

Operational Cost Optimization for Cloud Computing Data Centers Using Renewable Energy

Shaoming Chen, Samuel Irving, and Lu Peng

Abstract—The electricity cost of cloud computing data centers, dominated by server power and cooling power, is growing rapidly. To tackle this problem, inlet air with moderate temperature and server consolidation are widely adopted. However, the benefit of these two methods is limited due to conventional air cooling system ineffectiveness caused by recirculation and low heat capacity. To address this problem, hybrid air and liquid cooling, as a practical and inexpensive approach, has been introduced. In this paper, we quantitatively analyze the impact of server consolidation and temperature of cooling water on the total electricity and server maintenance costs in hybrid cooling data centers. To minimize the total costs, we proposed to maintain sweet temperature and available sleeping time threshold (ASTT) by which a joint cost optimization can be satisfied. By using real-world traces, the potential savings of sweet temperature and ASTT are estimated to be average 23% of the total cost, while 96% requests are satisfied compared to a strategy which only reduces electricity cost. The co-optimization is extended to increase the benefit of the renewable energy, and its profit grows as more wind power is supplied.

Index Terms—Cost optimization, data center, hybrid cooling, renewable energy.

I. INTRODUCTION

THE total cost of ownership (TCO) of cloud computing data centers consists of both one-time capital costs, incurred only at the beginning or upgrade stage of data centers, and recurring operational costs including electricity cost, maintenance cost, and salaries [6]. According to a recent report [5], the TCO is dominated by operational costs. As salaries are not a technical factor, we focus on the optimization of electricity and maintenance costs in this work.

The growth of the cost of electricity for supplying server power and cooling power outpaces expectations. In 2011, U.S. data centers spent about \$7.4 billion in electric power, with server power and cooling power contributing significantly to the total [33]. Several studies try to throttle the use of electricity in response to this increase, although a few of them consider the cost of server maintenance.

Prior works employ two methods for reducing energy cost: increasing the use of server consolidation and increasing the inlet air temperature.

Manuscript received May 7, 2014; revised September 21, 2014, December 19, 2014, April 15, 2015, and June 4, 2015; accepted July 26, 2015. This work was supported in part by NSF Grants CCF-1017961 and CCF-1422408.

The authors are with the Division of Electrical and Computer Engineering, School of Electrical Engineering and Computer Science, Louisiana State University, Baton Rouge, LA 70803 USA (e-mail: schen26@lsu.edu; sirvin1@lsu.edu; lpeng@lsu.edu).

Digital Object Identifier 10.1109/JSYST.2015.2462714

Server consolidation has been widely adopted to improve server energy efficiency by keeping active servers in high utilization and turning off overprovisioned servers [35]. Dynamic voltage and frequency scaling (DVFS) is also used to save server power [16]. However, the benefits of DVFS are shrinking due to increasing leakage power and because the voltage of processors is nearing the physical limit [27]. In addition, DVFS only reduces CPU power which only amounts to 30% of the consumed server power [33]. Server consolidation remains an effective and practical method to reduce server power consumption.

Increasing inlet air temperature is a common method for reducing the use of power for cooling. Raising inlet air temperatures by just one degree can reduce cooling energy consumption by 2%–5% [10]. However, the amount that inlet air temperature can be raised is very limited due to the requirement of keeping server temperature below the critical temperature. There are several prior works advocating the use of thermal-aware workload placement which distribute workloads according to a thermal map of data centers [29]. Unfortunately, these methods struggle to maintain energy efficiency of traditional air cooling when data centers are experiencing high utilization [35]. Therefore, a novel cooling system is required.

A hybrid cooling system is proposed as a practical and inexpensive solution to the problems of liquid cooling [18]. The system combines air and liquid cooling and has been deployed in data centers such as Aquasar, the first hot-water-cooled supercomputer prototype [44]. The hybrid cooling system uses water to cool down high-power-density components, such as processors and memory devices which dominate total heat dissipated in servers, while it uses air to cool down other auxiliary components which have a low power density. The hybrid cooling system can remove heat from a data center using less power than conventional air cooling.

In addition to the operational energy cost, the hardware maintenance cost is also considerable. According to a typical new multimewatt data center in the U.S., the cost of server repair and maintenance is approximately 50% of the costs of server power and cooling power [6]. Disks are the most frequently replaced components based on the empirical data of an HPC data center. The cost of disk maintenance can be increased by server consolidation due to the limited start–stop cycles of disks [11], since server consolidation frequently turns off servers or switches servers between the active state and the sleeping state. Additionally, higher inlet water temperature increases the cost of CPU and memory maintenance, since every 10 °C increase over 21 °C decreases the lifetime reliability of electronics by 50% [32]. Therefore, we can balance the

saving of the electricity costs and the increase of the costs of hardware maintenance by manipulating inlet water temperature and server consolidation.

On the other hand, the sustainability of data centers is becoming one of the top concerns of their owners, as three-year electricity bills of modern data centers grow over the server equipment cost [9]. The power sources are shifted toward renewable energies such as wind, solar, and tidal power, driven by soaring conventional energy price and the global warming. Wind power or tide power is integrated into our proposed optimization of electricity and server maintenance costs since wind energy is cheaper and widely used to power large-scale facilities [31].

The discussion of electricity and hardware maintenance costs drives us to propose our comprehensive framework covering these two costs. Integrating the models of electricity costs and hardware maintenance costs is nontrivial due to being studied separately by using different metrics. For example, the works focusing on electricity costs are likely to report their benefit in terms of power, while the works on hardware maintenance focus on expected lifetimes of hardware components. Although the two kinds of work are also studied in different scenarios, they interact with each other via inlet water temperature and server consolidation. Thus, how to fuse the models reasonably in a framework is our most challenging task. This framework distinguishes our work by optimizing these two costs together, while other prior works [37], [40] exclusively focus on electricity or hardware maintenance costs for data centers. Focusing on these two costs rather than one of them avoids categorizing our optimization as a suboptimal solution for the total cost.

The contributions of our work are shown in the following.

- 1) We set up analytical models for server power, cooling power, and hardware maintenance in a hybrid cooling data center for quantitative evaluation. We build a comprehensive framework which covers evaluations of these costs. To our knowledge, none of the existing methods have addressed this issue. This framework provides foundations to optimize the total cost in hybrid cooling data centers.
- 2) We propose a tradeoff between electricity and maintenance costs. In this paper, we show that the typical optimizations (high inlet water temperature and aggressive server consolidation) reduce the electricity costs but increase the maintenance costs.
- 3) To minimize the costs, we develop a joint optimization scheme based on server consolidation and dynamic optimal inlet water temperature. Our simulation results show that the method can gain considerable cost benefits.
- 4) We extend our cost optimization to exploit the benefits of two kinds of renewable energies: wind power and tide power. Based on our experiments, it increases the cost saving of the renewable energies, and this benefit grows as more renewable energies are supplied.

II. RELATED WORK

Prior works of the cost optimization of data centers fall into two categories: the optimization of electricity costs [37] and the optimization of hardware maintenance costs. The server power

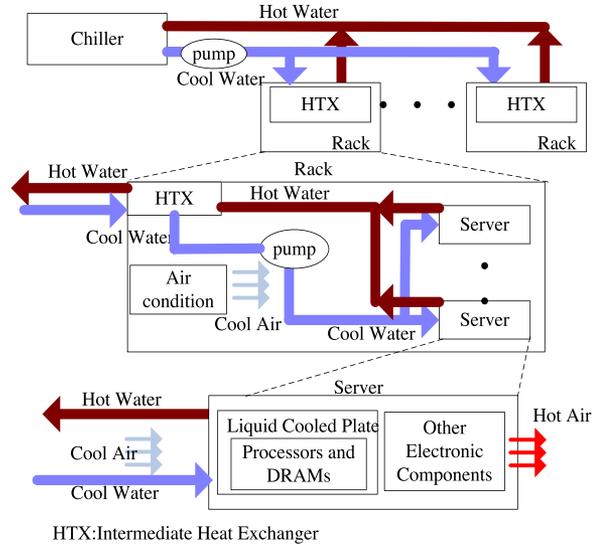


Fig. 1. Structure of hybrid cooling.

and the cooling power are the two main factors determining the electricity costs of data centers. Many studies focusing on the server power were addressed from server level [27], [28], rack level [36], and data center [11], [14], [25], and [41] to decrease the energy consumption of data centers. These focused on increasing server energy efficiency and reducing server idle power. On the other hand, Moore *et al.* [29] introduced thermal-aware workload placement to reduce cooling power in traditional air cooling data centers. On the contrary, other researchers employed advanced infrastructures of cooling systems to solve energy inefficiency of traditional air cooling [6], [19], [39]. However, all of these works just aimed at the reduction of either cooling power or server power.

