# CloudHeat: An Efficient Online Market Mechanism for Datacenter Heat Harvesting

SHUTONG CHEN, Key Laboratory of Services Computing Technology and System, Ministry of Education, School of Computer Science and Technology, Huazhong University of Science and Technology, China
ZHI ZHOU, Guangdong Key Laboratory of Big Data Analysis and Processing, School of Data and Computer Science, Sun Yat-sen University, China
FANGMING LIU*, Key Laboratory of Services Computing Technology and System, Ministry of Education, School of Computer Science and Technology, Huazhong University of Science and Technology, China
ZONGPENG LI, Wuhan University, China
SHAOLEI REN, University of California, Riverside, USA

Datacenters are major energy consumers and dissipate an enormous amount of waste heat. Simple outdoor discharging of datacenter heat is energy-consuming and environmental unfriendly. By harvesting datacenter waste heat and selling to the district heating system (DHS), both energy cost compensation and environment protection can be achieved. To realize such benefits in practice, an efficient market mechanism is required to incentivize the participation of datacenters. This work proposes CloudHeat, an online reverse auction mechanism for the DHS to solicit heat bids from datacenters. To minimize long-term social operational cost of the DHS and the datacenters, we apply a RFHC approach for decomposing the long-term problem into a series of one-round auctions, guaranteeing a small loss in competitive ratio. The one-round optimization is still NP-hard, and we employ a randomized auction framework to simultaneously guarantee truthfulness, polynomial running time, and an approximation ratio of 2. The performance of CloudHeat is validated through theoretical analysis and trace-driven simulation studies.

CCS Concepts: • **Information systems → Data centers**; • **Theory of computation → Mixed discrete-continuous optimization**; **Algorithmic mechanism design**; • **Hardware → Power estimation and optimization**;

---

Authors' addresses: Shutong Chen, Key Laboratory of Services Computing Technology and System, Ministry of Education, School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan, Hubei, 430074, China, shutongcs@hust.edu.cn; Zhi Zhou, Guangdong Key Laboratory of Big Data Analysis and Processing, School of Data and Computer Science, Sun Yat-sen University, 135 Xingang Xi Road, Guangzhou, Guangdong, 510006, China, zhouzhi9@mail.sysu.edu.cn; Fangming Liu*, Key Laboratory of Services Computing Technology and System, Ministry of Education, School of Computer Science and Technology, Huazhong University of Science and Technology, 1037 Luoyu Road, Wuhan, Hubei, 430074, China, fmliu@hust.edu.cn; Zongpeng Li, Wuhan University, School of Computer Science, 299 Bayi Road, Wuhan, Hubei, 430072, China, zongpeng@whu.edu.cn; Shaolei Ren, University of California, Riverside, Department of Electrical and Computer Engineering, 900 University Ave, Riverside, 92521, USA, sren@ece.ucr.edu.

---

The corresponding author is Fangming Liu (fmliu@hust.edu.cn). This work was supported in part by the National Key Research & Development (R&D) Plan under grant 2017YFB1001703, in part by NSFC under Grant 61722206 and 61761136014 and 392046569 (NSFC-DFG) and 61520106005, in part by the National 973 Basic Research Program under Grant 2014CB347800, in part by the Fundamental Research Funds for the Central Universities under Grant 2017KFKJXX009, in part by the National Program for Support of Top-notch Young Professionals in National Program for Special Support of Eminent Professionals, in part by NSFC under Grant 61571335 and 61628209, in part by Hubei Science Foundation under grant 2016CFA030 and 2017AAA125, and in part by the U.S. NSF under grants CNS-1551661 and ECCS-1610471.


2376-3639/2017/1-ART1 $15.00
https://doi.org/0000001.0000001

Additional Key Words and Phrases: Datacenter heat harvesting, online market mechanism, approximation algorithms

**ACM Reference Format:**
Shutong Chen, Zhi Zhou, Fangming Liu*, Zongpeng Li, and Shaolei Ren. 2017. CloudHeat: An Efficient Online Market Mechanism for Datacenter Heat Harvesting. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 1, 1, Article 1 (January 2017), 30 pages. https://doi.org/0000001.0000001

## 1 INTRODUCTION

Datacenters are among major energy consumers today around the world. In the U.S. alone, energy consumption by datacenters reached 91 million kWh in 2013 [31]. A massive amount of electricity flows through the IT infrastructure, and is finally transformed into thermal energy, discharged in the form of hot air or water. Such direct dissipation of waste heat is sometimes harmful to the environment (*e.g.*, hot water from a datacenter in Munich, German reportedly interfered the breeding of fish [6]). Furthermore, the heat removal process is also energy-consuming and can consume 12% of the total datacenter energy [4].

The sheer volume of energy consumption and waste heat from datacenters have been subject to negative views and criticisms. A subtle change to the picture recently emerged — the massive waste heat at datacenters can be turned into an opportunity, for residential and commercial heating. Specifically, the huge volume of waste heat produced by datacenters is now viewed as a valuable heat source, making datacenters an integral component of district heating systems (DHSs) widely established in many cold regions of the world.

Datacenter heat harvesting creates a sustainable "win-win" practice: datacenters cut cooling cost and reduce heat dissipation into the environment, while the DHS can (partially) meet its heat demand by harvesting datacenter heat at a modest cost. The inclusion of datacenters into the DHS is beyond a mere conception; it has witnessed pilot projects in various regions of the world, as attested to by the recent numerous installation sites [45].

While the popularity continues to grow, datacenter heat recovery has received very little attention from the research community and has not been well understood. In fact, datacenters currently contribute to DHS by passively and myopically producing heat, *without* taking into account the actual heat demand at runtime or the other datacenters' heat production. While simple market mechanisms exist today, they fail to balance datacenter heat supply and heat demand at runtime. For example, the open district heating project initiated by Fortum purchases heat from datacenters at market prices [29]. Hence, to balance supply and demand, the DHS manager must accurately predict datacenters' heat supply before setting price, but datacenters may not be controlled by the DHS manager and they have highly dynamic energy consumption incurred by the energy proportionality and varying workload. Thus, predictions of heat supply from datacenters are inaccurate and unreliable. Contract based mechanisms reserve heat supply from datacenters, and may work for long-term capacity planning. Yet they cannot accurately reflect the real-time DHS operating condition. For these reasons, the current practice of uncoordinated datacenter heat recovery is highly inefficient and results in frequent mismatches between heat supply and demand (or the desired level of supply). This mismatch leads to insufficient datacenter heat supplies and forces the DHS to resort to other more costly heat sources.

To meet the heat demand at a low cost through coordinated datacenter heat harvesting, we propose a novel market design, CloudHeat, based on reverse auction mechanisms. With CloudHeat, each datacenter first voluntarily submits its bid, *i.e.*, the amount of available heat and the operational cost of recovering such amount of heat to the DHS. If it wins this bid, the DHS pays it no less than the declared cost. Nonetheless, the market design in CloudHeat is rather challenging: datacenters are naturally self-interested and may not truthfully reveal their costs to the DHS. Furthermore,

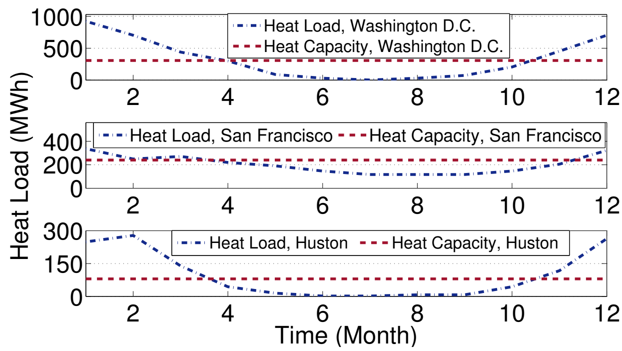


Fig. 1. The estimated heat demand of 50,000 residential homes (each with 1700 square-feet) at three selected locations in the U.S., and the peak-hour heat capacity of colocation datacenters with typical PUE = 1.5 [36, 55].

CloudHeat must work in an online fashion, given the time-varying operational environments and the heterogeneous heat sources in the heat harvesting market. A general auction mechanism such as Vickrey-Clarke-Groves (VCG) [9, 28, 57] becomes computationally prohibitive. To address these challenges, we apply the RFHC [62] sub-framework to decompose the long-term optimization into a series of one-round auction problems. In each round, the randomized auction sub-framework and primal-dual algorithm are leveraged to ensure CloudHeat's computational efficiency and truthfulness in expectation, with a small approximation ratio of 2 in social cost. The overall competitive ratio of the resulting online auction framework is bounded by $(2 + \frac{2\beta}{(w+1)\alpha})$, *i.e.*, the social operational cost obtained by CloudHeat is at most $(2 + \frac{2\beta}{(w+1)\alpha})$ times the optimal social operational cost in offline setting, where $w$ is the length of look-ahead window and $\alpha$, $\beta$ are system-dependent parameters.

In the rest of the paper, we first present the background for datacenter heat harvesting in Sec. 2, and formulate the market design in Sec. 3. Sec. 4 and Sec. 5 present the online algorithm sub-framework and the single round auction algorithm, respectively. Simulations are presented in Sec. 6, related works are discussed in Sec. 7, and Sec. 8 concludes the paper.

## 2 BACKGROUND

In this section, we provide some background information to facilitate the readers' understanding of datacenter heat harvesting. It will highlight the potential market size for datacenter heat harvesting, remarkable progress in the engineering domain as well as the necessity of designing an efficient market for datacenters' participation in heat recovery.

**Is it available to warm buildings by datacenter heat?** Firstly, datacenter heat is highly available and reliable: even during natural disasters, datacenters can still operate seamlessly due to their on-site backup power and strong physical protection, thus continuously producing heat. Further, the amount of datacenter heat is abundant and increasing. More concretely, Fig. 1 plots the comparison between the average hourly heat load of 50,000 homes and the peak-hour heat capacity[1] of the commercial colocation datacenters registered in [11], at three U.S. cities. It clearly demonstrates that, even without considering micro datacenters distributed ubiquitously, the heating capacity of only large commercial datacenters can satisfy a large portion of the total heat demand.

[1] In this work, we consider the power consumed by IT equipment is totally converted to heat and can be removed [49].

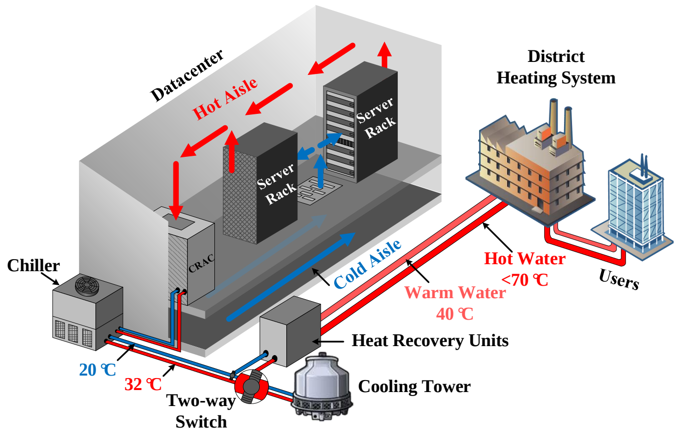


Fig. 2. A brief illustration of heat recovery in an air-cooled typical datacenter.

**Is it feasible to harvest heat from the datacenter?** Datacenter heat harvesting for district heating involves two key aspects, the first is to recover heat that satisfies the quality requirement of district heating, and the second is the transition of recovered heat from datacenters to the DHS. Note that the power consumed by IT infrastructures is the most important part of datacenter's power consumption [56], and almost all the IT power consumption is converted to heat. Although some running non-IT infrastructures can also produce heat, datacenter heat harvesting is utilized to recover waste heat produced by IT equipment (e.g., servers and switches), rather than non-IT infrastructures. In our work, we only focus on those waste heat produced by IT equipment. Furthermore, as the improvement in cooling efficiency focuses on reducing cooling power which belongs to the energy used by non-IT infrastructures, the amount of waste heat produced by IT equipment will not be affected.

