A Comprehensive Approach to Reduce the Energy Cost of Network of Datacenters

Baris Aksanli, Jagannathan Venkatesh, Tajana Rosing
Computer Science & Engineering Department
UCSD, La Jolla, USA
{baksanli, jvenkate, tajana}@ucsd.edu

Inder Monga
Energy Sciences Network
LBNL, Berkeley, USA
imonga@lbl.gov

***Abstract*—Several studies have proposed job migration over the wide area network (WAN) to reduce the energy of networks of datacenters by taking advantage of different electricity prices and load demands. Each study focuses on only a small subset of network parameters and thus their results may have large errors. For example, datacenters usually have long-term power contracts instead of paying market prices. However, previous work neglects these contracts, thus overestimating the energy savings by 2.3x. We present a comprehensive approach to minimize the energy cost of networks of datacenters by modeling performance of the workloads, power contracts, local renewable energy sources, different routing options for WAN and future router technologies. Our method can reduce the energy cost of datacenters by up to 28%, while reducing the error in the energy cost estimation by 2.6x.**

Keywords- Datacenter, energy, green energy, job migration

I. INTRODUCTION

Recent improvements in computer and network technologies and increasing computation demand have resulted in dramatic growth of large scale datacenters. Some companies, such as Google, Microsoft, deploy multiple datacenters distributed across the globe. A recent study indicates that the total power consumption of all datacenters in the world has increased by 56% from 2005 to 2010, with the total energy cost reaching billions of dollars [1]. The energy cost is expected to increase in the future as a result of higher cost of brown energy and the introduction of carbon emission taxes [2].

Previous publications focused on geographically separated datacenters, employed “follow the sun” or cheaper cost of brown energy strategies, [3], [4], [5]. They use WAN to increase system performance via load balancing [6], [7] improve energy efficiency by migrating jobs [3], [4], [5], or to determine new datacenter locations [8]. But, these studies do not model the energy cost of WAN, the long-term datacenter power contracts or the benefits of local renewable energy sources.

Datacenters often enter long-term power contracts with usage limits based on the expected peak power to limit the cost of energy. They are charged lower, fixed energy prices up to the contracted amount. Above the contracted power limit they pay energy at market prices which can be 5x more expensive [9]. Previous work neglected these power contracts. As a result, they miscalculate both the network usage and the potential energy cost savings.

Previous studies generally do not model local renewable energy sources due to their high variability. Some do account for expensive grid supplied renewable energy [10]. Renewable energy prediction can be used to decrease the variability of local green energy and increase its usage efficiency [11]. In our work, we include local renewable energy sources in the network of datacenters along with renewable energy prediction algorithms [11].

In this study, we develop a comprehensive algorithm that minimizes the energy cost of datacenters connected by WAN. Unlike previous work, our algorithm optimizes the datacenter and WAN parts of the system simultaneously, allowing feedback between them, thus improving the job migration decisions. Our model of network of datacenters leverages prediction of local wind and solar energy sources. We also include datacenter power contracts in our model and show that not accounting for these contracts can overestimate the energy savings by 2.6x. We analyze the impact of new technology and policies in datacenter WAN, such as energy-proportional routing and green energy aware routing policies. We show that using WAN to transfer workloads between datacenters can decrease the cost of energy by up to 28% with a negligible performance overhead compared to no data migration.

II. RELATED WORK

Moving jobs across WAN among different datacenters has been proposed by several studies. They all leverage the locality of the electricity prices, renewable energy availability and load demands due to different locations of datacenters. These studies take advantage of these differences to decrease electricity cost [12], [13], [4], [10], minimize carbon footprint [8] or maximize the load balance among the datacenters [7], [14].