To capture a broad scope of energy savings, several architects proposed approaches [5], [18], [33], [35] for optimization of server power and cooling power. For an example, Pelley *et al.* [33] set up a comprehensive framework of the total power of data centers to optimize server power and cooling power. Ahmad *et al.* [5] proposed a joint optimization of server power and cooling power with guaranteeing response time. However, all of these works did not consider the increment of the costs of hardware maintenance.

On the other hand, several papers discussed the issue related to hardware maintenance in data centers [23], [40], and [42]. Schroeder *et al.* [40] analyzed disk replacement rate based on the empirical data, which inspired researchers to study the reliability of hardware in servers.

Unlike the studies focusing on the optimization of electricity costs or hardware maintenance costs in data centers, our approach covered them both. Additionally, although Chen *et al.* [11] minimized the costs of energy and disk maintenance by combining DVFS and server consolidation, the author did not discuss the cooling cost and the maintenance costs of other components such as processors and memory in servers.

III. HYBRID COOLING

Fig. 1 shows the structure of hybrid cooling in modern data centers. The closed liquid loop between the chiller and the racks

is designed to remove heat dissipation from the racks. The cool water in the loop absorbs heat dissipation from the racks and returns back to the chiller with heat. In the closed liquid loop of a rack, the water is pumped into servers and cooled in the intermediate heat exchanger (HTX). The coolant water in a server flows through a liquid-cooled plate and takes away power dissipated by processors and memory devices. Other auxiliary components such as disks, power supply, and chipsets on the motherboard are still cooled by the air condition as traditional data centers since these components dissipate less power and, more importantly, exhibit lower power density compared to processors and DRAMs.

IV. COST MODELS

To optimize the electricity costs and the hardware maintenance costs, we set up the cost models which quantitatively estimate the impact of server consolidation and inlet water temperature on the costs when hybrid cooling is used.

A. Electricity Costs

The power of a typical data center includes server power, cooling power, and power distribution loss. Power distribution loss is denoted by P_{PDL} , which is equal to 10% of load power in our experiment [42]. In the following context, we address the models related to the server power and cooling power.

Server Power Model: P_{servers} consists of the aggregate power of active servers and the aggregate power of sleeping servers. The total power for servers is written as

$$P_{\text{servers}} = \sum_{i=1}^{\text{NAS}} P_{\text{Server}}(i) + \sum_{j=1}^{\text{NIS}} P_{\text{sleep}}(j). \quad (1)$$

Here, NAS and NIS denote the number of active and inactive servers which are in a deep sleep state and consume only 6 W of power per server [5]. For an active server, the total power consists of the power of the processors, the power of the memory, and the power of other components. The equation is listed as follows:

$$P_{\text{Server}} = \sum_{i=1}^{\text{NS}} P_{\text{Processor}}(i) + \sum_{j=1}^{\text{NM}} P_{\text{Memory}}(j) + P_{\text{Other}} \quad (2)$$

where NS and NM are denoted as the number of sockets and the number of DIMMs in a server. To simplify the equation, we assume that all servers in the data centers have the same number of sockets and the number of DIMMs.

For the power model of the components in a server ($P_{\text{Processor}}$, P_{Memory} , and P_{Other}), we adopt the linear power model shown as follows:

$$P = (P_{\text{TDP}} - P_{\text{idle}}) * U + P_{\text{idle}} \quad (3)$$

where P_{TDP} and P_{idle} indicate the maximum power and idle power of the components while U denotes server utilization. The linear model provides sufficient accuracy for modeling the server power at the data center level [14].

The configuration of the power model in a server is shown in Table I. For processors, its idle power amounts to 10% of

TABLE I
CONFIGURATIONS OF THE SIMULATED SERVER

Server Configurations			
Part	#	TDP(w)	Idle power(w)
Processor	2	150W	15W
Memory	8	10W	5W
Others	-	124W	73.6W
P_{sleep}	6W		
Hybrid Cooling Configurations			
Parameter		Value	
$T_{\text{inlet_water}}$ (°C)		25	
$T_{\text{inlet_air}}$ (°C)		25	
V_w (GPM)		1	
η_{pump}		70%	
ΔP_w (psi)		4.2	
K_s (¢/KWH)		9	
Thermal Reliability Configurations			
θ_{CP} (°C/W)		0.3	
θ_{MP} (°C/W)		4.75	
θ_p (°C/W)		0.03	
Maintenance Cost Configurations			
Start-stop cycles for disks		40000	
CPU maintenance price (\$)		300	
Disk maintenance price (\$)		200	
Memory maintenance price (\$)		150	

the TDP [12], while four HDD hard disks are assumed to be installed in the server to fit memory-intensive applications. The specification is derived from a typical server [12].

Cooling Power Model: According to the structure of the hybrid cooling, the cooling power can be divided into two parts: the liquid power and the air cooling power

$$P_{\text{cooling}} = P_{\text{liquid_cooling}} + P_{\text{air_cooling}}. \quad (4)$$

To estimate cooling power, $E = Q/\text{COP}$ is employed, where E denotes the energy to remove the heat dissipation Q from data centers and coefficient of performance (COP) which is defined as a metric to evaluate the efficiency of a cooling system [29]. According to prior studies [5], COP_{air} (coefficient of performance) can be derived in the following equation:

$$\text{COP}_{\text{air}} = (0.0068 \times T^2 + 0.0008 \times T + 0.458)$$

where T is the inlet air temperature.

The power of liquid cooling consists of the chiller power and the pump power [19]. The chiller efficiency for a typical chilled water system is written as $\text{COP}_{\text{liquid}} = E/Q$ [7]. $\text{COP}_{\text{cooled}}$ is written in terms of inlet water temperature: $\text{COP}_{\text{liquid}} = T * 0.18 - 0.4836$ based on the specification of the water-cooled

screw compressor chiller [2]. The water pump power is calculated by the following [19]:

$$P_{\text{pump}} = N \times \frac{V_w \times \Delta P_w}{\eta_{\text{pump}}} \quad (5)$$

where N is the number of servers and V_w is the water volume flow rate. ΔP_w denotes the water-side pressure drop based on the flow resistance. Finally, η_{pump} indicates the pump efficiency.

Overall, the cooling power of the data center is calculated as follows:

$$P_{\text{cooling}} = \frac{Q_{\text{liquid cooled}}}{\text{COP}_{\text{liquid}}(T_{\text{inlet_water}}) * t} + \frac{Q_{\text{air cooled}}}{\text{COP}_{\text{air}}(T_{\text{inlet_air}}) * t} + P_{\text{pump}} \quad (6)$$

where t is a time interval during which server components dissipate the heat $Q_{\text{liquid cooled}}$ and $Q_{\text{air cooled}}$. The heat $Q_{\text{liquid cooled}}$ is removed by the liquid cooling, while the heat $Q_{\text{air cooled}}$ is generated from the other components in the servers. Shown in Table I is the configuration of hybrid cooling derived from [19]. The pump power of a server is 0.6 W and is negligible compared to the chilling power.

Overall, the electricity cost of the data center is written as

$$EC = K_{\$}(P_{\text{servers}} + P_{\text{cooling}} + P_{\text{PDL}}). \quad (7)$$

Here, $K_{\$}$ is the commercial KWH billing rate which comes to 9 cents/KWH as the default value.

B. Costs of Hardware Maintenance

Arising temperature and frequent consolidation could accelerate the aging processes of components in servers. We focus on the maintenance costs of DRAM and CPU due to the high power density. In addition, we take the cost of disk maintenance into account, since their limited number of lifetime start–stop cycles is heavily impacted by frequent server consolidations, although hard disks have a low power density.

Thermal Model: We have set up thermal models to investigate the costs of processor and memory maintenance. The CPU temperature T_C is calculated as follows from [21]:

$$T_C = T_{\text{inlet}} + (\theta_{CP} + \theta_p) * Q_C. \quad (8)$$

Here, T_{inlet} is the inlet water temperature, and Q_C is the power dissipated by the CPU. The thermal resistance of the processor package and thermal interface material layer is denoted by θ_{CP} with a value derived from [19]. The thermal resistance of the cold plate which varies with water flow is denoted by θ_p , according to the specification of Lytron CP20 cold plates [19]. Regarding the reliability issue of CPU, there is a threshold temperature for processor chips as 90 °C [19].