For the first aspect, Fig. 2 illustrates heat recovery cycling in a typical air-cooled datacenter. The cooled air stems from the computer room air conditioning, first goes through the cold aisle, rises from the perforated floor, and flows through the racks where it picks up heat. The warm air then exits the rear of the racks and returns to the computer room air conditioning. The heat harvested by the computer room air conditioning is exchanged to the chiller outside the server room through the medium of water or air, and finally taken away by chilled water. However, the warm chilled water can not be directly fed to the DHSs, since its temperature is generally lower than 35℃ and inappropriate for district heating [18]. To enable heat harvesting for district heating, additional heat recovery units, such as industrial heat pumps are required to boost the return water (40℃) of the heating network to about 70℃ [3]. Currently, heat pumps specified for datacenter heat recovery have emerged in the market, for example, 3 heat pumps with a total heating capacity of 1.6 MW was connected in series to recovery the heat from the Bahnhof Thule datacenter in central Stockholm [23]. Furthermore, when district heating is not required (e.g., at hot days), then the heat recovery units can be shut down, and the hot water out of the chiller can be circulated to the cooling tower which is typically deployed in datacenters.

We should note that for the emerging liquid-cooled datacenter, the above paradigm for heat recovery is still applicable, and the only difference from air cooling is that the sever heat is absorbed by liquid but not air. Specifically, for liquid-cooled system with cold plate, water passes through a sealed liquid path to the cold plate which touches server components (mainly CPU, GPU, and memory) closely, and then water picks up heat and exits the server [40]. For immersion liquid-cooled

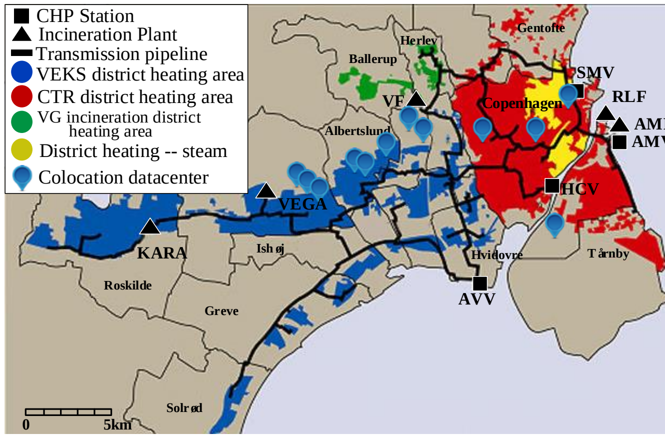


Fig. 3. The coverage of 4 district heating areas and the distribution of 12 colocation datacenters in the Greater Copenhagen. The transmission pipelines transmit $9.72 \times 10^6$ MWh of annual heat generated by CHP (combined heat and power) stations and incineration plants to serve 1.2 million citizens [11, 46].

system, the entire server is immersed in tailor-made liquid for cooling [51]. Like air-cooled system, the heat harvested by liquid is also exchanged to a chiller. Compared with air-cooled datacenter, these advance cooling techniques can increase the potential benefits of datacenter heat harvesting. For example, warm water is allowed to cool IT equipment in some advanced cooling systems [40]. When raising the temperature of feed water (which is used to pick up heat), the temperature of return water (which has picked up heat) will increase as well as the recovery efficiency and potential benefits, since less power will be used to boost the temperature of return water. In Sec. 6, we will discuss the influence of recovery efficiency on heat harvesting detailedly.

For feasibility of transmitting recovered heat from datacenters to the DHS, unlike many owner-operated datacenters (e.g., Google) in rural areas, the colocation datacenter[2] is a larger sector (over 40% [31]) and mostly locates in the densely-populated metropolitan area. In the U.S. alone, there are already nearly 2,000 large colocation datacenters, accounting for as five times the energy consumption as Google-type datacenters [53]. And in Europe, the number of colocation datacenter is about 1,500 [11]. These colocation datacenters in metropolis can be readily connected to the DHS. To further demonstrate the feasibility, we plot the coverage and pipelines of the 4 DHSs that serve the Greater Copenhagen, together with the colocation datacenters within this area, as shown in Fig. 3. A large proportion (8 out of 12) of the colocation datacenters are located within the coverage of DHSs [11, 46]. Moreover, all the colocation datacenters are located very close to the pipelines that transmit the heat, with a distance of at most 2.2 km. These observations indicate that in the Greater Copenhagen, datacenter waste heat can directly serve the nearby area or be readily transmitted into the DHS pipelines, without requiring much CapEx which is dominated by the distance and we will further discuss next.

**How big is the market?** Though datacenter heat recovery for district heating has been proved to be feasible, we acknowledge that it has only emerged recently in certain areas like Stockholm [45]. Nonetheless, there is potentially a large market of datacenter heat harvesting. To see this point, we study the weather conditions in metropolitans that have the largest amount of colocation datacenters, in America, Europe, and Asia, as shown in Table 1. From the weather data of Typical

[2]The colocation datacenter is a type of datacenter whose physical space can be rent by tenants to house their servers in a shared building.

Table 1. The weather conditions of metropolitans with large amount of colocation datacenters [11, 17, 20, 21, 24, 54].

| Locations | ⩽ 19℃ | ⩽ 12℃ | # of the colocation datacenter | Deployed DHSs |
|---|---|---|---|---|
| New York, US | 69.44% | 50.35% | 48 | ✓ |
| Seattle, US | 90.97% | 53.82% | 41 | ✓ |
| San Francisco, US | 91.32% | 32.64% | 18 | ✓ |
| Toronto, CA | 85.76% | 61.81% | 39 | ✓ |
| London, UK | 93.75% | 64.58% | 71 | ✓ |
| Frankfurt, DE | 89.24% | 57.64% | 51 | ✓ |
| Copenhagen, DK | 96.18% | 69.10% | 14 | ✓ |
| Stockholm, SE | 97.92% | 70.14% | 28 | ✓ |
| Tokyo, JP | 69.10% | 46.18% | 24 | ✓ |
| Hong Kong, CN | 29.17% | 0.00% | 47 | × |
| Singapore, SG | 0.00% | 0.00% | 22 | × |
| Los Angeles, US | 75.00% | 8.33% | 62 | × |
| Houston, US | 42.01% | 5.36% | 35 | × |

The air temperature conditions, 19℃ and 12℃, are in line with the typical base temperature of heating degree day — a building needs heating when the outside temperature below it [30], and the general temperature condition where heat is necessary [44], respectively.

Meteorological Year [20], we can see that, for many metropolitans such as New York, London, and Tokyo in America, Europe, and Asia, both long-period cold weather conditions and large numbers of datacenters (which are mostly larger than 20) are satisfied, meanwhile DHS is also highly popularized in those metropolitans. Thus, datacenter waste heat can be naturally suitable for them. For some cities such as LA and Houston in America, though there are a larger number of datacenters, DHSs have not been deployed. And even if there are deployed DHSs, it is still unpractical to recover datacenter waste heat in these cities due to the short period of cold days. While for Hongkong and Singapore that are at low latitudes and have long-period of warm days, heating is not required at all. In summary, though datacenter heat for district heating is not commonly applicable, it is promising for many metropolitans in Europe, some in America and few in Asia.

**Is the market profitable?** A report from the open district heating project shows that in Stockholm, the revenue from heat harvesting earned by Pionen, a colocation datacenter belonging to Bahnhof, surpasses £30,000 in 4 months [22]. However, a natural concern for the DHS is the profitability of datacenter heat harvesting, which incurs one-time infrastructure cost as well as recurring operating expense. Current estimates suggest that the infrastructure cost of building heating plant and heat recovery unit are both proportional to installed capacity and the cost of pipeline grows on average with its length [12, 23]. Fig. 4(a) shows the CapEx of installing gas boiler, biomass boiler, and heat pump all of which are with peak capacity of 15 MW, as well as the CapEx of installing heat recovery units with the same capacity and 3 km pipeline. We find that compared with installing new heating plant, datacenter heat harvesting is more economically efficient: the CapEx can be reduced by up to 32%, 54%, and 76% compared with installing the above three heating plants, respectively. This means the DHS can reduce its own heating plant construction and easily pass down its cost saving to datacenter operators to cover their CapEx.

As shown in Fig. 4(b), datacenter heat harvesting also shows its advantages in reducing heat supply cost. Without energy loss, the average cost of natural gas and heat recovery are 2.2404 cents/kWh and 1.0012 cents/kWh, respectively. And the CapEx of installing the heat recovery unit

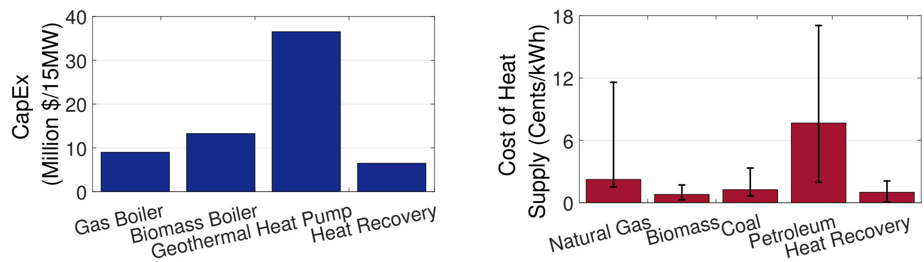


(a) CapEx of installing heating plant [12, 23].

(b) Cost of heat supply in the USA [19].

Fig. 4. Cost of heat harvesting.

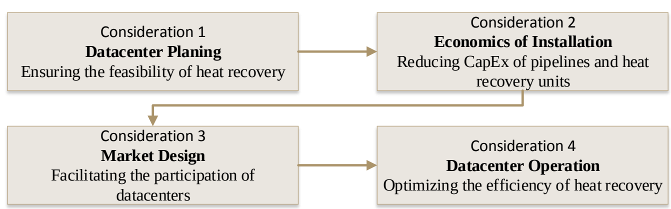


Fig. 5. Four considerations for the lifecycle of datacenter heat harvesting.

is about 6.1045 million dollars, as shown in Fig. 4(a). Hence, the return on investment of datacenter heat harvesting can be estimated by: $\frac{\$6.1045\times 10^6}{\frac{2.2404\text{ cents/kWh}-1.0012\text{ cents/kWh}}{10\text{ MW}}} \approx 5.6$ years. That is to say, the return on investment of datacenter heat harvesting can be less than 5.6 years if the DHS persistently purchases more than 10 MWh of heat from datacenter rather than using natural gas. Further, as discussed in Sec. 1, market efficiency and participation of datacenters should be guaranteed by a carefully designed market mechanism, which is just the focus of our work.

**The gap between the current system and the envisioned system.** Datacenter heat harvesting is a complex process that involves both engineering designs as well as market interactions with DHS, requiring the following considerations for wider adoption, as displayed in Fig. 5.

Consideration 1: *Suitable planning of datacenter.* Heat harvesting should be kept in mind at the very beginning of the lifecycle of datacenters. Specifically, when planning the location for a datacenter, easy integration into the DHS pipeline network should be considered as a new dimension besides the network access, climate, energy price, etc [5]. Furthermore, when designing the cooling hierarchy, suitable cooling technology (free-air cooling [27], or liquid cooling [34]) should be identified to efficiently capture and transmit the heat. Consideration 2: *Economical installation of auxiliary devices.* To make datacenter heat harvesting more appealing and profitable to both datacenters and DHSs, the CapEx of auxiliary infrastructures that mainly includes heat pumps and pipelines is expected to be lowered. Such cost reduction can be realized via technical advance in R&D [3] or drop in the manufacturing cost. Consideration 3: *Efficient market design.* Given the engineering feasibility of datacenter heat harvesting, an efficient market design is crucial to incentivize datacenters to participate in heat harvesting program while benefiting both datacenters themselves as well as the DHS. Consideration 4: *Optimization of datacenter operation.* The efficiency of datacenter heat harvesting can be further improved by optimizing datacenter operation. For example, by applying datacenter power management techniques such as DVFS, temperature adaption, sever right-sizing, etc to better balance the heat demand and heat supply [59]. Besides, cooling-aware workload placement can also be optimized to improve the efficiency heat capture and transmission [38].