Buchbinder et al. [12], Qureshi et al. [13] and Rao et al. [4] minimize the electricity cost by moving jobs to where the energy is cheaper. Their model lacks different energy types, detailed workload performance analysis and different routing options for datacenters. Le et al. [10] also minimizes the total energy cost with green energy included in their model. They assume a central dispatcher and do not model the network latency and cost. Hence, their model is not applicable to network of datacenters using WAN. Liu et al. [6] minimizes the brown energy usage by leveraging local green energy. However, they do not address the variability of green energy and do not have any network model or specify datacenter workloads they use. Aksanli et al. [7] use WAN to balance the load among datacenters. They include the network latency and job performance in their design but do not consider the cost of energy. In this study, we first outline the important aspects of datacenters and their WAN connectivity and show that a comprehensive approach is necessary to draw accurate energy saving estimations. We then design an algorithm to minimize the total cost of energy for various configurations.

III. Cost Minimization in Network of Datacenters

A. Datacenter and Backbone Network Models

Multiple datacenters provide several optimization opportunities, such as load balancing and cost minimization, due to the variation in green energy availability, brown energy pricing and load demand in different places. Although multiple datacenters increase the capacity and parallelism of the computation, a fast, reliable network is necessary to maintain performance of workloads during optimization. In this section, we first present our datacenter and backbone network models to capture the behavior of a large network of datacenters and then develop an algorithm to minimize the cost of distributed and renewable-powered datacenters connected with backbone network.

Figure 1: Network Topology; squares=datacenters, circles = routers

1) Backbone Network Model

We use a subset of LBNL ESnet's network topology, as shown in Figure 1, to model the datacenter WAN. It includes 5 datacenters and 12 routers distributed over the USA, where each connection link has a predefined capacity from 10Gbps to 100Gbps. We reserve 10% of this capacity for background traffic in our experiments. We compute the total network energy need with the router power and a fixed offset for the link power consumption. We estimate the router power consumption with a linear model based on bandwidth utilization [15].

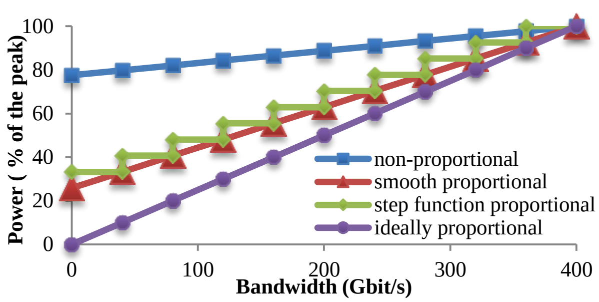

Figure 2: Power curves for different router power schemes

The power consumption of current routers is not proportional to the utilization, i.e. the idle power is very close to the peak regardless of the utilization. Since there have been several studies to design more energy proportional routers [7], we construct representative power proportionality models to account for future technologies. Figure 2 reflects 4 power curves representing different technologies: the non-proportional curve is an actual state-of-the-art router [16], and the step function is the result of turning line cards ON/OFF. This power curve is actually implementable with current technology [7]. Smooth proportionality is a linear correlation of the step function, and ideal proportionality assumes there is no idle power.

We implement two routing algorithms to find a path between two endpoints. The default algorithm is Dijkstra's Shortest Path Routing (SPR), which leads to static decisions. Our Green Energy Aware Routing algorithm (GEAR) [7] finds a path dynamically, based on green energy availability in router locations. Results section compares these routing algorithms based on their energy consumption and effects on job performance.

2) Datacenter Model

We model each datacenter in by the interaction between servers and the workloads they execute. We use measurements from our datacenter container on campus to capture the workload behavior and estimate the power consumption of a large scale datacenter. Our container includes 200 servers consisting of Nehalem, Xeon and Sun Fire servers. We run the workloads in multiple Xen virtual machines (VMs) and obtain the runtime characteristics of servers. Our measurements include resource, e.g. CPU and memory, utilization and power consumption values. We use these measurements to build and verify a baseline model for our simulations.

We classify the workload into two categories: service jobs with tight response time constraints and batch jobs with performance concerns but no explicit latency limits. The representative service jobs are RUBiS [17], a workload similar to eBay having 90^{th} percentile of response times at 150ms, and Olio [18], a social networking workload with response times ranging from 100ms up to multiple seconds. We run multiple Hadoop, a common MapReduce (MR) implementation [19], instances as batch jobs.