The temperature T_M for DRAM is given as follows:

$$T_M = T_{\text{inlet}} + (\theta_{MP} + \theta_p) * Q_{MP} \quad (9)$$

where Q_{MP} is the power dissipated by memory. The thermal resistance of the chip package of DRAM is denoted by θ_{MP} derived from [1]. There is a threshold temperature for DRAM

as 85 °C [24]. The characteristics of the thermal package of the DRAM, CPU, and cold plates are listed in Table I.

Thermal Reliability Model of Electronic Devices: We can predict the lifetimes of electronic devices based on the thermal reliability models of electronic devices. Chip temperature and power are the main factors to determine the lifetimes of electronic devices [13]. For memory, the lifetime prediction model is adopted [23]. Mean time to failure (MTTF) is widely used to represent the predicted lifetime of the electronic components for processors: $\text{MTTF} = 1/\lambda$. For the prediction of the lifetime of the processor and memory, λ is the number of failures per million hours and calculated according to Military Handbook MIL-HDBK-217F [38]

$$\lambda = (C_1 \pi_T + C_2 \pi_E) \pi_Q \pi_L \quad (10)$$

$$\pi_T = 0.1 \exp \left(\frac{-E_a}{8.617 \times 10^{-5}} \left(\frac{1}{T_p + 273} - \frac{1}{298} \right) \right). \quad (11)$$

Here, E_a is the effective activation energy (eV), and T_p is the temperature of electronic devices. The parameters ($C_1, C_2, \pi_E, \pi_L, \pi_Q$) are derived from [38]. We have scaled the lifetime of the CPU and memory according to recent studies [23]. The lifetime of the CPU is expected to be seven years when the chip temperature is 70 °C [42], while the expected lifetime of the 2-GB DRAM is five years when its temperature is 65 °C [23].

Maintenance Cost Models of the Processor and DRAM: We evaluate the costs of processors and DRAM maintenance based on their thermal reliability that is given as follows:

$$\text{RC} = \text{the cost of hardware maintenances/MTTF.}$$

For a time interval, MTTF is calculated based on their thermal reliability model with current chip temperature. The costs of a CPU, a disk, and a memory maintenance are \$300, \$200, and \$150, respectively, as shown in Table I, according to the maintenance ranging from \$300 to \$150 [6]. Based on the thermal reliability model, the cost of the CPU and memory maintenance in an active server is specified as follows:

$$\text{RC}_{\text{Server}} = \sum_{i=1}^{\text{NS}} \text{RC}_{\text{Processor}}(i) + \sum_{j=1}^{\text{DM}} \text{RC}_{\text{Memory}}(j). \quad (12)$$

Here, the costs of DRAM and CPU maintenance are increased by higher inlet water temperature. The auxiliary components are excluded from this model since they are still cooled down by air cooling. Their little heat dissipation, much lower power density, and fixed inlet air temperature result in their little cooling power and their stable maintenance cost.

Maintenance Cost Model of Hard Disk: The lifetime of hard disks is heavily impacted by server consolidations due to the limited number of lifetime start–stop cycles [15], while the impact of utilization and temperature is still unclear [34]. On the other hand, switching on/off servers incurs relatively little maintenance cost of other components such as processors and memory compared with that of hard disks. The cost of disk maintenance is computed by the following:

$$\text{RC}_{\text{Disk}} = \frac{\text{Price}}{\text{start} - \text{stop cycles}}. \quad (13)$$

As we know, the number of lifetime start–stop cycles for hard disks is 40 000 [11].

Overall, the cost of hardware maintenance of the data center is listed as follows:

$$RC = \sum_{n=1}^{ND} RC_{Disk} [NAS(t-1) - NAS(t)]^+ + \sum_{k=1}^{NAS} R_{Server}(k) [A]^+ = A \text{ if } A > 0 \text{ or } [A]^+ = 0 \text{ if } A \leq 0 \quad (14)$$

where ND and $NAS(t)$, respectively, denote the number of disks in a server and the number of active servers in the data center at time t . $[NAS(t) - NAS(t-1)]^+$ represents the number of servers which have been turned off.

Consequently, we have set up models for electricity cost and the cost of hardware maintenance to evaluate our approach which optimizes the total cost. The models have been validated with the costs of our campus data centers. We have listed all key notations for readers in Table II.

V. RENEWABLE ENERGY

We integrate renewable energies into our model such as wind power and tidal power as a supplementary energy source of data centers. The integration leads to a comparison of cost savings between wind power and tidal power. Wind power exhibits moderate variability but high unpredictability, while tide power is relatively easily predicted but varies in wider range. Based on our evaluation, wind power is more profitable than tidal power due to its relatively low variability, since its less fluctuation provides more available power to the data center with larger overlaps between the wind power and the power consumption. This comparison can help operators of data centers to select a kind of renewable energy for their data centers.

A. Wind Power

Wind power is captured by wind turbines which convert kinetic energy into mechanical energy used to produce electricity. Fig. 2 shows the output power of a typical wind turbine with respect to the wind speed [31]. The power is determined by three important wind speeds: cut-in wind speed, rated wind speed, and cutoff speed which are specific to a wind turbine. When the wind speed exceeds cut-in wind speed, the wind turbine starts to generate electricity. Its power grows as the wind speed increases until it reaches the rated wind speed. The relation between the power and the wind speed could be shown in the equation: $P = 0.5C_p\theta Av^3$, where C_p denotes the power efficiency, θ is the air density, A is the rotor swept area, and v is the wind speed. When the wind speed is between the cutoff wind speed and rated wind speed, the output power meets its maximum capacity. The power sharply drops to zero for protecting its blade assembly when the wind power exceeds the cutoff wind speed.

For most wind farm sites, the wind speed at most time is observed between the cut-in wind speed and the rated wind speed [31]. As a result, the output power is greatly sensitive to the wind speed due to their cubic relation. The resultant fluctuation of the power is shown in Fig. 3 of the wind power trace used in our experiment. Although the average power

TABLE II
KEY NOTATION

Notation	Definition
$P_{servers}$	The power consumption of all servers
P_{PDL}	The power distribution loss
$P_{Server}(i)$	The power consumption of an active server at time i
$P_{sleep}(i)$	The power consumption of a sleeping server at time i
$P_{Processor}(i)$	The power consumption of a processor at time i
$P_{Memory}(i)$	The power consumption of a memory module at time i
P_{Other}	The power consumption of auxiliary components in a server memory module at time i
NAS	The number of active servers
NIS	The number of sleeping servers
NS	The number of sockets in a server
NM	The number of DIMMs in a server
MINS	The minimal demanded number of active servers
K_s	commercial KWH Billing Rate
TC	The total costs
$P_{cooling}$	The cooling power
$P_{liquid_cooling}$	The liquid cooling power
$P_{air_cooling}$	The air cooling power
RC_{Server}	The maintenance cost for a server related to processors and memory
$RC_{Processor}(i)$	The maintenance cost for a processor
$RC_{Memory}(i)$	The maintenance cost for a memory
RC_{Disk}	The maintenance cost for a disk
ND	The number of disks in a server
T_C	The working temperature of a processor
T_M	The working temperature of a memory
T_{inlet_water}	The temperature of inlet water used to cool down processors and memory

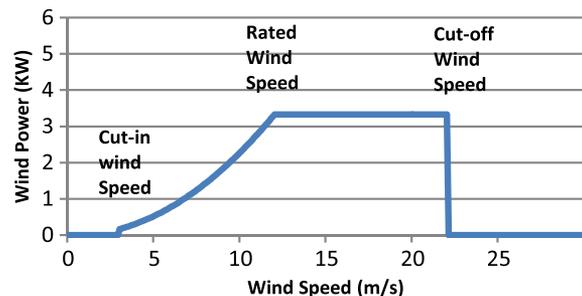


Fig. 2. Relationship between wind speed and power.

demand derived from Saskatchewan-HTTP trace is approximate to the total wind power in the example, a considerable mismatch is expected due to their unrelated factors for their fluctuation: diary human activities and local weather condition.

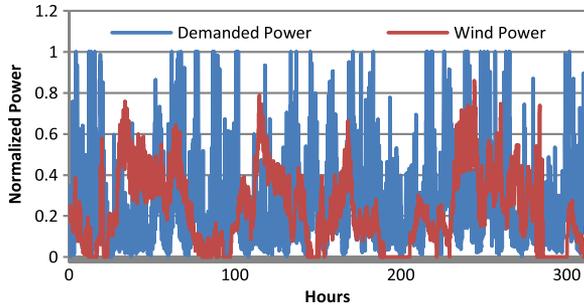


Fig. 3. Mismatch between wind power and power consumption of the data center.