Note that among the four considerations, datacenter planning, datacenter operation, and auxiliary device manufacturing can be considered as engineering problems that each individual datacenter addresses by leveraging the numerous recent advances [3, 59]. In contrast, the critical aspect — market design, which incentivizes and coordinates datacenters' participation in heat harvesting — has remained under-explored and become a crux for the sustainable growth of datacenter heat harvesting, and is the focus of our study.

## 3 ONLINE AUCTION MODEL FOR HEAT HARVESTING

In this section, we first present CloudHeat, an online reverse auction approach to datacenter waste heat harvesting, and then formulate the social operational cost minimization problem underlying CloudHeat. We list key notations in this paper in Table 2.

### 3.1 System Overview

In our work, we assume that there is one DHS trading with a number $M$ of datacenters equipped with heat recovery units to extract waste heat, in line with the fact that a DHS usually covers a moderate-size district, *e.g.*, Enwave at Toronto, CA and EnviroEnergy at Nottingham, UK [17]. For the special case that multiple DHSs trade to multiple datacenters, it can be accommodated by the more complicated double auction framework. For simplicity, we consider the case that one DHS trades to multiple datacenters.

When a certain level of heat demand arrives, the local DHS would solicit recovered heat from datacenters through the reverse auction. However, in cold days, the extracted heat from datacenters may not suffice to cover users' heat demand; heating plants operated by the DHS can be turned on to fulfill the remainder heat requirement. Heating plants often utilize fossil fuel such as coal and natural gas. These turn out to be the more expensive than recovered and recycled heat; the detail is discussed in Sec. 2.

### 3.2 The Online Reverse Auction

The proposed solution, CloudHeat, is based on a reverse auction that runs in a time-slotted fashion across a time frame of $T$ time slots. The length of one time slot can vary from one to several hours: for wholesale electricity market, the length of one time slot is one hour [60]. As illustrated in Fig. 6, at the beginning of each time slot $t \in \mathcal{T} = \{1, 2, \ldots, T\}$, the DHS who acts as the auctioneer receives the total heat demand $W(t)$ and solicits bids from datacenters. Then datacenter $i \in \mathcal{I} = \{1, 2, \ldots, M\}$ voluntarily submits a heat recovery bid $(c_i(t), u_i(t))$, where $u_i(t)$ is the amount of recovered heat datacenter $i$ can supply, and $c_i(t)$ is the incurred cost whose range is referred to Fig. 4(b). Consequently, the DHS determines the winning bids (*i.e.*, bids which are chosen to supply heat at the promised level) and the amount of self-produced heat of the current time slot $v(t)$, and announces the winning bids and payments $r(t)$. Our objective is to maximize the global efficiency of heat recycling by minimizing the social operational cost that consists of operational cost of datacenters and DHS, which is equivalent to maximizing the social welfare of the CloudHeat ecosystem (with payments canceling themselves). Note that profit maximization is also a realistic objective while is different to social welfare maximization in auction framework design, interested readers can refer to [43] for the auction framework of profit maximization.

**Operational cost of datacenters:** We use a 0-1 variable $x_i(t)$ to indicate whether datacenter $i$ wins its bid ($x_i(t) = 1$) or not ($x_i(t) = 0$). Let $r_i(t)$ denote the payment to datacenter $i$. Then, at time slot $t$, the operational cost of datacenter $i$ incurred by the waste heat auction is $c_i(t)x_i(t) - r_i(t)$.

Table 2. Summary of Notations

| Notations | Description |
|---|---|
| $K$ | # of homogeneous heating plants of DHS |
| $V$ | The capacity of each heating plant |
| $\mathcal{T}$ | The set of time slots |
| $\mathcal{I}$ | The set of volunteer colocation datacenters in the reverse auction |
| $W(t)$ | The total heat demand at time slot $t$ |
| $u_i(t)$ | The amount of recovered heat datacenter $i$ can supply at time slot $t$ |
| $c_i(t)$ | The datacenter $i$'s cost of recovering heat at time slot $t$ |
| $v(t)$ | The amount of self-produced heat of DHS at time slot $t$ |
| $r_i(t)$ | The payment for datacenter $i$ at time slot $t$ |
| $x_i(t)$ | Binary variable indicating whether datacenter $i$ wins or not at time slot (t) |
| $y(t)$ | # of running heating plants at time slot $t$ |
| $p_h(t)$ | The fuel cost for ont unit heat of heating plants at time slot $t$ |
| $\alpha$ | The sunk cost of maintaining a heating plant in its active state per time slot |
| $\beta$ | The start-up cost of turning on a heating plant |

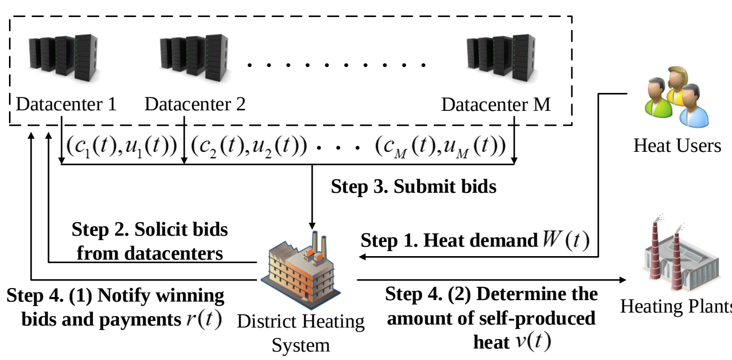


Fig. 6. A system overview of CloudHeat.

Note that as discussed in Sec. 2, the datacenter can circulate the hot water out of the chiller to the cooling tower instead of heat recovery units if its bid is not accepted, and thus no operational cost of heat recovery will be incurred. Later in Sec. 5, we will show that our carefully-designed payment mechanism ensures the individual rationality, *i.e.*, datacenter $i$ receives a payment which is no lower than the cost $c_i(t)$ if it wins the bid.

**Operational cost of DHS:** The DHS deploys a number of $K$ homogeneous heating plants, and its total operational cost at time $t$ consists of four components: 1) the payment to datacenters: $\sum_i r_i(t)$. 2) The switching cost of heating plants: $s(y(t-1), y(t)) = \beta[y(t) - y(t-1)]^+$, where $[\cdot]^+ = \max\{0, \cdot\}$, $y(t)$ is the number of running heating plants at time slot $t$, and $\beta$ denotes the start-up cost of turning on a heating plant. Start-up cost typically involves the heating up cost, the time-amortized capital and additional maintenance costs resulted from each start-up. 3) The maintenance cost of running a number of $y(t)$ heating plants: $\alpha y(t)$, here $\alpha$ denotes the sunk cost of maintaining a heating plant in its active state per time slot. 4) The fuel cost of self-produced heat: $v(t)p_h(t)$ where $p_h(t)$ represents the fuel cost for one unit heat of the heating plants. Then, the total operational cost of DHS is $v(t)p_h(t) + \sum_i r_i(t) + \alpha y(t) + s(y(t-1), y(t))$.

Given the operational cost of each datacenter and the DHS, we are now in a position to formally formulate the social operational cost minimization problem (CMP) that determines the winning bids of the CloudHeat auction and the number of running heating plants.

$$\mathtt{CMP}: \min \quad \sum_{t\in\mathcal{T}} \Big\{ \sum_{i\in\mathcal{I}} c_i(t)x_i(t) + v(t)p_h(t) + \alpha y(t) + s(y(t-1), y(t)) \Big\} \tag{1}$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{I}} u_i(t)x_i(t) + v(t) \geqslant W(t), \forall t \in \mathcal{T} \tag{2}$$

$$v(t) \leqslant V \cdot y(t), \ \forall t \in \mathcal{T} \tag{3}$$

$$x_i(t) \in \{0,1\}, \ \forall i \in \mathcal{I}, \ t \in \mathcal{T} \tag{4}$$

$$y(t) \in \{0,1,...,K\}, v(t) \geqslant 0, \ \forall t \in \mathcal{T} \tag{5}$$

Here $V$ is the capacity of each heating plant. The constraint (2) indicates that the total amount of heat demand should be satisfied by the extracted datacenter waste heat and self-produced heat. The constraint (3) guarantees the self-produced heat do not exceed the total heat capacity. Note that as the variable $v(t)$ is continuous and bounded by the constraint (3), it is not suitable to represent $v(t)$ as a combination of $x_i(t)$ and reduce the CMP as an integer linear programming problem.

For some traditional heat plants such as large centralized coal- or oil-fired heating plants, heat production is constrained by the start-up time and ramping-up/down rate [39]. Fortunately, for the natural gas fired heating technology have become widely adopted in modern district heating, the start-up time can be reduced to as short as 1 minute, and the heating plants can ramp from start command to peak within $3-6$ minutes [42], which is very small when compared to the length of one time slot. Thus, for simplicity, we do not consider the above operation constraints for heating plants. Interested readers can refer to [39] for the modeling of these constraints in the scenario of micro-grid.

**Why is reverse auction suitable for our problem?** Compared to other market approaches such as (i) reference-based pricing [29] that may under-price or over-price datacenter heat, (ii) the pre-contract approach that fails to capture the time-varying operation environment, and (iii) supply function bidding that is vulnerable to untruthful information [32], CloudHeat which builds upon reverse auction has the following advantages. It enables market efficiency and agility through pricing datacenter heat based directly on realtime supply-demand. By matching the heat demand with datacenters with the lowest cost of heat recovery, it reduces the chance of over-pricing and under-pricing, and hence reduces system-wide social cost and increases waste heat utilization efficiency.

However, the market mechanism design for CloudHeat is indeed challenging. The first challenge arises from the *online* nature of CMP, as we can see that the switching cost temporally couples the social operational cost at different time slot $t$ across the time frame $\mathcal{T}$. Thus, given limited knowledge of further information on heat demand and datacenter bids, how can we make online decisions to achieve *economic efficiency*, *i.e.*, minimize the long-term social operational cost? Furthermore, note that when $y(t) = 0$ for all $t \in \mathcal{T}$, CMP is reduced to a minimum knapsack problem and is NP-hard [8]. Thus CMP is also an NP-hard problem in general, as it subsumes the case of $y(t) = 0$. And as there can be tens of participating datacenters, as discussed in Sec. 2, the complexity of solving CMP is high. This rules out direct application of the classic VCG auction that requires exactly solving the underlying social operational cost minimization problem multiple times and hence is computationally prohibitive in practice. Then how can we design a *computationally* *efficient* approximation algorithm to preserve the social efficiency, yet still achieving *truthfulness (in expectation)* and *individual rationality*?

In the following two sections, we propose an online framework that focuses on two problems: at time slot $t$, 1) how many heating plants should be opened; and 2) which winner should be chosen. The framework is summarized in Algorithm 3, where the former problem is solved by the randomized fixed horizon control algorithm (line 8 to line 13 in Algorithm 3) with an approximation algorithm (Algorithm 2, where Algorithm 1 is used as a component) in Sec. 4, and the later is determined by the random auction mechanism (Algorithm 4) in Sec. 5.

## 4 THE ONLINE ALGORITHM FRAMEWORK

In this section, to address the challenges of solving CMP with limited future information, we design an online algorithm framework to decompose the long-term auction problem into a series of one-round auctions.

### 4.1 Offline Algorithm for CMP

We first study CMP in the offline setting, where the heat demand $W(t)$, production cost of heating plants $p_h(t)$, as well as the bids of each datacenter $(c_i(t), u_i(t))$, $\forall i \in \mathcal{I}$, over the time frame $\mathcal{T}$ are given at the beginning of time slot $t = 1$.