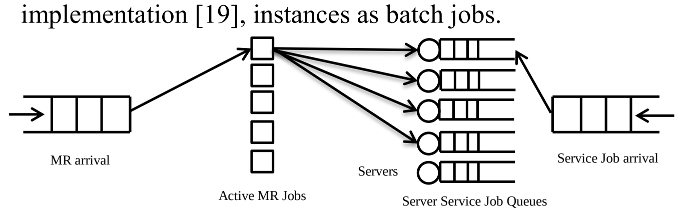

Figure 3: Job arrival processes in the datacenter model

Figure 3 shows the decoupled job interarrival to the system and how we schedule each different job. MR jobs are placed into a global active job queue whereas a service job is assigned to a server immediately. Each MR job consists of multiple tasks. A number of tasks are dispatched from an active job and sent to different servers. We use the scheduling policy described in [11] in order to guarantee the quality of service (QoS) of service jobs. This policy sets the maximum number of batch jobs that can simultaneously run on a server to limit the interference effects between service and batch jobs. We use separate lognormal distributions to model the interarrival process of each type of job and find the parameters of the distributions from measured data [11].

We model each server as an M/M/8 queue and calculate its power consumption with a linear CPU-utilization based equation [20]. We scale the aggregate server power consumption with power usage effectiveness (PUE) metric to account for power related overheads, such as cooling, distribution losses, etc. We use a PUE value of 1.15 in our simulations [21]. The error between our simulations and measurements is 3% for power consumption, 6% for service job QoS and 8% for MR job performance.

B. *Cost Minimization Algorithm*

The peak hours and electricity prices of datacenters have variations because of geographic distribution of datacenters. A datacenter with higher electricity prices can send some of its workload to another datacenter to take advantage of the lower prices there, with performance overhead due to network delay [4], [12], [13]. The live VM migration over high speed WAN has made this idea feasible, as it reduces the effects of network delay with fast data transmission [22]. Network elements should use the local green energy supply as much as possible to reduce their overall energy cost. But, today’s static routing decisions may not always minimize the brown energy consumption. A comprehensive cost minimization problem targeting network of datacenters needs to model the energy cost of both datacenter and network elements accurately to minimize the total cost and monitor the network delays in order to manage the performance of the workloads.

The main decision variable of the job migration is the energy prices. Previous studies mostly use the market prices to decide a data transfer. But, datacenters often make power contracts to decrease their energy cost. These contracts include a fixed price up to the contracted amount and change depending on the both the location and the expected demand of a datacenter. The demand exceeding the contracted amount is charged with market prices, which are generally much higher than the contracted price, up to 5x [9]. Therefore, a data transfer should not increase the power consumption beyond the power contract to avoid the high market prices. If these power contracts are not modeled, we can make incorrect data transfer decisions and cause power threshold violations.

In addition to the utility power, we include local renewables, such as solar and wind, to generalize the source of energy. Our algorithm uses accurate prediction methods to manage the unreliable nature of the renewable energy. It divides the timeframe into epochs of equal length. We set this length to 30 min because of the length of the prediction interval and that average batch job completion time is around 30 min [11]. The goal of our algorithm is to find which workloads can be relocated in each interval.

We prioritize green energy over brown energy as the green energy is generated and used locally and its cost is already paid at the time of deployment. Thus, we first target to maximize the green energy usage by transferring workloads to datacenters which have extra green energy available. Our next goal is to take advantage of cheap brown energy prices. These prices are fixed up to a contracted amount due to datacenter power contracts. Thus, we also maintain the utilization level in datacenters to avoid power threshold violations. If the utilization levels of all the datacenters are above their threshold values, we use the market prices in that interval to decide the data transfers.

Figure 4 outlines our algorithm. We calculate the energy need of each datacenter at the beginning of each interval, based on the load rate from the previous interval and the incoming load rate. For each interval we get the fixed and variable brown energy cost, peak power thresholds and estimate the green energy availability by using prediction. Then, each datacenter calculates the extra green energy based on its energy need and green energy availability. We send workloads from datacenters with the highest utilization to those with the largest amount of available green energy. Each sender selects jobs from its migrateable workloads to transfer until either there are no migrateable workloads left or there is no extra green energy in the receiver. In this study, batch jobs are the only migrateable workloads due to tight response time constraints of the service jobs. We continue this process until every datacenter is analyzed. Next, if there are any remaining workloads, the algorithm moves workloads from the datacenters with higher energy costs to those with the cheapest brown energy. This process continues until the sender does not have any migrateable workloads left or the receiver has reached its peak power threshold. Any remaining jobs after this point are scheduled based on the market energy prices and datacenter resource availability.