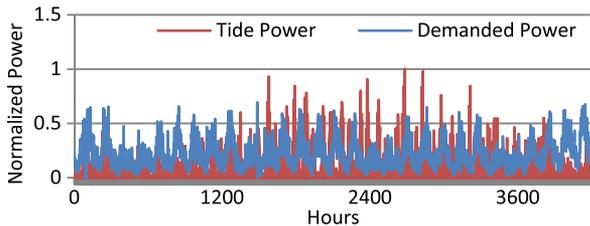


Fig. 4. Mismatch between tide power and power consumption of a data center.

This mismatch leads to low wind power usage or requires a huge capacity of energy storage to reshape the wind power. However, the energy storage incurs additional capital costs and wastes wind energy, since required batteries are considerably expensive and waste energy due to energy conversions. When wind power is used as a supplementary energy resource for the data center, a conventional power grid also powers the data center unless its demand is less than the wind power.

B. Tidal Power

Similar to wind power generated from the kinetic energy of air flow, tidal power is produced from a tidal stream of sea water which is relatively predictable due to a known tide table. This advantage potentially increases the usage of tidal power for data centers, since data centers have to reserve considerable energy for unpredictable wind power. It also lowers the capacity of energy storage in data centers which incurs noticeable capital cost, which further reduces the total cost for data centers. On the other hand, employing tidal power for data centers is hindered by its considerable variance and pattern which is unrelated with human activities shown in Fig. 4. They pose a serious challenge on boosting the usage of tidal power, since unbearable amount of energy storage is required to sync tidal power and demanded power from data centers. Alternatively, the usage of tidal power can be increased in the design: when tidal power surpasses demanded power, the exceeded power can be used to reduce hardware maintenance cost by maintaining lower temperature of cooling water and decreasing the number of server consolidation. Otherwise, the usage of tidal power is saturated since it is fully used by data centers.

The power availability of a tidal stream with a tidal velocity V is estimated by using the equation: $P_{\text{tide}} = (1/2)\rho AV^3$, where ρ and A stand for water density and the area swept by rotor blades, respectively [17]. The power generated from a tidal stream generator can be estimated by using the equation:

$P_m = C_p P_{\text{tide}}$, where C_p represents the efficiency of conversion from kinetic energy into electrical energy [22].

VI. COST OPTIMIZATION IN DATA CENTERS

We formulate the total cost in (15) based on (7) and (14) with the constraints. Focusing on the operational cost of data centers, we pick up a typical specification for our heuristic data center shown in Table I. There are two important decision variables $T_{\text{inlet_water}}$ and NAS, while other variables are determined by available servers, server performance, and characteristics of traces, which are also treated as parameters. For example, NS denotes the total number of servers, while MINS denotes the minimum required number of active servers which is determined by traces. Our objective is to minimize the total cost with the constraints

$$\min \left\{ \text{TC} = \sum_{n=1}^{\text{ND}} \text{RC}_{\text{Disk}} * [\text{NAS}(t-1) - \text{NAS}(t)]^+ + \sum_{i=1}^{\text{NAS}} \text{RC}_{\text{Server}}(i) + K_{\$} * (P_{\text{servers}} + P_{\text{cooling}} + P_{\text{PDL}}) \right\} \quad (15)$$

subject to

$$T_C \leq 90^\circ\text{C} \text{ and } T_M \leq 85^\circ\text{C} \text{ MINS} \leq \text{NAS} \leq \text{NS}.$$

The space of feasible solutions of this discrete optimization is too large, resulting in that exhaustively searching the global optimal solution is impossible. To optimize the total cost of electricity and hardware maintenance, we proposed to trace the local optimal solution by dynamically manipulating $T_{\text{inlet_water}}$ and NAS corresponding to the fluctuation of workloads.

$T_{\text{inlet_water}}$ and NAS are tuned simultaneously to minimize the total cost since they interact with each other. An optimal $T_{\text{inlet_water}}$ can be found for the given average utilization of servers to minimize the cooling and maintenance costs of the processors and memory modules. The optimal $T_{\text{inlet_water}}$ is affected by NAS which determines the average server utilization for a given amount of user requests. On other hand, an optimal NAS can be derived for the given average server utilization and $T_{\text{inlet_water}}$ to minimize maintenance and electricity costs.

A. Overview of the Cost Optimization System

For the manipulation of T_{inlet} and NAS, we proposed a structure shown in Fig. 5. In this structure, there are four modules: Workload Prediction, Server Monitor, Server Manager, and Thermal Manger, working together to reduce the total cost. The workload prediction collects request history and predicts future request trend and the future minimum required number of active servers. The server monitor collects the temperature and utilization information of servers and estimates the cost of hardware maintenance. The thermal manager generates $T_{\text{inlet_water}}$ based on the average server utilization, while the server manager selects NAS and the resultant average server utilization according to the predicted future minimum required number of active servers and $T_{\text{inlet_water}}$. $T_{\text{inlet_water}}$ and NAS are selected to minimize the total cost via the interactions of the two managers.

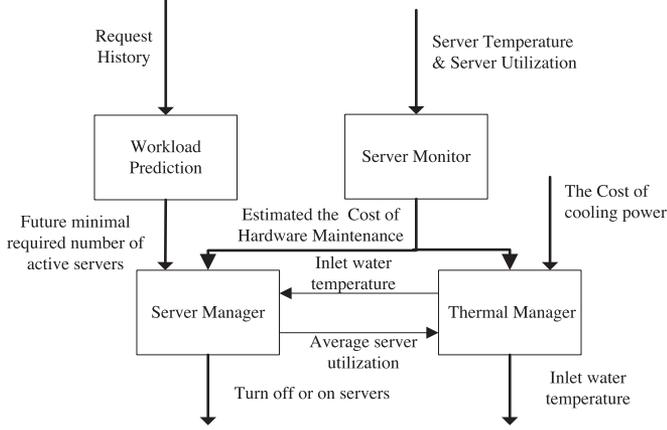


Fig. 5. Overview of the cost optimization system.

B. Optimal Inlet Water Temperature

To investigate the impact of the inlet water temperature on the total cost, we divide the total cost into two parts: the cost of cooling power and CPU and memory maintenance which are affected by the inlet water temperature, and the other costs which are unaffected denoted by C

$$TC = K_{\$} * P_{cooling} + \sum_{i=1}^{NAS} RC_{Server}(i) + C. \quad (16)$$

As the inlet water temperature increases, $P_{cooling}$ decreases based on the function of COP, while RC_{Server} increases according to (8)–(12). Note that $P_{cooling}$ and RC_{Server} also depend on the utilization of each server. The optimal configuration can be obtained for a given utilization.

C. Adjusting the Number of Active Servers

The other substantial variable NAS is facilitated by server consolidation which lively migrates jobs cross servers, with the upper bound of available servers and the lower bound of service level agreement. Generally, the server manager weighs the one-time costs and time-dependent benefits for turning off a server. They are investigated as follows.

One-Time Costs of Server Consolidation: It is well known that server consolidation could save the electricity cost. Unfortunately, it increases the cost of disk maintenance, according to (14). Furthermore, servers waste energy during the transitions between the active state and the sleeping state. We formulate the cost for server consolidation denoted by C_{cs} . The cost C_{cs} per server is calculated as follows:

$$C_{cs} = \sum_{j=1}^{ND} RC_{Disk} + P_{max} * T_T * K_{\$} \quad (17)$$

where T_T is the time of the two transitions (switching from active to sleep and back) including two job migrations (20 s for one [11]) and two transitions between the active state and the sleeping state (5 s for ACPI S3 state [26]). Therefore, T_T is estimated to be 50 s, which is relatively small compared with the 5 min it takes to change the state of a server in our

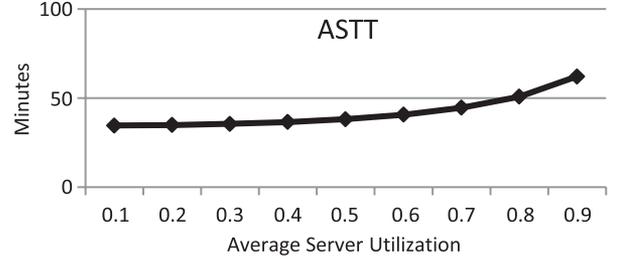


Fig. 6. Varying ASTT with the average server utilization.

```
//NAS : the Current Number of Active Servers
//FMRNAS: Future Minimum Required Number of Active Servers
if NAS < FMRNAS [T]
NAS = FMRNAS [T]
Else
// Turn off servers
While NAS > Max(FMRNAS [T,T+ASTT]){
//turn off a server
NAS= NAS-1
Recalculate ASTT
}
```

Fig. 7. Algorithm based on varying ASTT-ANN.

experiment. P_{max} and $K_{\$}$, respectively, represent the maximum power for a server and the commercial KWH billing rate.