Note that if the number of running heating plants $y(t)$ over time frame $\mathcal{T}$ is given, the solution of CMP can be obtained in each time slot independently. Moreover, the switching cost is the only term that jointly depends on the past state (*i.e.*, $y(t-1)$) and the current state (*i.e.*, $y(t)$). Based on this observation, we reformulate CMP to the following equivalent form:

$$\begin{aligned} \min \quad & \sum_{t \in \mathcal{T}} C_{y(t)}(\boldsymbol{x}(t), v(t)) + \sum_{t \in \mathcal{T}} \Big( \alpha y(t) + s(y(t-1), y(t)) \Big) \\ \text{s.t.} \quad & y(t) \in \{0, 1, \dots, K\}, \ \forall t \in \mathcal{T} \end{aligned}$$

where at any time slot $t \in \mathcal{T}$,

$$\begin{aligned} C_{y(t)}(\boldsymbol{x}(t), v(t)) \triangleq \min \quad & \sum_{i \in \mathcal{I}} c_i(t) x_i(t) + v(t) p_h(t) \\ \text{s.t.} \quad & \sum_{i \in \mathcal{I}} u_i(t) x_i(t) + v(t) \geqslant W(t), \\ & 0 \leqslant v(t) \leqslant V \cdot y(t), \\ & x_i(t) \in \{0, 1\}, \ \forall i \in \mathcal{I} \end{aligned} \tag{6}$$

The reformulated problem can be viewed as a shortest path problem in a directed graph $G$, from $y(0) = 0$ to $y(T) = \{0, 1, \dots, K\}$. Each vertex $y(t) = \{0, 1, \dots, K\}, t \in \mathcal{T}$ denotes the possible number of running heating plants the DHS may choose at time $t$. An edge $(y(t-1), y(t))$ represents the process that DHS turns on $y(t) - y(t-1)$ (when $y(t) - y(t-1) \geqslant 0$) or turns off $y(t-1) - y(t)$ heating plants (when $y(t-1) - y(t) > 0$) from time slot $t-1$ to $t$. And the weight of edge $(y(t-1), y(t))$, denoted by $d(y(t-1), y(t))$, can be regarded as the social operational cost at time slot $t$ and calculated by:

$$d(y(t-1), y(t)) \quad = \quad C_{y(t)}(\boldsymbol{x}(t), v(t)) + \alpha y(t) + s(y(t-1), y(t)). \tag{7}$$

Note that given the weight of each edge, the shortest path, *i.e.*, the minimum social operational cost over the time frame $\mathcal{T}$, as well as the optimal number of running heating plants of each time slot can be computed with Dijkstra's algorithm [39]. Unfortunately, the weight of each edge can not be easily obtained, as the term $C_{y(t)}(\boldsymbol{x}(t), v(t))$ in $d(y(t-1), y(t))$ corresponds to the mixed integer linear programming problem (6) that is NP-hard. Specifically, as mentioned in Sec. 3, at any time slot $t \in \mathcal{T}$, when $y(t) = 0$, i.e., there is no running heating plants, the problem (6) is reduced to a minimum 0-1 knapsack problem, a well-known NP-hard problem. The problem (6) is also NP-hard in general, as it subsumes the case of $y(t) = 0$. Moreover, in the case of $y(t) = 0$, the problem (6) only has 0-1 variable $x_i(t)$ which indicates whether DHS purchases heat from datacenter $i$. Hence the problem (6) can be solved by the dynamic programming technique. However, in the case of $y(t) \neq 0$, i.e., when there are running heating plants, the dynamic programming technique for 0-1 knapsack problem is not applicable to the problem (6). The basic idea of this technique is to label each item and divide the 0-1 knapsack problem into several subproblems based on labeled items. But there is a continuous variable $v(t)$ in the problem (6) which cannot be divided and labeled. As discussed in Sec. 3, when $y(t) \neq 0$, continuous variable $v(t)$ is utilized to represent the amount of heat produced by heating plants, and hence there are continuous variable $v(t)$ and 0-1 variable $x_i(t)$ in the problem (6). The continuous variable $v(t)$ makes the dynamic programming inapplicable.

Given the complexity challenge, we resort to efficient approximation algorithms that compute near-optimal solutions. We apply the linear programming relaxation strengthening technique [7] and propose a primal-dual based approximation algorithm to compute an approximate solution to the problem (6).

**Are prior primal-dual algorithms directly applicable to the problem (6)?** The prior primal-dual algorithms [60, 63] have considered that the buyer has self-produced capacity to cover all demands. They are, however, not applicable to the problem (6), since the capacity of heating plants is limited in our work and directly applying those solutions may result in violation of the constraint (2). To overcome this limitation, in this work, we introduce a dual variable for the constraint (3) to guarantee the feasibility and near-optimality of the solution. For simplicity of presentation, in the following discussion on the primal-dual algorithm, we omit the time index $t$ due to its independence.

Consider a subset $\mathcal{X} = \{i_1, i_2, \ldots\} \subseteq \mathcal{I}$ that satisfies $\sum_{i \in \mathcal{X}} u_i < W$. $\mathcal{X}$ can be regarded as a collection of chosen bids whose total heat supply doesn't meet the heat demand $W$ yet. We have to continue to choose bids from $\mathcal{I} \setminus \mathcal{X}$ to cover the deficit $\Delta\mathcal{X} = W - \sum_{i \in \mathcal{X}} u_i$. For datacenter $i \in \mathcal{I} \setminus \mathcal{X}$, let $u_i(\mathcal{X}) = \min\{u_i, \Delta\mathcal{X}\}$ denote how much contribution it provides in covering the remainder demand $\Delta\mathcal{X}$. After introducing $\Delta\mathcal{X}$ and $u(\mathcal{X})$, we obtain the linear program relaxation (LPR) of the problem (6) by relaxing the binary variables $x_i$:

$$
\begin{aligned}
\mathsf{LPR}: \ \min \quad & \sum_{i \in \mathcal{I}} c_i x_i + p_h v \\
\text{s.t.} \quad & \sum_{i \in \mathcal{I} \setminus \mathcal{X}} u_i(\mathcal{X}) x_i + v \geqslant \Delta\mathcal{X}, \ \forall \mathcal{X} \subseteq \mathcal{I} : \Delta\mathcal{X} > 0 \\
& 0 \leqslant v \leqslant V \cdot y, \ 0 \leqslant x_i, \ \forall i \in \mathcal{I}
\end{aligned}
$$

The first constraint states that for any set in $\{\mathcal{X} | \mathcal{X} \subseteq \mathcal{I}, \Delta\mathcal{X} > 0\}$, the difference between the heat demand $W$ and the amount of the chosen heat supply $\sum_{i \in \mathcal{X}} u_i$ has to be covered by the remainder bids. Further note that every feasible solution of the problem (6) is also feasible to LPR. Now we introduce the dual variables $z(\mathcal{X})$ and $\eta$ corresponding to the constraints $\sum_{i \in \mathcal{I} \setminus \mathcal{X}} u_i(\mathcal{X}) x_i + v \geqslant \Delta\mathcal{X}$ and $v \leq V \cdot y$, respectively. Then we obtain the dual problem of LPR [52]:

$$
\begin{aligned}
\mathsf{LPRD}: \ \max \quad & \sum_{\mathcal{X} \subseteq \mathcal{I} : \Delta\mathcal{X} > 0} \Delta\mathcal{X} z(\mathcal{X}) - V y \cdot \eta \\
\text{s.t.} \quad & \sum_{\mathcal{X} \subseteq \mathcal{I} : \Delta\mathcal{X} > 0, i \notin \mathcal{X}} u_i(\mathcal{X}) z(\mathcal{X}) \leqslant c_i, \ \forall i \in \mathcal{I} && (8) \\
& \sum_{\mathcal{X} \subseteq \mathcal{I} : \Delta\mathcal{X} > 0} z(\mathcal{X}) \leqslant p_h + \eta, && (9)
\end{aligned}
$$

$$z(\mathcal{X}), \eta \geqslant 0, \; \forall \mathcal{X} \subseteq \mathcal{I} : \Delta\mathcal{X} > 0 \tag{10}$$

Following the idea of primal-dual optimization, we can construct a mixed integer solution to LPR and a feasible solution to its dual LPRD, via continuously increasing the dual variable $z(\mathcal{X})$. For each feasible $z(\mathcal{X})$, once a dual constraint (*i.e.*, (8) or (9)) becomes tight (a constraint $ax \leqslant b$ is tight if $ax = b$), the corresponding primal variable (*i.e.*, corresponding to $x_i$ or $v$) is determined. The increase of $z(\mathcal{X})$ can not be terminated until the constraint (2), *i.e.*, total heat demand $W$, is satisfied. Based on this idea, we design a feasible, 2-approximation solution shown in Algorithm 1 to the problem (6).

We first discuss Algorithm 1's complexity. As the termination condition of the `while` loop is $\Delta S > 0$, the loop runs at most $M + 1$ times when $\sum_{\mathcal{X} \subseteq \mathcal{I}: i \notin \mathcal{X}, \Delta\mathcal{X} > 0} u_i(\mathcal{X}) z(\mathcal{X}) = c_i$ is satisfied for all $i \in \mathcal{I}$ as well as $\sum_{\mathcal{X} \subseteq \mathcal{I}} z(\mathcal{X}) = p_h$, $v = 0$. And in the `while` loop, `if` statement runs at most $M$ times. We conclude that the total time complexity is $O(M^2)$.

The feasibility and approximation ratio are established in Lemma 1 and Theorem 1, respectively.

**ALGORITHM 1:** 2-approximation Primal-Dual Algorithm

Initialization: input $c_i, u_i, p_h, V, y$, and $W$, $\forall i \in \mathcal{I}$. Let $x_i = 0$, $\forall i \in \mathcal{I}$, $z(\mathcal{X}) = 0$, $\forall \mathcal{X} \subseteq \mathcal{I}$, $v, \eta = 0$, and $S = \emptyset$ be the set of chosen bids.
If $\sum_{i=1}^{M} u_i + Vy < W$, there exists no feasible solution for (6). Set the corresponding objective value to infinity.
**While** $\Delta S > 0$ **do**
  Increase dual variable $z(\mathcal{X})$ continuously, if $v > 0$ then increase $\eta$ at the same rate simultaneously, until the dual constraint (8) or (9) become tight;
  **if** $\sum_{\mathcal{X} \subseteq \mathcal{I}: i \notin \mathcal{X}, \Delta\mathcal{X} > 0} u_i(\mathcal{X}) z(\mathcal{X}) = c_i$ **then**
  $x_i = 1; S = S \bigcup \{i\}$;
  **end if**
  **if** $\sum_{\mathcal{X} \subseteq \mathcal{I}} z(\mathcal{X}) = p_h$ and $v = 0$ **then**
  $v = \min\{\Delta S, Vy\}$; $\Delta S = \Delta S - v$;
  **end if**
**end while**

Lemma 1. *The solution given by Algorithm 1 is feasible to the problem (6), the linear relaxation* LPR *and the dual problem* LPRD.

Proof. *i*) Note that for any $i \in \mathcal{I}$, the variables $x_i$ is initialized to be 0 and set to 1 when the corresponding dual constraint is tight, hence the constraint $x_i \in \{0, 1\}$ is satisfied. The variable $v$ is initialized to 0 and may be updated to no more than the heating capacity $Vy$, hence the constraint $v \leq V \cdot y$ is also satisfied. Furthermore, since the iteration stops only in the case of $\Delta S \leqslant 0$, i.e., $\sum_{i \in S} u_i(\mathcal{X}) x_i + v \geqslant W$ when the iteration stops, the constraint $\sum_{i \in \mathcal{I}} u_i(t) x_i(t) + v(t) \geqslant W(t)$ is satisfied. In summary, the solution given by Algorithm 1 is feasible to (6).

*ii*) We can verify that the feasible solution to the problem (6) is feasible to LPR as well.