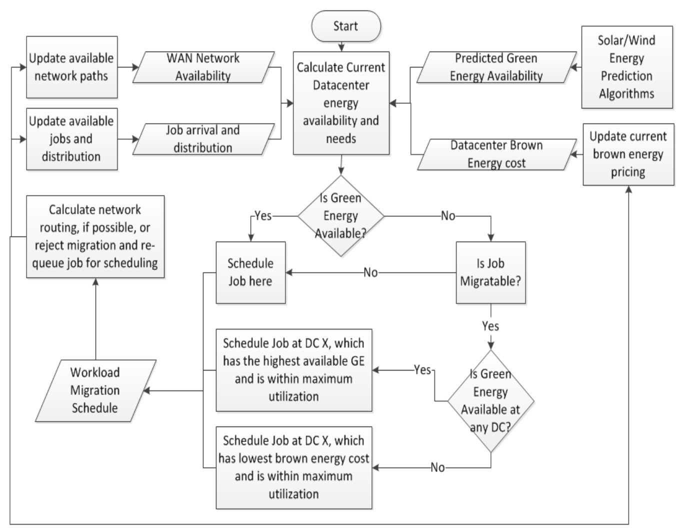

Figure 4: Flowchart of the cost minimization algorithm

We complete the datacenter part of the algorithm, which computes the transfer matrix representing data transfers, at the end of this iterative process. The next step is the networking part that takes the transfer matrix as input and calculates the path between each sender/receiver datacenter pair and how much bandwidth to be allocated on each path. The default routing algorithm used in today’s WAN is shortest path routing (SPR). We compare SPR to our green energy aware routing algorithm (GEAR). SPR always selects the shortest path between two points whereas GEAR allocates the path with the least brown energy need. The implementation details of SPR and GEAR are given in [7].

We include a load threshold for datacenters that is different than the power contract thresholds. It is used to reserve some resource capacity for emergency cases, such as hardware failures. Thus, we verify the load ratio of the receiving datacenter at the end of the networking part. If the ratio is higher than the load threshold of the datacenter, we cancel the incoming transfers. Similarly, if we cannot select a path for a job migration in the transfer matrix due to network unavailability, the job is rescheduled to either run on a different remote center or its original owner. We iterate this process until all datacenters are checked.

We also model the energy proportionality of both servers and routers in the network of datacenters. Along with the increased energy proportionality, our work is the first to consider local green energy with prediction, fixed power contracts, the impact of different routing algorithms, and green energy in WAN, all together in a cost minimization problem for network of datacenters.

Algorithm complexity: Our algorithm executes at most in time $O(n^2)$ in each step, where n is the number of datacenters. The first part of the algorithm converges faster because when the green energy availability is high, fewer workloads need to be migrated. In contrast, there are fewer migrateable jobs when fewer datacenters have green energy available. The second part of the algorithm can take longer to converge since and nearly all jobs could potentially migrate based on fluctuating energy prices. In this case, the number of migrateable jobs is limited by the available workloads and peak power thresholds.

The time required to apply the results of the algorithm in each epoch is the sum of running time of the algorithm and the communication delay between datacenters. The estimated running time of the algorithm is negligible for a network consisting of tens of datacenters. The communication delay is also low due to high speed (up to 100 Gbps) of the WAN links. The information each datacenter needs to broadcast in each epoch consists of green energy availability, brown energy prices, and the number and size of jobs that can be relocated. The amount of data required to represent this information is relatively small, and the time needed for small data transfers between datacenters are on the order of milliseconds [23]. Since we assume that all the data between datacenters is anyway replicated for fault tolerance needs, the actual transfer only requires any most recent changes to the job's data set, a relatively small fraction of the overall data. As the latency of batch jobs is on the order of minutes, this additional delay is insignificant.