Time-Dependent Benefits of Server Consolidation: The reward of server consolidation depends on the length of server sleeping time once turning off. In other words, the benefit is determined by the length of the period of turning off servers without violation of user level agreement. The length of this period is referred to as available sleeping time (AST), which indicates the maximal server sleeping time. Ideally, the benefits of turning off a server should be calculated as the integral of $B_{sleeping}$ during its AST, where $B_{sleeping}$ denotes the benefit of turning off a server for a minute, which changes with the average server utilization of servers. The precise value of $B_{sleeping}$ is difficult to capture and depends on the number of active servers in the following intervals. We estimate the benefit of turning off a server as $B_{sleeping} \times AST$. Here, $B_{sleeping}$ can be obtained from the derivative of (15) excluding the maintenance cost of the disk at the current interval with respect to NAS.

ASTT: To weight the costs and benefits, we define the AST threshold (ASTT) as follows:

$$ASTT = \frac{C_{cs}}{B_{sleeping}}. \quad (18)$$

When the available sleep time of a server is longer than ASTT, a server should be turned off. $B_{sleeping}$ and ASTT are recalculated since the increasing average server utilization changes the optimal inlet water temperature and $B_{sleeping}$. This step will be repeated unless there is no benefit to turning off a server. Based on the parameters, we obtain the value of ASTT with the average server utilizations of active servers shown in Fig. 6; the ASTT should be approximately 50 min when our model excludes the cooling cost and the maintenance costs of the CPU and memory modules.

ASTT-ANN: Varying ASTT Based on Artificial Neural Network: We design an algorithm shown in Fig. 7 based on the concept of ASTT. Generally, the algorithm turns off/on

servers based on the estimations of the costs and benefits, and the minimum number of active servers in the following intervals. The minimum number of active servers is predicted using artificial neural networks [20]. The array of artificial neural networks is employed to predict the minimum number of active servers from 5 min ahead (predicted minimum number of servers $[T]$) to 65 min ahead (maximal predicted minimum number of servers $[T : T + 60]$) at intervals of 5 min. The prediction for 5 min ahead is used to guide the system to turn on more servers to satisfy the QoS requirement, while the others are used to decide the number of active servers which should be turned off as shown in Fig. 7. We use 50% of the data set from each trace to tune the weights of the artificial neural networks using offline training. Each artificial neural network is a feedforward neural network with one hidden layer and a quadratic cost function. To improve the QoS, we modify the quadratic cost function of the feedforward neural network which predicts the minimum number of active servers for 5 min ahead, since any deficit of the predicted value leads to violence in the QoS. The modified cost function is presented as follows:

$$\text{Cost} = \sum (\hat{y} - y)^2 * a$$

$$a = 100 \text{ if } \hat{y} < y; a = 1 \text{ if } \hat{y} \geq y \quad (19)$$

where Cost denotes the total cost while \hat{y} and y stand for the predicted value and the real value, respectively. a is introduced to penalize cases where the predicted value is less than the real value. Generally, large a values improve the QoS but reserve more servers, thereby increasing the total cost. The value of a is chosen based on our empirical results to balance the improvement of QoS and the overhead of power consumption. In the following section, the model of a data center is built to quantitatively evaluate the benefit of sweet inlet water temperature and varying ASTT.

D. Co-Optimization With Wind Power or Tidal Power

For the unreliable renewable energies such as wind and tidal power, the proposed optimization is designed to increase its benefit. Rather than merely targeting at electricity costs, the optimization reduces the server maintenance costs at the expense of increased power consumption. The cost of such overhead could be avoided when the renewable power is larger than the electrical demand of data centers. It could be explained by the modified objective

$$\min \left\{ \text{TC} = \sum_{n=1}^{\text{ND}} \text{RC}_{\text{Disk}} * [\text{NAS}(t-1) - \text{NAS}(t)]^+ \right. \\ \left. + \sum_{i=1}^{\text{NAS}} \text{RC}_{\text{Server}}(i) + K_{\S} * (P_{\text{servers}} + P_{\text{cooling}} \right. \\ \left. + P_{\text{PDL}} - P_{\text{Renewable}}) \right\} \quad (20)$$

where $P_{\text{Renewable}}$ denotes the renewable power at time t , which can be wind power or tide power. There are two scenarios regarding the comparison between the renewable power and the power demand of data centers.

```
//Predictor
//NAS : the Current Number of Active Servers
//FMRNAS: Future Minimum Required Number of Active Servers

If Renewable Power > power Consumption & Predictor <M
Predictor = Predictor + 1
If Renewable Power < power Consumption & Predictor >0
Predictor = Predictor - 1

if NAS < FMRNAS [T]
NAS = FMRNAS [T]
Else
// Turn off servers
while NAS > Max(FMRNAS [T,T+ ASTT])&& (Predictor<M/2){
// Turn off a server
NAS= NAS-1
Recalculate ASTT
}
```

Fig. 8. Algorithm of renewable power ASTT-ANN.

- 1) $P_{\text{Renewable}} \geq (P_{\text{servers}} + P_{\text{cooling}} + P_{\text{PDL}})$: Power over sufficient period (POS period). With over sufficient renewable power, the only concern of this optimization is to reduce the cost of server maintenance costs by lowering the inlet water temperature and stopping turning off active servers. The power consumption of data centers could be increased as long as it is less than the renewable power.
- 2) $P_{\text{Renewable}} < (P_{\text{servers}} + P_{\text{cooling}} + P_{\text{PDL}})$: Power insufficient period (PI period). When the renewable power partially compensates the power consumption of data centers, ASTT-ANN can reduce the electricity costs and server maintenance costs together by adjusting the inlet water temperature and the number of active servers. Since the derivative of the total cost in their factor is not affected by the renewable power, our method still reach the optimal point to minimize the total costs at each interval.

Disk Replacement Cost: Predicting the comparison between the renewable power and the power demand in the following intervals is substantial to reduce disk replacement costs by exploiting the benefit of the renewable power. The disk replacement cost is amortized over the saving of the electricity costs in the server sleeping time. The saving could be reduced if the sleeping time includes some POS periods. Consequently, the longer AST is demanded to compensate the disk replacement cost, since electricity saving can only be gained in the PI periods. The portion of the POS periods in the following time becomes the key to reduce disk replacement cost with the renewable power. To further reduce disk replacement cost, we design a POS predictor which is similar to the classical CPU branch predictor.

Wind Power ASTT-ANN: ASTT-ANN as well as sweet temperature is extended to fully exploit the benefit of the wind power based on the aforementioned discussion. The optimization of sweet temperature is intuitive; the inlet water temperature tracks the optimal value to balance the CPU and memory replacement costs in PI periods; otherwise, it is fixed at the lowest temperature to minimize the server maintenance cost. The modified ASTT-ANN also shows distinct policies in different periods to minimize the electricity cost and the replacement costs of disks shown in Fig. 8. During POS periods, turning off

active servers is prohibited to avoid incurred replacement cost; otherwise, the original ASTT-ANN still works. For capturing the immediately following POS period, we design a predictor based on the recent history, which is widely used in CPU branch prediction in Fig. 8. M is chosen to be 8, since we discovered that it is the optimal value for our five traces. This modified co-optimization is referred to as Wind Power ASTT-ANN (WP-ASTT-ANN), which reduces electricity and server maintenance costs by utilizing the wind power.

Tidal Power ASTT-ANN: Sharing the same underlying idea, the co-optimization with tidal power known as TP-ASTT-ANN exploits the benefit of tidal power during POS periods. The benefit of the exploitation relies on how to accurately predict the length of a POS period in the following intervals, which can guide the algorithm to reduce unnecessary server consolidations and thus the total cost. In contrast to wind power, the prediction of tidal power is more accurate since it exhibits a relatively predictable pattern which is forecasted based on the predicted sea level in our experiment. On the other hand, the demanded power is still unforeseeable and predicted based on its history as WP-ASTT-ANN.