*iii*) For the dual LPRD: Since the dual variable $\eta$ increases as the same rate of $z(\mathcal{X})$ after $v > 0$, we observe that the constraint (9) is always tight and never violated after we update $v > 0$. Moreover, the increasing $z(\mathcal{X})$ doesn't increase the left-hand side of (8) corresponding to $i \in \mathcal{I}$, hence the constraint (8) is always satisfied. □

Theorem 1. *Algorithm 1 is a 2-approximation algorithm to the problem (6),* i.e., *the cost computed by Algorithm 1 is at most 2 times the optimal cost in the problem (6).*

We prove the theorem in Appendix A.1.

Having solved the NP-hard problem, we are now ready to compute the aforementioned edge weight $d(y(t-1), y(t))$ and present the approximation algorithm to `CMP` in the offline setting, as shown in Algorithm 2.

**ALGORITHM 2:** Approximate Dynamic Algorithm for `CMP` in the offline scenario

1: Initialization: input $c_i(t), u_i(t), p_h(t), V(t), y(t)$, and $W(t)$, $\forall i \in \mathcal{I}, \forall t \in \mathcal{T}$. Let $y(0) = 0$.
2: **for** $t = 1, 2, \ldots T$ **do**
  **for** $y(t) = 0, 1, \ldots, K$ **do**
  Calculate the $d(y(t-1), y(t))$, where $C_{y(t)}(x(t), v(t))$ is given by Algorithm 1.
  **end for**
**end for**
3: Find the shortest path among the paths from $y(0)$ to $y(T) = \{0, 1, \ldots, K\}$, respectively. The constituent vertices $y(1), \ldots, y(T)$ of the shortest path are exactly the number of running heating plants over time frame $\mathcal{T}$, and the sum of weights of this path is the total social operational cost.

As discussed above, the time complexity of Algorithm 1 is $O(M^2)$, hence the complexity of `for` loops is $O(\Gamma K M^2)$. And the complexity of shortest path finding—we utilize Dijkstra's algorithm in this work—is $O((K^2 + \Gamma K)\log(\Gamma K))$. In summary, the complexity of Algorithm 2 is $O(\Gamma K M^2 + (K^2 + \Gamma K)\log(\Gamma K))$. Here $\Gamma$ represents the number of time slots during which Algorithm 2 has knowledge of heat demand, fuel cost, and datacenters' bids. In offline scenario, all the information over time frame $\mathcal{T}$ is given at the beginning of $\mathcal{T}$, hence $\Gamma = T$ where $T$ represents the number of time slots over $\mathcal{T}$. However, in online scenario, the online framework has to make decisions with limited near-future information. As Algorithm 2 is a component of the online framework, which is discussed in Sec. 3, $\Gamma$ is shorter than $T$.

The complexity seems high but is affordable in practice — $M$, $K$, and $\Gamma$ is small in real condition, making the computing overhead acceptable: As listed in Table 1, the number of colocation datacenters $M$ in one city is no more than a few dozen. Moreover, Fig. 3 shows that the number of heating plants $K$ is no larger than 10. In offline scenario, $\Gamma$ is usually 24 in one-day time frame. However, $\Gamma$ is much smaller in real condition. In online scenario, the value of $\Gamma$ is limited by the level of prediction accuracy, as discussed in the following Sec. 4.2. Moreover, as shown in the following Fig. 15, a much longer look-ahead window would achieve little performance improvement. In our work, $\Gamma$ is no larger than 7 in online scenario. In Sec. 6, we will discuss the run time of the framework and show that the complexity of Algorithm 2 is affordable in practice.

THEOREM 2. *The approximation ratio of Algorithm 2 to* `CMP` *in the offline setting is* 2.

PROOF. Let $d^*(y(t-1), y(t))$ be the optimal weight of edge $(y(t-1), y(t))$, $C^*_{y(t)}(\boldsymbol{x}(t), v(t))$ be the optimal value of the problem (6). Note that $C_{y(t)}(\boldsymbol{x}(t), v(t))$ is calculated by Algorithm 1, we have

$$\begin{aligned} d(y(t-1), y(t)) &= C_{y(t)}(\boldsymbol{x}(t), v(t)) + \alpha y(t) + s(y(t-1), y(t)) \\ &\leqslant 2C^*_{y(t)}(\boldsymbol{x}(t), v(t)) + \alpha y(t) + s(y(t-1), y(t)) \\ &\leqslant 2\Big\{C^*_{y(t)}(\boldsymbol{x}(t), v(t)) + \alpha y(t) + s(y(t-1), y(t))\Big\} \\ &\leqslant 2d^*(y(t-1), y(t)) \end{aligned} \tag{11}$$

Let the sum of weights of shortest path over time frame $[1, T]$ is $\sum_{t=1}^{T} d^*(y(t-1), y(t))$, and hence the weights given by Algorithm 2 be no more than $2\sum_{t=1}^{T} d^*(y(t-1), y(t))$. (If Algorithm 2 choose

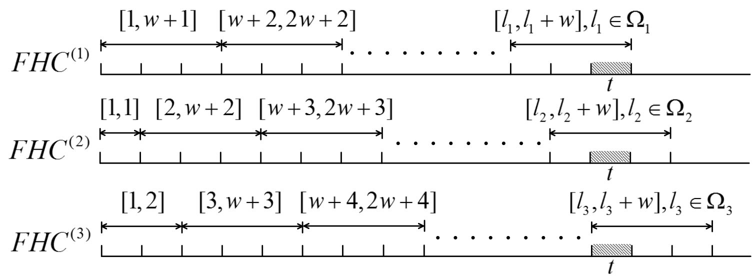


Fig. 7. An illustration of Fixed Horizon Control (FHC).

another path rather than the actually shortest, the summation of weights of this path should be less than $2\sum_{t=1}^{T} d^*(y(t-1), y(t))$, otherwise the actually shortest path must be chosen.)

Above all, the approximation ratio of Algorithm 2 is 2. □

## 4.2 Online Algorithm to CMP

We now present the online algorithm that decomposes the long-term CMP over time frame $\mathcal{T}$ into a series of deterministic problems over its look-ahead window, based on the technique of RHFC [62] and the offline algorithm proposed in Sec. 4.1.

In practice, through time series forecasting or artificial neural networks, the near-future information (*e.g.*, heat demand, fuel cost and the power consumption of datacenters) within a look-ahead window consisting of $w$ time slots, can be predicted to a good level of accuracy. In order to fully utilize such near-future information to make better online decisions, at each time slot $t$, we assume that the future bids, till to time slot $t + w$, have been submitted to the DHS by each datacenter in advance, based on the prediction of its power consumption. That is, at time slot $t$, we can utilize $w + 1$ time slots' information in total.

We begin by describing the fixed horizon control (FHC) briefly. As depicted in Fig. 7, in the first time frame $[1, w + 1] \subseteq [1, T]$, we could initialize $w + 1$ time slots starting from $t = 1, 2, \ldots, w + 1$, respectively, the algorithm starting from time slot $p \in \{1, 2, \ldots, w + 1\}$ is denoted by $FHC^{(p)}$, and denote the set of starting time slots of $FHC^{(p)}$ as $\Omega_p = \{l | l \mod (w + 1) = p, l \in \mathbb{Z}, 1 \leqslant l \leqslant T\}$, where $p \in \mathbb{Z}$, and $p \in [1, w + 1]$. For any time slot $l \in \Omega_p$, $\forall p \in [1, w + 1]$, given the future information which can be precisely predicted ahead of $w$ time slots and $y(t - 1)$, we can solve the following social operational cost minimization by Algorithm 2, and get $\boldsymbol{x}^{(p)}(t)$, $y^{(p)}(t)$, $v^{(p)}(t)$ from $t = l$ to $t = l + w$:

$$\begin{aligned} \min \quad & \sum_{t=l}^{l+w} C_{y(t)}(\boldsymbol{x}(t), v(t)) + \sum_{t=l}^{l+w} \Big(\alpha y(t) + s(y(t-1), y(t))\Big) \\ \text{s.t.} \quad & y(t) \in \{0, 1, \ldots, K\},\ \forall t \in [l, l + w] \end{aligned} \tag{12}$$

The RFHC used in our online solution Algorithm 3 below is adapted from FHC for determining the number of running heating plants at each time slot. We randomly choose an integer variable $p \in [1, w + 1]$ with equal probability to initialize look-ahead window, and utilize $FHC^{(p)}$ to set $y(v) = y^{(p)}(v)$ for every time slot $v \in [1, T]$. Unlike the FHC, RFHC randomly selects the start time of look-ahead window to defeat strategic adversaries who choose a small number of running heating plants at the beginning of each look-ahead window to inflate switching cost. Moreover, RFHC achieves low expected worst case performance as we will prove in Theorem 3.

**ALGORITHM 3:** The Online Algorithm Framework

```
Initialization: let ν = 1, y(0) = 0, and randomly choose p ∈ [1, w + 1] with equal probability
while ν ⩽ T do
  if ν = 1 then //Submit bids
    Datacenters submit their bids of time slot ν = 1 to ν = w + 1;
  else
    Datacenters submit their bids of time slot ν + w;
  end if
  //Determine the number of running heating plants at time ν by RFHC as follows
  if ν = 1 and p ⩾ 2 then
    Obtain x^(p)(t), y^(p)(t) and v^(p)(t), t = 1,...,p − 1 by solving the problem (12) from t = 1 to t = p − 1 through Algorithm 2;
  else if p ⩽ ν ⩽ T and ν ∈ Ω_p then
    Obtain x^(p)(t), y^(p)(t) and v^(p)(t), t = ν, ν + 1,...,min{ν + w, T} by solving the problem (12) from t = ν to t = min{ν + w, T} through Algorithm 2;
  end if
  The DHS determines winners and designs payments of time slot ν with given y(ν) for the one-round auction (6) in Sec. 4.1, by applying the randomized auction framework in Sec. 5;
  ν = ν + 1;
end while
```

In Algorithm 3, we have decomposed the long-term online auction problem into a series of one-round auction problems at each time slot. By further applying the truthful and 2-approximation randomized auction framework that will be presented in Sec. 5, we further show the performance of the online Algorithm 3 as follows:

Theorem 3. *Algorithm 3 is a* $(2 + \frac{2\beta}{(w+1)\alpha})$*-competitive solution to* CMP *in the online setting,* i.e., *the social operational cost obtained by Algorithm 3 is at most* $(2 + \frac{2\beta}{(w+1)\alpha})$ *times the optimal social operational cost in offline setting.*

We prove the theorem in Appendix A.2.

## 5 THE RANDOMIZED AUCTION FRAMEWORK

In this section, we design a truthful 2-approximate auction mechanism for the one-round auction problem, based on the randomized auction framework. For simplicity, we omit the time index $t$ due to its independence.

Note that utilizing the VCG auction mechanism requires selecting the set of items (i.e., bids and the amount of self-produced heat) with maximum valuation. As mentioned in Sec. 4.1, the problem (6) is NP-hard and we cannot obtain the set of items which minimizes the objective value in polynomial time. That is, the VCG auction is not applicable to our work.

### 5.1 Randomized Auction Mechanism

The basic idea of the randomized auction framework is as follows. After computing the optimal fractional solution to LPR, we take a linear programming duality based decomposition technique to decompose the fractional solution into a convex combination of feasible solutions to the original problem. Then, one of the feasible solutions is chosen randomly, with its weight in the convex combination taken as the probability. Finally, we compute the truthful payments to the winners according to a well-established payment rule [2].

The auction algorithm is shown in Algorithm 4, and the details of the mechanism are given blow:

**ALGORITHM 4:** A Randomized Auction Mechanism

1: *Calculate the fractional optimal solution.*

Obtain the fractional optimal solution $(\boldsymbol{x}^*, v^*)$ of LPR.

2: *Decompose the optimal fractional solution.*

Decompose the optimal fractional solution $(\boldsymbol{x}^*, v^*)$ to a set of feasible mixed integer solution $(\boldsymbol{x}^k, v^k)$, $k \in \mathcal{K}$ of CMP with convex decomposition technique.