IV. Methodology and Results

A. Methodology

We use an event-based simulation framework to analyze and compare the results of our solution to the cost minimization problem. The inputs to our simulator are derived from measurements on our datacenter container and data obtained from industrial deployments. This section discusses how we construct the simulation environment, including the datacenter loads, simulation parameters, green energy availability, and brown energy prices.

Datacenter load: We use a year of traffic data for two Google products, Orkut and Search, reported in the Google Transparency Report [24] to represent latency-centric service jobs and reproduce the waveform in Figure 3 in [25] to represent MR workloads to model the throughput-oriented batch jobs in datacenters. In Figure 5, we show a sample workload combination with Social Networking and Search jobs representing service jobs, and MR workloads to represent batch jobs. The maximum load ratio is around 80% with average of 45%. We use these data to find the parameters of the statistical workload models described in Section III.2 and feed them into our simulator. We list these parameters in Table I. We also only migrate batch jobs, which corresponds to 15-20% of the total workload, due to the tight response time constraints of service jobs.

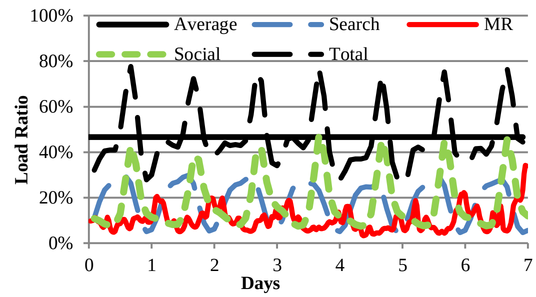

Figure 5: Datacenter workload mixture

Brown and Green Energy Costs: Datacenters make power contracts to decrease the cost of energy. These power contracts are generally adjusted to meet the load demand and if the datacenter exceeds the contracted amount of power, it is charged with higher market prices. We obtain sample contract pricing for the Midwest, the east and the west coasts from [10]. We use the California ISO [26] database to obtain brown energy prices for various California locations to represent variable market prices and time-shift and scale those values for the other locations based on published averages [27]. Figure 6 shows summer and winter daily pricing values for brown energy in comparison to contracted costs.

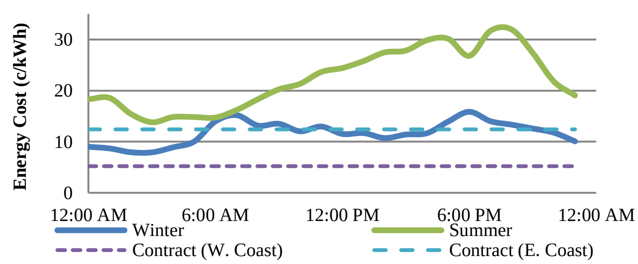

Figure 6: Daily brown and amortized green energy cost (c/kWh)

We represent the local green energy costs as a fixed offset to our cost model, which is amortized over the lifetime of an installation, including both the capital and the maintenance costs. We use data from [28] to obtain the capital and operational expenses of several solar and wind farms, amortized over their lifetimes, as representative solar and wind costs per interval.

Green Energy Availability: We obtain solar energy data from the UCSD Microgrid and wind data from a wind farm in Lake Benton, MN, made available by the National Renewable Energy Laboratory. Both solar and wind databases have data from a number of years that are sampled at minimum 1 second intervals. We scale these measured results to published average data for other locations in our experiments [29], [30].

Table II shows our simulation parameters and Table II presents our network topology and green energy availability. Green energy is scaled to 80% of peak datacenter and router energy needs.

