Center. In: Proceedings of the 10th International Symposium on Pervasive Systems, Algorithms and Networks (IS-PAN 2009), Kao-Hsiung, Taiwan, December 14-16 (2009)

11. Sharma, R.K., Bash, C.L., Patel, C.D., Friedrich, R.J., Chase, J.S.: Balance of Power: Dynamic Thermal Management for Internet Data Centers. IEEE Internet Computing 9(1), 42–49 (2005)
12. Qureshi, A., Weber, R., et al.: Cutting the Electric Bill for Internet-Scale Systems. In: SIGCOMM 2009, Barcelona, Spain, August 17-21 (2009)
13. Moran, M.J., Shapiro, H.N.: Fundamentals of Engineering Thermodynamics. Wiley, New York (1995)
14. Oakridge National Laboratory (ORNL) Heat Pump Model, `http://www.ornl.gov/~wlj/hpdm/MarkVII.shtml`
15. International Energy Agency, `http://www.iea.org`
16. European Climate Assessment & Dataset (ECA&D), `http://eca.knmi.nl/`
17. European Commission: Eurostat, `http://epp.eurostat.ec.europa.eu/`