VII. EXPERIMENTAL SETUP

A. Data Center

Recalling the models related to the costs of electricity and hardware maintenance, we combined them with server performance model and real traces to simulate our prototype data center which consists of 1024 servers cooled by hybrid cooling.

Server Performance Model and Response Time Analysis: We assume that a server in our data center provides a 2.6-GB/s service rate and the mean of the response times should be bound by 6 ms for SLA [11]. To calculate the FMRNAS at a time interval, we use the GI/G/m model [8] to determine how many servers can satisfy a demand based on the following equation:

$$\begin{aligned} \widetilde{W} &= \frac{1}{\mu} + \frac{P_m}{\mu(1-\rho)} * \left(\frac{C_A^2 + C_B^2}{2m} \right) \\ P_m &= \rho^{\frac{m+1}{2}} \text{ if } \rho \leq 0.7 \\ P_m &= \frac{\rho^m + \rho}{2} \text{ if } \rho > 0.7 \end{aligned} \quad (21)$$

where \widetilde{W} is the mean response time. $1/\mu$ is the mean service time of a server. $\rho = \lambda\varphi/mf$ is the average utilization of servers. λ, φ, C_A , and C_B are derived from trace characteristics [5]. We use this performance server and response time model to acquire the minimum required number of active servers at every time slot. For a time interval, we choose 5 min as the minus unit [5].

B. Traces

We use five traces downloaded from the Internet traffic Archive [43]: Clarknet-HTTP, NASA-HTTP, Saskatchewan-HTTP, UC Berkeley IP, and WorldCup. Their lengths range from 14 to 30 days, and all of the trace files cover several peak requests. We have scaled the traces to meet our data center performance.

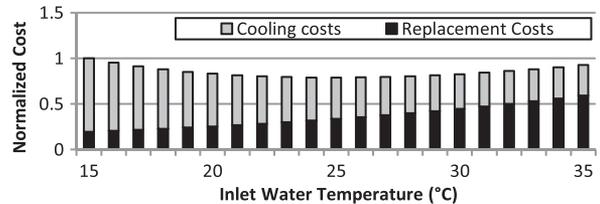


Fig. 9. Impact of inlet water temperature on the costs of cooling power and hardware maintenance.

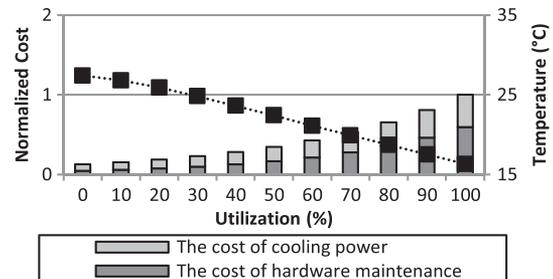


Fig. 10. Variation of sweet temperature and these costs corresponding to the utilization of the data center.

C. Renewable Power Trace

We calculated the wind power based on the relation between the wind speed and the output power of wind turbines [31], with the specific parameters such as power efficiency from [3]. The 14-day wind speed trace is derived from [4]. Since the average power consumption cross Web traces are different due to their distinct patterns, we scale the average wind power to match the average power consumption for data centers for each trace. To scrutinize the benefit of our optimization, the average wind power is scaled to 50%, 75%, 100%, 125%, and 150% of the average power demand in each trace, which are referred to as 50%, 75%, 100%, 125%, 150% wind power (WP). On the other hand, we derive 14 traces of tide power from [30]. Similarly, the average tidal power is scaled to the same portions of the average power, which are presented as 50%, 75%, 100%, 125%, and 150% tidal power (TP). The extra power for data centers comes from the conventional power grid when the renewable power is less than the power demand.

VIII. RESULTS

In this section, we compare our optimization with other suboptimal solutions to reflect our potential benefit in the experiments. For example, the aggressive server consolidation (ASTT = 5 min) could be considered as a typical case which prior works use to reduce electricity cost. Additionally, warmer cooling water might be a good example to demonstrate that prior works reduce cooling cost without the awareness of hardware maintenance cost.

A. Optimization Based on Sweet Temperature

As illustrated in (16), when the server power is fixed, the total cost is only related to cooling and hardware maintenance. Fig. 9 illustrates the impact of the inlet water temperature changing from 15 °C to 35 °C on the cooling cost and the

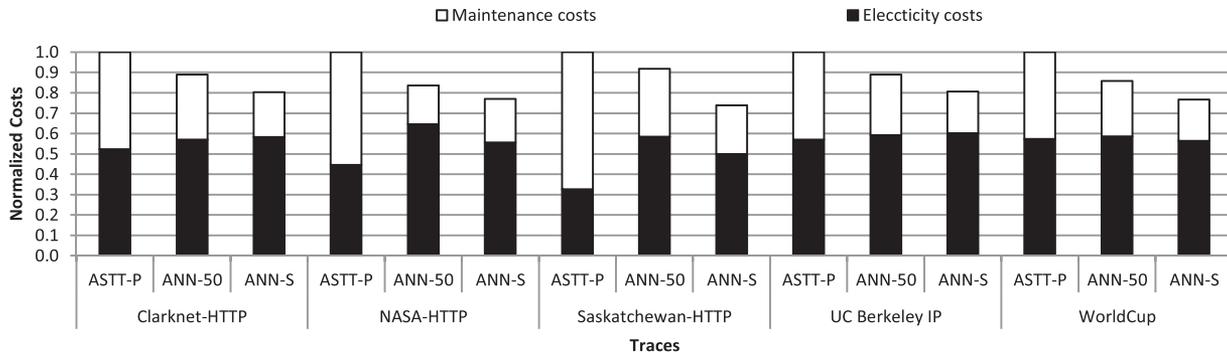


Fig. 11. Breakdown costs of ASTT-P with ASTT of 5 min and fixed inlet water temperature of 25 °C (ASTT-P), ASTT-ANN with ASTT fixed at 50 min and fixed inlet water temperature of 25 °C (ANN-50), and ASTT-ANN with varying ASTT and sweet temperature (ANN-S) in five traces.

cost of hardware maintenance of our data center with 30% utilization. These costs are normalized against the total costs when the inlet water temperature is 15 °C. Increasing inlet water temperature reduces cooling power especially when the temperature is below 25 °C. However, a high inlet water temperature increases the cost of the hardware maintenance of the CPU and memory. As observed from Fig. 8, we can find an optimal inlet water temperature (25 °C in this case) which minimizes the total cost when utilization is fixed at 30%. In the following context, we will refer the sweet temperature to the optimal inlet water temperature. This observation justifies that a high inlet water temperature is reasonable in data centers when the current average server utilization is low (below 30%). Otherwise, a high inlet water temperature could hurt the cost of hardware maintenance during the high utilization.

Fig. 10 shows the cooling and hardware maintenance costs of our data center when the average server utilization varies from 0% (all servers are powered on with no workload) to 100% (all servers are powered on with workloads). The right vertical axis of the figure illustrates sweet temperatures for different utilizations. In the figure, the total costs for all utilizations are the lowest for the data center cooled by water at corresponding sweet temperatures. When the utilization of the data center is low, warm inlet water temperature offers more benefit since the cost of cooling power is larger than the cost of hardware maintenance (e.g., in our simulation result, the cost of cooling power is 1.65 times the cost of hardware maintenance when the utilization is 10%). On the other hand, as the data center utilization increases, we must keep a cold chilling water to cool down the heating hardware and slow the growth of hardware maintenance especially when their temperatures are close to the critical temperatures. Consequently, to minimize the total costs, inlet water temperature should be dynamically adjusted according to the data center utilization.

B. Optimization of Electricity Costs and Hardware Maintenance Costs Based on ASTT

We compare three optimizations here: one which only minimizes the electricity costs and is presented by ASTT-P (ASTT = 5 min) with an inlet water temperature of 25 °C, one which optimizes the electricity costs and disk maintenance cost and is demonstrated by ASTT-ANN (ASTT = 50 min) with an inlet water temperature of 25 °C, and one which minimizes

the electricity costs and hardware maintenance costs and is exhibited by ASTT-ANN with varying ASTT and sweet inlet water temperature.

Fig. 11 shows the breakdown costs of the three optimizations for five traces. All of the costs are normalized against the total costs of ASTT-P (ASTT = 5 min) with an inlet water temperature of 25 °C. The baseline uses perfect prediction instead of artificial neural networks to avoid the considerable costs of prediction errors, since this overhead may exaggerate the benefits of the other two optimizations. The baseline achieves the lowest electricity costs for the five traces with the highest total costs. Electricity costs are minimized by adjusting the number of active servers at the edge of the QoS requirements without awareness of hardware maintenance costs, since servers are aggressively turned on and off at the minimal 5-min ASTT.