TABLE I: Simulation parameters (ST: service time, IAT: interarrival time)

Parameter	Value	Parameter	Value
Avg. Web Req. IAT/client	5ms	Avg. # tasks/MR job	70
Avg. Web Request ST	20ms	Avg. throughput level per MR job	0.35
Service Request SLA	150ms	#servers/datacenter	1000
Mean MR Job IAT	2 min	#datacenters	5
Mean MR Task ST	4 min	#routers	12
Idle Server Power	212.5W	Idle Router Power	1381W
Peak Server Power	312.5W	Peak Router Power	1781W
Single link capacity [7]	100Gbps	Avg. batch VM size	8 GB [25]

TABLE II: Renewable energy and network configuration for each location

Location	Node	Type	Location	Node	Type
Chicago	DC+Router	Wind	Kansas	Router	-
Atlanta	DC+Router	Solar	Nashville	Router	Wind
New York	DC+Router	Wind	El Paso	Router	Solar
San Diego	DC+Router	Solar	Cleveland	Router	Wind
San Francisco	DC+Router	Both	Houston	Router	Solar
Denver	Router	-	Washington DC	Router	-

B. *Results*

In this section, we present the results of our cost minimization solution. We first compare the results with the case where there is no data transfer. Then, we analyze the effects of power contracts. The power contracts are applied based on a percentage of the peak power. For example, with an 85% power contract, a datacenter pays the contracted power price for each Watt up to 85% of the peak power and market price for each Watt exceeding that limit. We also evaluate different energy proportionality options for both servers and routers, along with the outcomes of GEAR.

No job migration: In this case we assume that each datacenter executes its own jobs, i.e. does not transfer data to another, and uses the locally available green energy. This case is the baseline for our comparison, as it represents the nominal brown energy cost and the base performance of batch jobs. We study different combinations of power contracts and server energy proportionality. The first and the third columns for each different power contract configuration in Figure 7 correspond to this case with non-proportional and energy proportional servers, respectively. We schedule service and batch jobs simultaneously [11]. The response time constraints of service jobs are always met (service QoS ratio is always less than one) where the average MR job completion time is 22.8 min. In addition, 59% of the overall green energy supply is consumed by the datacenters locally with the given green energy profile. We can use migration to improve the green energy usage level and decrease the total energy cost.

Cost minimization using migration: The main goal of our cost minimization framework is to maximize green energy and inexpensive brown energy usage, while minimizing the overall cost of energy and meeting performance constraints. Each datacenter can transfer workloads depending on the cost of energy in the remote center and the availability of local green energy. We also show the results of the case where servers are more energy proportional to quantify the benefits of our policy for future systems.

We obtain 66% of the green energy usage efficiency, 7% more than the baseline. This increase is significant given the high idle power of servers and the fact that we can move only 15-20% of the total workload. There is no performance overhead for service requests and the average job completion time for MR jobs is only 4.5% worse than the baseline.

Figure 7 shows the effects on total cost of adding energy proportionality and different types of power contracts to our model. The left axis shows the total cost that is normalized against the no migration case with an 85% power contract and non-proportional servers. The second and the fourth columns correspond to the energy cost values when job migration is allowed with non-proportional and energy proportional servers, respectively. The previous work neglects the power contracts, thus overestimating the total cost of energy by 28% even when there is no job migration. We obtain 19% energy savings with only market prices whereas the savings are 12% and 8% with a 70% and 85% power contract, respectively. Similarly, the previous work overestimates the cost savings by up to 2.3x when not considering power contracts. When we use energy proportional servers, the savings increase: 28% without a power contract; 18% and 11% with a 70% and 85% power contract. The savings are overestimated by up to 2.6x with energy proportionality. As seen, with future technologies, the error due to incorrect energy pricing increases. Figure 7 also shows that our algorithm can achieve up to 18% and 28% savings with and without a power contract, respectively.

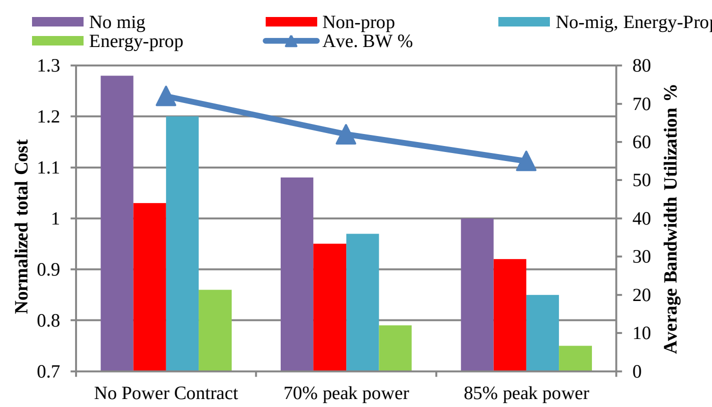

Figure 7: Normalized cost and average bandwidth utilization with different power contracts and energy proportionality