3: *Select winners and calculate payments.*

Randomly select winners, which is indicated by $\boldsymbol{x}^k$, with probability $\lambda_k$ from set $\mathcal{K}$.
The payment for datacenter $i$ is

$$r_i = \begin{cases} c_i + \frac{\int_{c_i}^{\zeta} \min\{\gamma x_i^*(c, c_{-i}), 1\}}{\min\{\gamma x_i^*(c_i, c_{-i}), 1\}} & , \quad x_i = 1 \\ 0 & , \quad x_i = 0 \end{cases}$$

where $\zeta = p_h u_i + \lceil \frac{u_i}{V} \rceil (\alpha + \beta)$, here $\lceil x \rceil$ indicates the smallest integer no less than $x$.

**Step 1:** *Calculating the fractional optimal solution.* Solve the LPR, obtain the optimal fractional winner decision $\boldsymbol{x}^*$, and the optimal amount of self-produced heat $v^*$.
**Step 2:** *Decomposing the optimal fractional solution.* As the optimal fractional solution is infeasible for DHS to choose winner, we decompose it into a convex combination of feasible solutions each with a fractional weight that sums up to 1. This step requires a computationally efficient decomposition algorithm with a good approximation ratio to obtain the feasible solutions, satisfying:

$$\sum_{i \in \mathcal{I}} c_i x_i + p_h v \leqslant \gamma \Big( \sum_{i \in \mathcal{I}} c_i x_i^* + p_h v^* \Big), \tag{13}$$

where $\gamma$ is the approximation ratio.

Our basic idea is to employ the convex decomposition technique to obtain a set of feasible solutions $(\boldsymbol{x}^k, v^k)$ and $\lambda_k$, $\forall k \in \mathcal{K}$, such that $\sum_{k \in \mathcal{K}} \lambda_k \boldsymbol{x}^k \leqslant \gamma \boldsymbol{x}^*$, $\sum_{k \in \mathcal{K}} \lambda_k v^k \leqslant \gamma v^*$ and $\sum_{k \in \mathcal{K}} \lambda_k \leqslant 1$ are all satisfied. Moreover, to ensure the decomposition is feasible, the constraints $x_i^k \leqslant 1$ and $v^k \leqslant Vy$ should also be satisfied. Furthermore, we have $\sum_{k \in \mathcal{K}} \lambda_k x_i^k \leqslant \min\{\gamma x_i^*, 1\}$ and $\sum_{k \in \mathcal{K}} \lambda_k v^k \leqslant \min\{\gamma v^*, Vy\}$ since $\sum_{k \in \mathcal{K}} \lambda_k \leqslant 1$. The weight $\lambda_k$, $k \in \mathcal{K}$, can be computed by solving the following linear program:

$$\begin{aligned} \max \quad & \sum_{k \in \mathcal{K}} \lambda_k \\ \text{s.t.} \quad & \sum_{k \in \mathcal{K}} \lambda_k x_i^k = \min\{\gamma x_i^*, 1\}, \ \forall i \in \mathcal{I} \\ & \sum_{k \in \mathcal{K}} \lambda_k v^k \leqslant \min\{\gamma v^*, Vy\}, \\ & \sum_{k \in \mathcal{K}} \lambda_k \leqslant 1, \ \lambda_k \geqslant 0, \ \forall k \in \mathcal{K} \end{aligned} \tag{14}$$

where the exact decomposition $\sum_{k \in \mathcal{K}} \lambda_k x_i^k = \min\{\gamma x_i^*, 1\}$ guarantees that the winning probability of datacenter satisfies the truthfulness condition in Theorem 6.

Though the above linear program can be solved directly, the exponential number of variables make the computation process slow. Instead, by introducing the dual variables $\boldsymbol{\mu}$, $\sigma$, and $\tau$, we formulate the dual problem as:

$$\begin{aligned} \min \quad & \sum_{i\in\mathcal{I}} \min\{\gamma x_i^*, 1\}\mu_i + \min\{\gamma v^*, Vy\}\sigma + \tau \\ \text{s.t.} \quad & \sum_{i\in\mathcal{I}} x_i^k \mu_i + v^k\sigma + \tau \geqslant 1,\ \forall k \in \mathcal{K},\ \sigma,\tau \geqslant 0. \end{aligned} \tag{15}$$

Note that the dual problem has an exponential number of constraints. Applying Algorithm 1, we obtain a separation oracle for the dual problem and then get a polynomial number of separating hyperplanes. Hence we can solve the dual in polynomial time utilizing the ellipsoid method and get a polynomial number of feasible solution to (6). Furthermore, the primal decomposition linear programming (14) can be solved in polynomial time since it reduces to a linear program with a polynomial number of variables corresponding to its dual solution.

Theorem 4. *The convex decomposition problem (14) can be solved within polynomial time and the optimal objective value is* $\sum_{k\in\mathcal{K}} \lambda_k = 1$.

We prove the theorem in Appendix A.3.

**Step 3:** *Selecting winners and calculating payments.* Let $P_i(c_i)$ be the probability datacenter $i$ wins in the auction with cost $c_i$, and the expectation of $x_i$ is $E[x_i] = P_i(c_i) \times 1 + (1 - P_i(c_i)) \times 0 = P_i(c_i)$. Moreover, since we utilize the convex decomposition technique, we have $P_i(c_i) = E[x_i] = \sum_{k\in\mathcal{K}} \lambda_k x_i^k = \min\{\gamma x_i^*, 1\}$. Let $c_{-i}$ be all bids except for bid $i$. The payment for the winning datacenter $i$ is given by:

$$r_i = c_i + \frac{\int_{c_i}^{\zeta} P_i(c)dc}{P_i(c_i)} = c_i + \frac{\int_{c_i}^{\zeta} \min\{\gamma x_i^*(c, c_{-i}), 1\}dc}{\min\{\gamma x_i^*(c_i, c_{-i}), 1\}},$$

where $\zeta = p_h u_i + \lceil \frac{u_i}{V} \rceil(\alpha + \beta)$ ($\lceil x \rceil$ indicates the smallest integer no less than $x$), which will be discussed in the following Theorem 6. And $x_i^*(c, c_{-i})$ is the optimal solution to (6) where the cost of $i$'s bid is $c_i$ and others' are $c_{-i}$.

### 5.2 Performance Analysis

To prove the payment scheme achieves truthfulness, we first introduce a sufficient and necessary condition for a reverse auction to be truthful.

Theorem 5. *[2] A randomized auction with bids* $\mathbf{c}$ *and payment* $\mathbf{r}$ *is truthful in expectation iff: i)* $P_i(c_i)$ *is monotonically nonincreasing in* $c_i$, $\forall i \in \mathcal{I}$. *ii)* $\int_0^{\infty} P_i(c)dc < \infty$. *iii) The expected payment satisfies* $E[r_i] = c_i P_i(c_i) + \int_{c_i}^{\infty} P_i(c)dc$, $\forall i \in \mathcal{I}$.

Based on this condition, we now prove the performance of the proposed auction mechanism, in terms of individual rationality, social efficiency and truthfulness, in the following Theorem 6.

Theorem 6. *The proposed reverse auction mechanism with payment* $r_i = c_i + \frac{\int_{c_i}^{\zeta} P_i(c)dc}{P_i(c_i)}$, $\forall i \in \mathcal{I}$ *achieves:*

*i) individual rationality;*

*ii) 2-approximation in social operational cost;*

*iii) truthfulness in expectation.*

We prove the theorem in Appendix A.4.

Note that as there are multiple options for datacenter to meet its overall electricity demand. The unit price of power varies across datacenters and across the temporal domain. For example, power price in the electricity wholesale market, known as *locational marginal pricing*, is set based on the

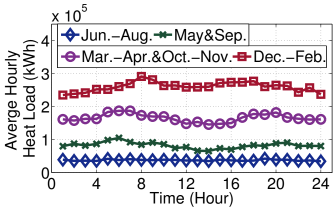


Fig. 8. The average hourly heat load for four different seasons in a district heating system.

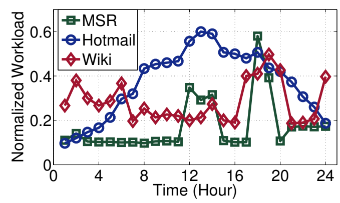


Fig. 9. The three normalized workloads trace data for different datacenters.

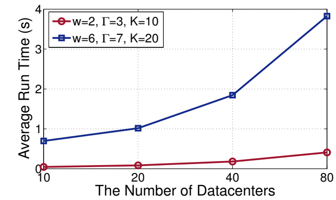


Fig. 10. The average run time of the online framework.

value of the power at the specific location and delivering time [48]. Furthermore, on-site renewable generation (*e.g.*, wind and solar power) and energy storage (*e.g.*, battery and super-capacitor [37]) make the power cost of datacenters different. There exist spatial and temporal variations in the unit cost of heat harvesting, and hence it is necessary to guarantee the truthfulness of our market mechanism.

## 6 PERFORMANCE EVALUATION

### 6.1 Simulation Setup

We simulate a geographical region where a DHS harvests waste heat from 8 datacenters. Each datacenter's capacity is set to $m = 6 \times 10^4$ homogeneous servers. The static power and computing power of those servers are set to $P_s = 100$ W and $P_c = 150$ W, respectively, reflecting state-of-art levels [60]. The DHS is equipped with $K = 10$ homogeneous heating plants, each with a capacity of $V = 50$ MW. Following the recent report on fuel cost and heating plants [19], we set fuel cost $p_h = 2.24$ cents/kWh, maintenance cost $\alpha = 7.5$ \$/hour, and switching cost to the fuel cost of running each heating plant at full capacity for 5 hours, *i.e.*, $\beta = \$3925$ [62]. The simulation runs on a server which consists of 2×Xeon 2.30 GHz CPU with 16 cores and 64 GB of RAM.

**Heat demand:** We use the one-day hourly heat demand traces, from a DHS in Sweden [25], as shown in Fig. 8. It shows significant variations of heat demand across different seasons. The level of heat demand affects the amount of harvested waste heat and the number of running heating plants. And hence it has an effect on the level of social operational cost as well as datacenters' payment and utility. As a result, in our evaluation, we not only show the difference among different datacenters, but also illustrate the difference among different seasons. For simplicity, we denote the seasons "Jun.-Aug.", "May&Sep.", "Mar.-Apr.&Oct.-Nov.", and "Dec.-Feb." as "Season 1" to "Season 4" in the following discussion respectively.

**Workloads:** We collect 3 types of one-day hourly workload traces from three representative internet services, Hotmail, Wikipedia (Wiki), and Microsoft Research (MSR) [60], respectively, as depicted in Fig. 9. We normalize these workloads with respect to each datacenter's capacity and duplicate them while scaling up or down about 20% randomly to generate these datacenters' workloads. Datacenters 1, 4, and 7 process workload of Hotmail, datacenters 2, 5, and 8 process workload of Wiki and the other process workload of MSR.

**Datacenters' bids and operational cost:** We consider all of the power consumed by IT equipment is converted to waste heat and can be removed [49]. We consider all of the power consumed by IT equipment is converted to waste heat and can be removed. In our simulations, we use the power model $u = m(P_s + P_c \frac{\lambda}{m\mu})$ to compute the power consumption of IT equipment, here $\mu$ is the service rate of server and $\lambda$ is the workload arrival rate. The parameter $\mu$ is set according to [50]. We further assume all datacenters equip with heat recovery units whose COP (coefficient of performance, the ratio of the heat it recovers to the power it consumes) are all 4.0 [10]. We set the electricity price $p_e$ based on the hourly real-time locational marginal pricing data of PJM [47] on January 1, 2017. The operational cost of recovering waste heat $u$ is $c = \frac{u}{COP} \times p_e$.