Compared to the baseline, ASTT-ANN (ASTT = 50 min) and an inlet water temperature of 25 °C achieve a normalized total cost of 0.87 (geometric mean of the five benchmarks). This configuration reduces disk maintenance costs considerably. The high inlet water temperature reduces the cooling costs but increases the maintenance costs of the CPU and memory modules. Here, 50-min ASTT is the optimal ASTT if the models only include the disk maintenance cost but exclude the maintenance costs of the CPU and memory modules.

Finally, ASTT-ANN with varying ASTT and sweet inlet water temperature yields a normalized total cost of 0.77 (geometric mean of the five benchmarks) via minimizing the sum of electricity and hardware maintenance costs. It consistently reduces the total costs for the five traces compared to the other two optimizations.

We also show the success request ratios and average prediction errors to evaluate the performance of our predictions. To prove that the ASTT-ANN with varying ASTT can guarantee the QoS, Fig. 12 shows the request success ratio which stands for the percentages of requests serviced under the constraint of the QoS. All of the ratios are above 96%, which satisfies our goal. The QoS can be protected by the modified cost function in (19) even if the prediction performance is fair. The cost function leads the systems to reserve servers by increasing the probability that the predicted value exceeds the real value.

The average prediction errors and the probabilities of overestimation for 5 min ahead and 50 min ahead are shown in Fig. 13. The employed artificial neural networks predict the maximal and minimum number of active servers from 5 min ahead to

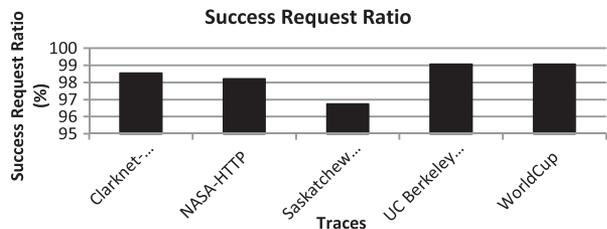


Fig. 12. Request success ratio of ASTT-ANN with varying ASTT.

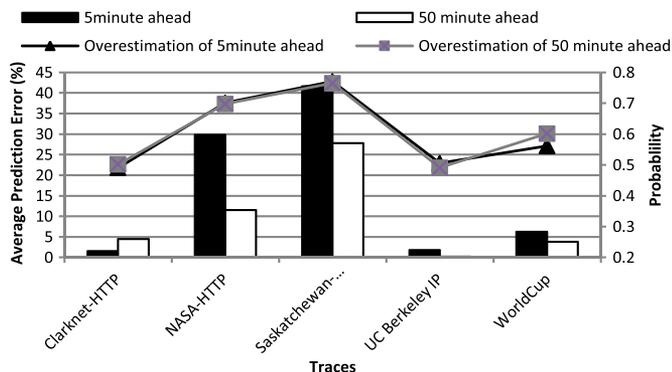


Fig. 13. Average prediction error and overestimation probabilities of the minimum number of active servers for 5 min ahead and 50 min ahead.

65 min ahead. We pick up two of them to show their prediction accuracies. The prediction errors are bounded within 10% for the three traces, except NASA-HTTP and Saskatchewan-HTTP which show intense fluctuations of the minimum number of active servers. Since underestimation is penalized in (19), the predictor is more eager to overestimate the required number of servers when prediction errors have been large. This leads the optimization to reserve more servers and thereby maintain a high QoS at the cost of more energy consumption. The higher probabilities of overestimation are observed in Fig. 13 for the two traces. Although the performances of predictions are fair for the traces, ASTT-ANN with varying ASTT can still yield considerable benefits for them.

C. WP-ASTT-ANN

The benefit of WP-ASTT-ANN is revealed by the comparison between Figs. 14 and 15. Fig. 14 shows the normalized costs in five traces of the simulated data center powered by 50% WP, 75% WP, 100% WP, 125% WP, and 150% WP with the baseline which merely targets electricity costs. The total costs are normalized against those of the baseline without the wind power. The shrinking marginal profit of increasing the wind power could be observed from that the total costs of 50% WP, 75% WP, 100% WP, 125% WP, and 150% WP are 0.77, 0.7, 0.66, 0.63, and 0.61 in geometric mean, respectively. This trend is confirmed by the results of five traces. Fig. 15 also shows this normalized costs but with WP-ASTT-ANN. The similar decrease of the marginal profit could be observed from that the total costs of 50% WP, 75% WP, 100% WP, 125% WP, and 150% WP are 0.62, 0.53, 0.48, 0.44, and 0.40 in geometric mean, respectively. The benefit of WP-ASTT-ANN grows as the wind power increases based on the facts that with 50% WP,

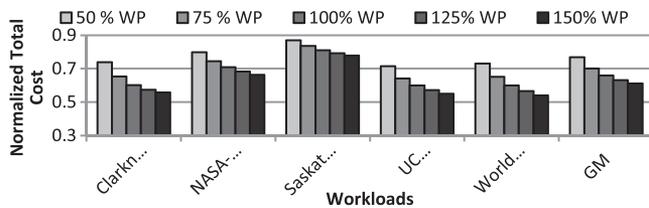


Fig. 14. Normalized costs in five traces of the simulated data center powered by 50% WP, 75% WP, 100% WP, 125% WP, and 150% WP.

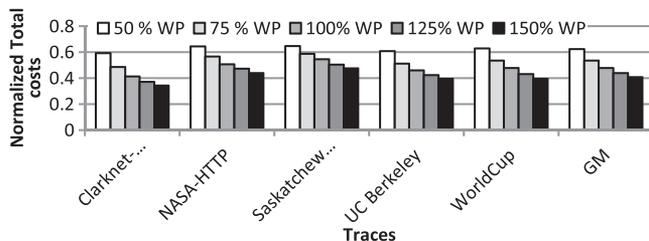


Fig. 15. Normalized costs in five traces of the simulated data center powered by 50% WP, 75% WP, 100% WP, 125% WP, and 150% WP and optimized by WP-ASTT-ANN.

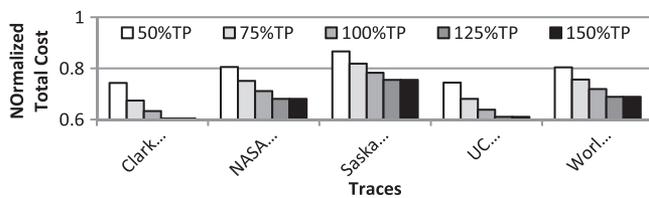


Fig. 16. Total cost normalized against the cost without tidal energy contributed by 50% TP, 75% TP, 100% TP, 125% TP, and 150% TP.

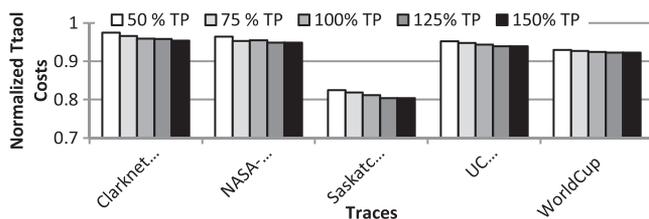


Fig. 17. Reduced total cost normalized against the cost without TP-ASTT-ANN contributed by 50% TP, 75% TP, 100% TP, 125% TP, and 150% TP.

75% WP, 100% WP, 125% WP, and 150% WP are 0.15, 0.17, 0.18, 0.19, and 0.21 compared with Fig. 14.

D. TP-ASTT-ANN

The total cost is reduced when the amount of tidal power is increased from 50% to 150%, shown in Fig. 16. Similar to wind power, the marginal benefit shrinks quickly, and the total cost is steady when the amount of tidal power exceeds 125% demanded power. For example, the total cost of Saskatchewan-HTTP is only reduced by 25% total cost even if the average tidal power is 1.5 times the demanded power. It is caused by the larger variance of tidal energy shaped by the orbit of the Earth–Moon system. The low usage can be moderately mitigated by the TP-ASTT-ANN which aims at boosting the usage of tidal power shown in Fig. 17. It yields a moderate benefit for all traces, except Clarknet-HTTP which

exhibits a relatively large variance and leads to low accuracy of predicting subsequent power demand. Additionally, tidal power hardly provides long POS periods compared to wind power, since its frequent limited availability results from its noticeable variance. This drawback of tidal energy limits the benefit of TP-ASTT-ANN compared with WP-ASTT-ANN.