The right axis in Figure 7 shows the average bandwidth utilization and how it changes with different power contracts. Energy proportionality of the servers does not affect the migration decisions and thus bandwidth utilization values are similar compared to the non-proportional case. In contrast, the power contracts decrease the network utilization by up to 24%. This is because they create a more balanced energy cost scheme across datacenters. Furthermore, the contracted power level is inversely proportional with the network utilization. We observe 55% average bandwidth utilization with the 85% power contract and 62% utilization with the 70% contract. Datacenters reach the contracted power threshold earlier with a lower contract and start sending more jobs to each other to avoid expensive market prices. Therefore, neglecting the power contracts can result in overestimated network usage.

Using a green energy aware network: We next study the effects of using green energy in routers along with GEAR. Instead of selecting the shortest path between two datacenters, GEAR chooses the path with the least brown energy need. The average bandwidth utilization increases with GEAR as compared to SPR: 61% vs. 55% when using 85% peak power limit. The main reason is that GEAR sometimes selects a longer path, which results in higher utilization of the network. In Figure 8, we show the comparison between SPR and GEAR in terms of total and brown energy consumption of the routers. The brown energy percentage that GEAR uses 62% of the router peak power, compared to 65% of SPR, i.e. even though GEAR has larger overall power consumption, it lowers the brown energy usage. The difference between GEAR and SPR becomes more visible with increasing router energy proportionality, up to 65% with current proportionality schemes. GEAR also has 3% higher network delay as a result of sometimes choosing a longer path, whose effect on batch job completion time is unnoticeable.

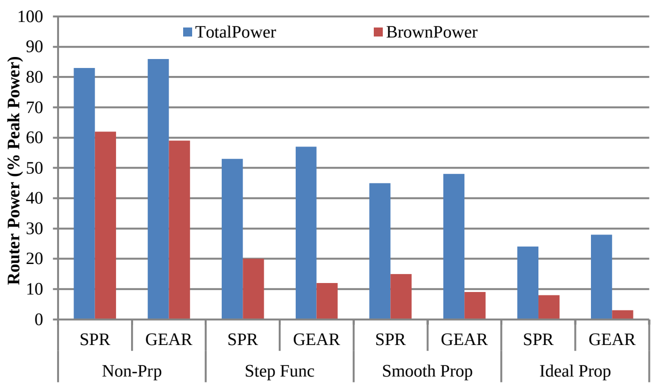

Figure 8: Comparison between SPR and GEAR in terms of total and brown energy consumption of routers with different router energy proportionality

V. CONCLUSION

Energy efficiency of network of datacenters has gained a lot in importance. Previous work migrates workloads over WAN to leverage the load demand, electricity price and green energy availability variations in different locations. But, these large scale systems have a lot of different aspects to be modeled that previous work neglects. In this paper, we show that using inaccurate datacenter energy price models can overestimate the cost savings by up to 2.6x. Our paper is the first to show the effects of fixed power contracts and energy proportionality on the cost savings of datacenters. We also use a dynamic green energy aware routing algorithm that can decrease the brown energy consumption of the network with a negligible performance overhead. Finally, our cost minimization algorithm reduces the datacenter energy cost by up to 28%.

ACKNOWLEGMENT

This work was sponsored in part by the Multiscale Systems Center (MuSyC), Energy Sciences Network (ESnet), NSF ERC CIAN, CNS, NSF IRNC TransLight - StarLight, Microsoft and Google.

REFERENCES

[1] J. Koomey, "Growth in data center electricity use 2005 to 2010," *Analytics Press*, 2011.

[2] J. Mankoff, R. Kravets and E. Blevis, "Some computer science issues in creating a new sustainable world," *IEEE Computer*, 2008.