**Benchmarks:** To validate the effectiveness of CloudHeat, we compare its performance with that of another four benchmarks: 1) the offline mixed integer programming (MIP) optimum, for which `CMP` with offline setting is exactly solved by CPLEX. 2) The offline randomized auction, for which the full future information across the time frame $\mathcal{T}$ is assumed. 3) The offline no auction, for which the DHS does not harvest heat from the datacenters and minimize its operational cost in an offline manner. 4) The online greedy strategy with $w = 2$, for which the DHS chooses those heat resources with lowest costs until the demand is covered in each time slot, and switches on $n$ heating plants at time slot $t$ only when all $n$ heating plants are also needed in time slot $t + 1$ and $t + 2$, or switches off $n$ heating plants only when all $n$ heating plants are all unnecessary in both time slot $t + 1$ and $t + 2$.

### 6.2 Evaluation Results

**Average run time:** We first study the run time of our framework. The computing overhead of the proposed framework is mainly related to $\Gamma$ which represents the number of time slots during which the framework has knowledge of heat demand, fuel cost, and datacenters' bids ($\Gamma = w + 1$), the number of heating plants $K$, and the number of volunteer datacenters $M$. We set $\Gamma = 3$ and $K = 10$, and vary the number of datacenters from 10 to 80 since its value shows large discrepancy between different cities, as listed in Table 1. The red line in Fig. 10 shows the average run time of the online framework under the different number of datacenters. We find that the framework only takes less than 0.4 seconds to finish. We also increase the length of look-ahead window $w$ to 6 (i.e., $\Gamma = 7$) and increase the number of heating plants $K$ to 20, which are so large that they rarely happen in real condition. We observe that the framework takes no more than 4 seconds. For wholesale power market, the length of one time slot is usually set to one hour or larger. Hence the run time of the online framework is affordable in practice.

**Social operational cost reduction**: Fig. 11 compares the social operational cost achieved by the proposed online randomized auction (whose length of a look-ahead window is 2), as well as that achieved by the other four benchmarks. We have the following observations: first, compared to the offline no auction benchmark, the proposed auction-based solution CloudHeat can effectively reduce the social operational cost, demonstrating the economical efficiency of CloudHeat. Second, compared with online greedy strategy, CloudHeat also shows a great effectiveness in reducing operational social cost when the future information has limited predictability. Third, compared with the offline mixed integer programming optimum, *i.e.*, the theoretical optimal social operational cost, CloudHeat only incurs a small loss in social operational cost and performs much better than theoretical bound given by Theorem 3, which indicates that our solution can provide a small competitive ratio in practice. Finally, comparing the offline mixed integer programming optimum to offline randomized auction, we find that their performance is very close, indicating that randomness introduced in CloudHeat would not hurt the social efficiency too much.

**Meeting heat demand:** To understand how the heat demand is satisfied by datacenter heat harvesting, in Fig. 12, we plot the amount of heat harvested from each datacenter together with the amount of self-produce heat by the DHS for different seasons. We find that in Season 1 and Season 2, the relatively low heat demand can be almost entirely served by harvesting the cheap datacenter waste heat. However, when the heat demand bursts in Season 3 and Season 4, the heating plants will be turned on to cover the remainder heat demand after all the datacenter waste heat has been harvested. These demonstrate that by harvesting datacenter waste heat, the heat demand can be satisfied in an adaptive and cost-efficient manner.

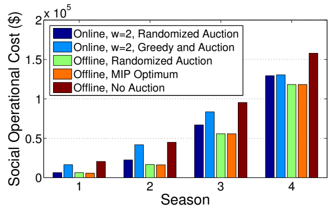


Fig. 11. Comparisons of social operational cost with different algorithms in different seasons.

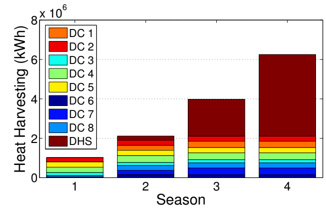


Fig. 12. Comparison of heat harvesting by all datacenters and DHS in different seasons.

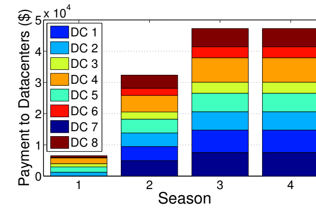


Fig. 13. Comparisons of the payments to all datacenters in different seasons.

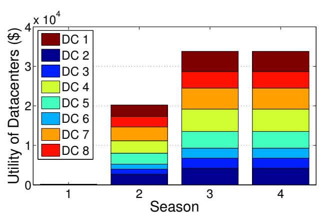


Fig. 14. Comparison of the utilities received by all datacenters in different seasons.

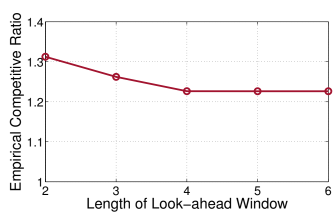


Fig. 15. Comparison of empirical average competitive ratio with different length of look-ahead window.

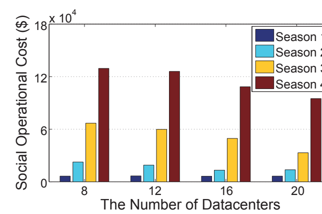


Fig. 16. Comparison of the social operational cost with different number of participant datacenters.

**Datacenters' utilities:** We further study the daily payment and utility received by each datacenter in different seasons, as illustrated in Fig. 13 and Fig. 14, respectively. As expected, each datacenter receives a non-negative utility from the waste heat auction, demonstrating the individual rationality ensured by CloudHeat (proved in Theorem 6). Furthermore, we observe that the disparity of total utilities in different seasons is much more significant than that of the total payments in different seasons, the rationale behind this difference is automatically adapting the price based on the realtime supply-demand unbalance. For example, in Season 1 when there is little heat demand, then the DHS buys small amount of heat from datacenter with low price, thus datacenters receive low utilities. While in Season 4 when the heat demand peaks, the DHS buys all the heat with high price, giving high utilities to the datacenters. Interestingly, we also observe that though the heat demand in Season 4 far overweighs that in Season 3, the total payments and utilities of the datacenter in Season 3 and Season 4 are almost the same with the fact that the DHS buys all the datacenter waste heat when the heat demand bursts.

**Influence of the length of look-ahead window:** Intuitively, the length of look-ahead window, *i.e.*, how long of the accurate future information we possess, has a remarkable influence on the performance of the online auction. Fig. 15 plots the average competitive ratio under different lengths of look-ahead window. Excitingly, we find that when varying the length from 2 to 6, the competitive ratio never exceeds 1.32, demonstrating the great empirical performance of the online auction. Basically, a longer look-ahead window would lead to better performance.

**Influence of waste heat supply:** We next study the influence of the amount of datacenter waste heat supply, by varying the number of participated datacenters. As we can see from Fig. 16, as the number of participated datacenter increases, the social operational cost reduces dramatically in seasons with large amounts of heat demand. While for Season 1 when there is little heat demand, the social operational cost does not vary much as the number of participated datacenter changes,

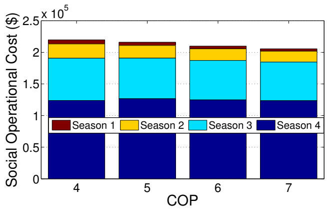


Fig. 17. Comparison of the social operational cost with different COP.

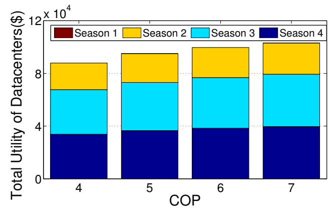


Fig. 18. Comparison of the total utility received by datacenters with different COP.

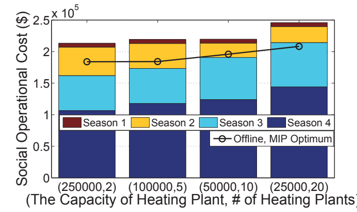


Fig. 19. Comparison of the social operational cost with different capacity and the number of heating plant.

since that the amount of heat demand can be satisfied with few datacenters. This suggests that, in practice, the social operational cost in cold days can be reduced by incentivizing more datacenters to participate in waste heat harvesting.

**Influence of recovery efficiency:** We further study the influence of recovery efficiency, by varying value of COP. We find that the improvement of recovery efficiency has a positive impact on social operational cost reduction and datacenters' utility improvement. As shown in Fig. 17 and Fig. 18, as COP increases from 4 to 7, the social operational cost decreases by 6% and it also provides a 17% increase in datacenters' total utility. This suggests that recovering heat with power-efficiency recovery units not only reduces datacenters' power consumption but also boosts their utilities in heat harvesting.

**Influence of heating plant's capacity:** We finally study the influence of heating plant's capacity. We decrease the capacity from $250,000$ W to $25,000$ W, and the number of heating plants is varied accordingly to ensure total capacity is constant (500 MW in total). As illustrated in Fig. 19, we find that the larger capacity can result in less social operational cost in Season 3 and Season 4. As in these seasons with large heat demand, heating plants with low capacity are sensitive to the variation of heat demand with limited length of look-ahead window. Specifically, if the capacity of heating plant is low, some heating plants are likely to be opened and closed frequently with the amount of heat demand varying. However, when the level of the amount of heat demand is low, such as in Season 2, less number of heating plants are sufficient to meet the heat demand. Moreover, opening heating plants with lower capacity causes lower start-up cost. Hence the social operational cost of Season 2 decreases when the capacity of heating plant reduces. For Season 1, the social operational cost is constant as cheap waste heat provided by datacenters is sufficient to meet total heat demand. We also find that although the number of heating plants increases tenfold (*i.e.*, the number of Algorithm 2's vertices increases hundredfold), the competitive ratio of our online framework does not vary significantly.

## 7 RELATED WORK

There have been many studies to minimize datacenter energy consumption and energy cost, including energy proportionality [26], capacity provisioning [35], workload management [1, 64], and the usage of energy storage [41] or renewable energy [13–16]. In contrast with these studies that view datacenter energy consumption as a negative, we turn datacenters into a socially valuable asset by coordinating datacenter waste heat harvesting to supplement DHS heat supplies. Additionally, our study also complements the existing energy efficiency literature by reducing datacenter's cooling energy through heat harvesting.

Waste heat recovery has been quickly emerging in the literature. Leveraging thermoelectric generator, Lee et al. [33] propose several real systems which harvest heat from CPU, and reuse the waste heat to power the fan or thermoelectric cooler to cool CPU simultaneously. Ebrahimi et al. [18] review state-of-the-art datacenter heat recovery techniques, demonstrate the feasibility and benefits of the coordination between datacenters and DHS. Liu et al. [36] propose “data furnace”, an approach that provides heat to residential buildings through running servers placed locally. Ward et al. [58] design a scaled-down prototype to recover waste heat from datacenters to heat local facilities. The last three studies all focus on improving and/or demonstrating the heat harvesting from a single datacenter and can be used by participating datacenters in our study. Nonetheless, our work is complementary and focuses on market designs to coordinate the heat harvesting from multiple datacenters.

Market-based energy efficiency programs, such as datacenter demand response [60, 61, 63, 66] and bilateral power trade between smart grids and datacenters [65], also have opportunities for cost reduction. Most mechanisms rely on reference-based prices [29], supply bidding functions [32], and reverse auction [60, 63]. These studies are not applicable to datacenter heat harvesting, because of limitations like imbalance of supply and demand, non-truthful behaviors, computational complexity, and/or lack of long-term performance guarantees. For example, leveraging *supply function bidding*, Islam et al. [32] propose an incentive mechanism for handling power emergencies, but the supply function bidding-based mechanism cannot guarantee truthfulness. Further, its efficiency can be manipulated by the market power of each tenant, and the fairness can not be guaranteed when market power of tenants is disparate.

*Reverse auctions* represent another promising market mechanism. Zhang et al. [60] propose a reverse auction mechanism for demand response in colocation datacenters. Zhou et al. [63] extend it to an online version in the scenario of smart grids with limited energy storage. Although our underlying auction design shares similarity with [60, 63], our study focuses on datacenter heat harvesting and differs from [60, 63] in: i) we propose a new primal-dual algorithm for the auction design; ii) we use a different technique towards online bid processing.