IX. CONCLUSION

The quick growth of electricity bill drives owners of data centers to employ server consolidation and the high temperature of data center. However, the traditional air cooling system offers limited benefit of these two approaches due to its low energy efficiency of cooling power especially. We have built a comprehensive framework which covers the costs of server power, cooling power, and hardware maintenance. Based on the models, we introduce a joint optimization of the costs of electricity and server maintenance. The approach gains 23% savings of the total cost and guarantees the response time of more than 96% requests. In the future, our framework will incorporate elaborated reliability models for state-of-the-art servers and power managements which are also important for minimizing costs of data center owners.

REFERENCES

- [1] DRR3 DRAM. Retrieved from Sep. 29, 2013. [Online]. Available: <http://www.micron.com/products/dram>
- [2] Catalog of Water-Cooled Screw Compressor Chillers. Retrieved from Sep. 29, 2013. [Online]. Available: <http://chillertrader.com/manufacturer/McQuay>
- [3] 2.5 MW Series Wind Turbine. (n.d.). Retrieved from Sep. 29, 2013. [Online]. Available: http://site.genengy.com/prod_serv/products/wind_turbines/en/2xmw/index.htm
- [4] Center for Operational Oceanographic Products and Services. Retrieved from Sep. 29, 2013. [Online]. Available: <http://tidesandcurrents.noaa.gov>
- [5] F. Ahmad and T. N. Vijaykumar, "Joint optimization of idle and cooling power in data centers while maintaining response time," *SIGPLAN Notices*, vol. 45, no. 3, pp. 243–256, Mar. 2010.
- [6] L. A. Barroso and U. Hözl, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. San Rafael, CA, USA: Morgan and Claypool, 2009.
- [7] M. H. Beitelmal and C. D. Patel, "Model-based approach for optimizing a data center centralized cooling system," Hewlett Packard Lab., Palo Alto, CA, USA, Tech. Rep. No. HPL-2006-67, 2006.
- [8] G. Bolch *et al.*, *Queueing Networks and Markov Chains: Modeling and Performance Evaluation With Computer Science Applications*. New York, NY, USA: Wiley, 1998.
- [9] K. Brill, *The Economic Meltdown of Moore's Law and the Green Data Center*. Uptime Inst., Seattle, WA, USA, 2007.
- [10] "Summer Time Energy-Saving Tips for Businesses," California Energy Commission, Sacramento, CA, USA, Retrieved from Sep. 29, 2013. [Online]. Available: consumerenergycenter.org/tips/business/summer.html
- [11] Y. Chen *et al.*, "Managing server energy and operational cost in hosting centers," in *Proc. SIGMETRICS*, 2005, pp. 303–314.
- [12] J. Chang *et al.*, "Green server design: Beyond operational energy to sustainability," in *Proc. Int. Conf. HotPower*, 2010, pp. 1–8.
- [13] "Guidelines to Understanding Reliability Prediction," Eur. Power Supply Manuf. Assoc., Northants, U.K., Retrieved from Sep. 29, 2013. [Online]. Available: http://www.epsma.org/pdf/MTBF%20Report_24%20June%202005.pdf
- [14] X. Fan *et al.*, "Power provisioning for a warehouse-sized computer," in *Proc. ISCA*, 2007, pp. 13–23.
- [15] J. G. Elerath, "Specifying reliability in the disk drive industry: No more MTBF," in *Proc. Annu. Rel. Maintain. Symp.*, 2000, pp. 194–199.
- [16] M. Elnozahy *et al.*, "Energy-efficient server clusters," in *Proc. 2nd Workshop Power Aware Comput. Syst.*, Feb., 2002, pp. 179–197.
- [17] J. Hardisty, *The Analysis of Tidal Stream Power*. Chichester, U.K.: Wiley, 2009.
- [18] W. Huang *et al.*, "TAPO: Thermal-aware power optimization techniques for servers and data centers," in *Proc. IGCC*, 2011, pp. 1–8.
- [19] D. C. Hwang *et al.*, "Energy savings achievable through liquid cooling: A rack level case study," in *Proc. ITHERM*, Jun. 2010, pp. 1–9.
- [20] J. Heaton, "Introduction to the Math of Neural Networks," Heaton Res., Inc., Chesterfield, MO, USA, Apr. 3, 2012.
- [21] Intel® Core™2 Duo Processor E8000 and E7000 Series and Intel® Pentium® Processor E5000 Series Thermal and Mechanical Design Guidelines. Retrieved from Sep. 29, 2013.
- [22] M. Kuschke *et al.*, "Modeling of tidal energy conversion systems for smart grid operation," in *Proc. IEEE Power Energy Soc. Gen. Meet.*, Detroit, MI, USA, 2011, pp. 1–2.
- [23] S. Li *et al.*, "System implications of memory reliability in exascale computing," in *Proc. High Perform. Comput., Netw., SC*, Nov. 12–18, 2011, pp. 1–12.
- [24] J. Lin *et al.*, "Thermal modeling and management of DRAM memory systems," in *Proc. ISCA*, 2007, pp. 312–322.
- [25] M. Lin *et al.*, "Dynamic right-sizing for power-proportional data centers," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 1098–1106.
- [26] Linux Documentation. latest release 2.6, Retrieved from Sep. 29, 2013. [Online]. Available: <http://www.kernel.org/doc/documentation/power/states.txt>
- [27] D. Meisner *et al.*, "PowerNap: Eliminating server idle power," in *Proc. ASPLOS*, 2009, pp. 205–216.
- [28] D. Meisner and T. F. Wenisch, "DreamWeaver: Architectural support for deep sleep," in *Proc. ASPLOS*, 2012, pp. 313–324.
- [29] J. Moore *et al.*, "Making scheduling 'cool': Temperature-aware workload placement in data centers," in *Proc. USENIX*, 2005, p. 5.
- [30] National Oceanic and Atmospheric Administration Retrieved from Jun. 6, 2014. [Online]. Available: <http://tidesandcurrents.noaa.gov/>
- [31] M. Patel, *Wind and Solar Power Systems*. Boca Raton, FL, USA: CRC Press, 1999.
- [32] M. K. Patterson, "The effect of data center temperature on energy efficiency," in *Proc. IOTHERM*, 2008, May 2008, pp. 1167–1174.
- [33] S. Pelley *et al.*, "Understanding and abstracting total data center power," in *Proc. WEED*, 2009, pp. 1–6.
- [34] E. Pinheiro *et al.*, "Failure trends in a large disk drive population," in *Proc. FAST*, 2007, p. 2.
- [35] A. Qouneh *et al.*, "A quantitative analysis of cooling power in container-based data centers," in *Proc. IISWC*, 2011, pp. 61–71.
- [36] P. Ranganathan *et al.*, "Ensemble-level power management for dense blade servers," in *Proc. ISCA*, 2006, pp. 66–77.
- [37] R. Raghavendra *et al.*, "No 'power' struggles: Coordinated multi-level power management for the data center," *Archit. Support Programm. Languages Oper. Syst.*, vol. 36 no. 1, pp. 48–59, Mar. 2008.
- [38] Reliability Prediction of Electronic Equipment. Military Handbook. Retrieved from Sep. 29, 2013. [Online]. Available: <http://snebulos.mit.edu/projects/reference/MIL-STD/MIL-HDBK-217F-Notice2.pdf>
- [39] B. A. Rubenstein *et al.*, "Hybrid cooled data center using above ambient liquid cooling," in *Proc. ITHERM*, Jun. 2–5, 2010, pp. 1–10.
- [40] B. Schroeder and G. A. Gibson, "Disk failures in the real world: What does an MTTF of 1,000,000 hours mean to you?" in *Proc. FAST*, 2007, pp. 1–16.
- [41] S. Srikantaiah *et al.*, "Energy aware consolidation for cloud computing," in *Proc. Conf. HotPower*, 2008, p. 10.
- [42] J. Srinivasan *et al.*, "Lifetime reliability: Toward an architectural solution," *Micro*, vol. 25, no. 3, pp. 70–80, May 2005.
- [43] Traces in the Internet Traffic Archive. Retrieved from Sep. 29, 2013. [Online]. Available: <http://ita.ee.lbl.gov/>
- [44] S. Zimmermann *et al.*, "Aquasar: A hot water cooled data center with direct energy reuse," *Energy*, vol. 43, no. 1, pp. 237–245, Jul. 2012.

Authors' photographs and biographies not available at the time of publication.