[3] B. Bergler, C. Preschern, A. Reiter and S. Kraxberger, "Cost-Effective Routing for a Greener Internet," in *CPSCom*, 2010.

[4] L. Rao, X. Liu, M. Ilic and J. Liu, "Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment," in *Infocom*, 2010.

[5] A. Sankaranarayanan, S. Somsubhra and A. Federova, "Global cost diversity aware dispatch algorithm for heterogeneous data centers," in *ICPE*, 2011.

[6] Z. Liu, M. Lin, A. Wierman, S. Low and L. Andrew, "Greening geographical load balancing," in *Sigmetrics*, 2011.

[7] B. Aksanli, T. Rosing and I. Monga, "Benefits of Green Energy and Proportionality in High Speed Wide Area Networks Connecting Data Centers," in *DATE*, 2012.

[8] A. Mohsenian-Rad and A. Leon-Garcia, "Energy-information transmission tradeoff in green cloud computing," in *Globecom*, 2010.

[9] S. Govindan, A. Sivasubramaniam and B. Urgaonkar, " Benefits and limitations of tapping into stored energy for datacenters," in *ISCA*, 2011.

[10] K. Le, R. Bianchini, T. Nguyen, O. Bilgir and M. Martonosi, "Capping the brown energy consumption of Internet services at low cost," in *International Green Computing Conference*, 2010.

[11] B. Aksanli, J. Venkatesh, L. Zhang and T. Rosing, "Utilizing green energy prediction to schedule mixed batch and service jobs in data centers," in *HotPower*, 2011.

[12] B. Buchbinder, N. Jain and I. Menache, "Online Job-Migration for Reducing the Electricity Bill in the Cloud," *Networking*, 2011.

[13] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs, "Cutting the electric bill for internet-scale systems," *SIGCOMM Copt. Comm. Rev.*, p. 2009.

[14] L. Rao, X. Liu, M. Ilic and J. Liu, "MEC-IDC: joint load balancing and power control for distributed Internet Data Centers," in *ICCPS*, 2010.

[15] P. Mahadevan, P. Sharma, S. Banerjee and P. Ranganathan, "A Power Benchmarking Framework for Network Devices," in *Networking*, 2009.

[16] R. Tucker, J. Baliga, R. Ayre, K. Hinton and W. Sorin, "Energy Consumption in IP Networks," in *ECOC Symposium on Green ICT*, 2008.

[17] "RUBiS," [Online]. Available: http://rubis.ow2.org/ .

[18] Apache, [Online]. Available: http://incubator.apache.org/olio/.

[19] "Hadoop," [Online]. Available: http://hadoop.apache.org/ .

[20] D. Economou, S. Rivoire and C. Kozyrakis, "Full-system power analysis and modeling for server environments," in *MOBS*, 2006.

[21] Google, "Google Summit," 2009. [Online]. Available: http://www.google.com/corporate/datacenter/events/dc-summit-2009.html.

[22] F. Travostino, P. Daspit, L. Gommans, C. Jog, C. De Laat, J. Mambretti, I. Monga, B. Van Oudenaarde, S. Raghunath and P. Wang, "Seamless live migration of virtual machines over the MAN/WAN," *Future Generation Computer Systems*, 2006.

[23] T. Kraska, G. Pang, M. Franklin and S. Madden, "MDCC: Multi-Data Center Consistency," 2012.

[24] "Google," http://www.google.com/transparencyreport/traffic/.

[25] Y. Chen, A. Ganapath , R. Griffith and R. Katz, "The case for evaluating MapReduce performance using workload suites," 2011.

[26] "California ISO," 2012. http://oasis.caiso.com/.

[27] "U.S. Energy Information Administraion. Electric Power Monthly.," http://205.254.135.7/electricity/monthly/.

[28] R. Miller, "Green Data Centers. Data Center Knowledge," 2011. http://www.datacenterknowledge.com/archives/category/infrastructure/green-data-centers/.

[29] "NREL Solar Maps," 2012. http://www.nrel.gov/gis/solar.html.

[30] "NREL Wind Maps," 2012. http://www.nrel.gov/gis/wind.html.