## 8 CONCLUSION

Datacenter waste heat harvesting is envisioned as a promising approach for mitigating operational availability and reliability issues faced by modern district heating systems. This work studied how district heating systems could harvest waste heat from datacenters at the minimum long-term social operational cost. To incentivize participation from datacenters, we proposed CloudHeat, a first-of-its-kind reverse online auction based solution. By blending the advantages of randomized fixed horizon control (RFHC) and a randomized auction framework, CloudHeat guaranteed polynomial running time, truthfulness (in expectation), and bounded long-term social operational cost.

## A APPENDIX

### A.1 Poof of Theorem 1

Proof. We prove this theorem by four cases as follows:

Case 1: If there is no set $\mathcal{X} \subseteq \mathcal{I}$ such that $\Delta\mathcal{X} > 0$, *i.e.,* the total heat demand can be met by any of participate datacenter's bid. In this case, we can obtain the optimal solution by comparing the cost between: $i$) purchasing recovered heat from datacenters who submit bids and $ii$) producing $W$ units of heat by heating plants, and finding the minimum one.

Then we discuss the case where the number of set $\{\mathcal{X} | \mathcal{X} \subseteq \mathcal{I}, \Delta\mathcal{X} > 0\}$ is no less than 1.

Case 2: If $v = 0$, it means that the constraint (9) never goes tight, and thus $\eta = 0$. The objective value of LPR is given by:

$$\sum_{i \in \mathcal{I}} c_i x_i = \sum_{i \in S} c_i = \sum_{i \in S} \sum_{\mathcal{X} \subseteq \mathcal{I}: i \notin \mathcal{X}, \Delta\mathcal{X} > 0} u_i(\mathcal{X}) z(\mathcal{X}) = \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \sum_{i \in S \setminus \mathcal{X}} u_i(\mathcal{X}).$$

According to the definition of $u_i(\mathcal{X})$, we have

$$\sum_{i \in S \setminus \mathcal{X}} u_i(\mathcal{X}) \leqslant \sum_{i \in\ S \setminus \{l\}} u_i - \sum_{i \in \mathcal{X}} u_i + u_l(\mathcal{X}) < W - \sum_{i \in \mathcal{X}} u_i + u_l(\mathcal{X}) = \Delta\mathcal{X} + u_l(\mathcal{X}) \leqslant 2\Delta\mathcal{X}, \tag{16}$$

where $l$ denotes the last bid added to the set $S$, hence we have $\sum_{i \in S \setminus \{l\}} u_i < W$. Plug the inequality (16) in the above objective function, we have

$$\sum_{i \in \mathcal{I}} c_i x_i = \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \sum_{i \in S \setminus \mathcal{X}} u_i(\mathcal{X}) \leqslant 2 \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Delta\mathcal{X} \leqslant 2OPT_{LPR} \leqslant 2OPT,$$

where the $OPT_{LPR}$ is the optimal objective value of LPR, and the $OPT$ is the optimal objective value of the problem (6).

Case 3: If $v > 0$ and $\eta = 0$, it means that the constraint (9) goes tight and then the iteration stops, *i.e.,* the maximum heating capacity is no less than the remainder demand $\Delta S$. Hence the objective value of LPR consists of two parts: $\sum_{i \in \mathcal{I}} c_i x_i$ and $p_h v$. For the latter part, we have:

$$p_h v = \Delta S \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}). \tag{17}$$

Notice that for all $\mathcal{X} \subseteq \mathcal{S} : \Delta\mathcal{X} > 0$ and $z(\mathcal{X}) > 0$, we have $\Delta\mathcal{X} \geqslant \Delta S$. Hence

$$p_h v \leqslant \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Delta\mathcal{X} \leqslant OPT_{LPR}. \tag{18}$$

For the former part, same as the Case 2, we have

$$\begin{aligned}\sum_{i \in \mathcal{I}} c_i x_i = \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \sum_{i \in S \setminus \mathcal{X}} u_i(\mathcal{X}) &\leqslant \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Big(\sum_{i \in S} c_i - \sum_{i \in \mathcal{X}} c_i\Big) \\ < \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Big(W - \sum_{i \in \mathcal{X}} c_i\Big) &= \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Delta\mathcal{X} \leqslant OPT_{LPR}.\end{aligned}$$

Sum the above two inequalities, we have $\sum_{i \in \mathcal{I}} c_i x_i + p_h v \leqslant 2OPT_{LPR} \leqslant 2OPT$.

Case 4: If $v > 0$ and $\eta > 0$, it means that the constraint (9) goes tight once while the maximum heating capacity cannot meet the remainder demand $\Delta S$ (*i.e.,* $Vy < \Delta S$). Similar to Case 3, the the objective value of LPR consists of two parts: $\sum_{i \in \mathcal{I}} c_i x_i$ and $p_h v$. For the former part, we have

$$\sum_{i \in \mathcal{I}} c_i x_i = \sum_{i \in S} c_i = \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \sum_{i \in S \setminus \mathcal{X}} u_i(\mathcal{X}) \leqslant \sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) \Bigg(\sum_{i \in S \setminus l} u_i - \sum_{i \in \mathcal{X}} u_i + u_l(\mathcal{X})\Bigg),$$

where the definition of $l$ is same as the above case. Then for the latter part, since the dual variable $\eta$ increases as the same rate of $z(\mathcal{X})$ after (9) becomes tight, we have

$$p_h v = \Bigg(\sum_{\mathcal{X} \subseteq \mathcal{I}: \Delta\mathcal{X} > 0} z(\mathcal{X}) - \eta\Bigg) Vy. \tag{19}$$

Summing the above two inequality, we have

$$
\begin{aligned}
\sum_{i\in\mathcal{I}} c_i x_i + p_h v &\leqslant \sum_{\mathcal{X}\subseteq\mathcal{I}:\Delta\mathcal{X}>0} z(\mathcal{X})\Big(\sum_{i\in S\setminus l} u_i + Vy - \sum_{i\in\mathcal{X}} u_i + u_l(\mathcal{X})\Big) - Vy\cdot\eta \\
&< \sum_{\mathcal{X}\subseteq\mathcal{I}:\Delta\mathcal{X}>0} z(\mathcal{X})\Big(W - \sum_{i\in\mathcal{X}} u_i + u_l(\mathcal{X})\Big) - Vy\cdot\eta \\
&\leqslant \sum_{\mathcal{X}\subseteq\mathcal{I}:\Delta\mathcal{X}>0} z(\mathcal{X})\Big(\Delta\mathcal{X} + u_l(\mathcal{X})\Big) - Vy\cdot\eta \\
&\leqslant \sum_{\mathcal{X}\subseteq\mathcal{I}:\Delta\mathcal{X}>0} z(\mathcal{X})\Big(\Delta\mathcal{X} + \Delta\mathcal{X} - Vy\Big) - Vy\cdot\eta \\
&\leqslant 2\sum_{\mathcal{X}\subseteq\mathcal{I}:\Delta\mathcal{X}>0} z(\mathcal{X})\Delta\mathcal{X} - 2Vy\cdot\eta = 2OPT_{LPR} \leqslant 2OPT,
\end{aligned}
$$

where the fourth step uses the fact that the contribution of the last bid is $u_l(\mathcal{X}) = W-\sum_{i\in\mathcal{X}} u_i - Vy = \Delta\mathcal{X} - Vy$, and the fifth step uses the fact that the number of set $\{\mathcal{X}|\mathcal{X}\subseteq\mathcal{I}, \Delta\mathcal{X}>0\}$ is no less than 1.

Above all, the Algorithm 1 achieves 2-approximation ratio to the problem (6). □

## A.2 Proof of Theorem 3

Lemma 2. *Let $cost(OPT)$ denote the offline optimal social operational cost of CMP, $cost(FHC^{(p)})$ is the cost obtained by $FHC^{(p)}$ through Algorithm 2. We have $cost(FHC^{(p)}) \leqslant 2\Big\{cost(OPT) + \sum_{l\in\Omega_p} s(y^{*(p)}(l-1), y^*(l-1))\Big\}$, $\forall p \in [1, w+1]$, where $y^{*(p)}(l-1)$ is the optimal solution of $FHC^{(p)}$ and $y^*(l-1)$ is the optimal solution of CMP at time $l-1$.*

Proof. Note that the problem (12) is NP-hard and the solution is obtained by Algorithm 2. Let $cost(FHC^{*(p)})$ be the optimal social cost given by $FHC^{(p)}$, we first prove $cost(FHC^{*(p)}) \leqslant cost(OPT) + \sum_{l\in\Omega_p} s(y^{*(p)}(l-1), y^*(l-1))$. Denote $\boldsymbol{x}^{*(p)}(t)$, $y^{*(p)}(t)$, $v^{*(p)}(t)$ as the optimal solution for the problem (12).

Note that $\sum_{t=l}^{l+w} C_{y^{*(p)}(t)}(\boldsymbol{x}^{*(p)}(t), v^{*(p)}(t)) + \sum_{t=l}^{l+w}\Big(\alpha y^{*(p)}(t) + s(y^{*(p)}(t-1), y^{*(p)}(t))\Big)$ is the local optimum objective value over $[l, l+w]$ and hence no larger than that of the following strategy: varying the number of heating plants $y^{(p)}(l-1)$ to $y^{(*)}(l-1)$ and following the offline optimum solution over $[l, l+w]$, we have

$$
\begin{aligned}
&cost(FHC^{*(p)}) \\
&= \sum_{l\in\Omega_p}\Bigg\{\sum_{t=l}^{l+w}\Big(C_{y^{*(p)}(t)}(\boldsymbol{x}^{*(p)}(t), v^{*(p)}(t)) + \alpha y^{*(p)}(t)\Big) + \sum_{t=l}^{l+w} s(y^{*(p)}(t-1), y^{*(p)}(t))\Bigg\} \\
&\leqslant \sum_{l\in\Omega_p}\Bigg\{\sum_{t=l}^{l+w} C_{y^*(t)}(\boldsymbol{x}^*(t), v^*(t)) + \alpha y^*(t) + \sum_{t=l+1}^{l+w} s(y^*(t-1), y^*(t)) \\
&+ s(y^{*(p)}(l-1), y^*(l-1)) + s(y^*(l-1), y^*(l))\Bigg\} \\
&= \sum_{t=1}^{T}\Big(C_{y^*(t)}(x^*(t), v^*(t)) + \alpha y^*(t) + s(y^*(t-1), y^*(t))\Big) + \sum_{l\in\Omega_p} s(y^{*(p)}(l-1), y^*(l-1)) \\
&= cost(OPT) + \sum_{l\in\Omega_p} s(y^{*(p)}(l-1), y^*(l-1)).
\end{aligned}
$$

Recall that $cost(FHC^{(p)})$ is calculated by Algorithm 2, according to the Theorem 2, we have for any $p \in [1, w+1]$,

$$cost(FHC^{(p)}) \leqslant 2cost(FHC^{*(p)}) \leqslant 2\Big\{cost(OPT) + \sum_{l \in \Omega_p} s(y^{*(p)}(l-1), y^*(l-1))\Big\}.$$

□

The proof of Theorem 3 is as follows:

Proof. Recall that we select the $FHC^{(p)}$ to calculate the approximate solution of over $[1, T]$ with equivalent probability. Furthermore, as mentioned in Theorem 6, the proposed reverse auction mechanism achieves 2-approximation ratio to the problem (6), which is the same as the ratio $\frac{cost(FHC^{(p)})}{cost(FHC^{*(p)})}$ to (6). Hence, the expected cost obtained by Algorithm 3 is:

$$\frac{1}{w+1}\sum_{p=1}^{w+1} cost(FHC^{(p)}) \leqslant \frac{2}{w+1}\sum_{p=1}^{w+1}\Big(cost(OPT) + \sum_{l \in \Omega_p} s(y^{*(p)}(l-1), y^*(l-1))\Big)$$
$$= 2cost(OPT) + \frac{2}{w+1}\sum_{p=1}^{w+1}\sum_{l \in \Omega_p} s(y^{*(p)}(l-1), y^*(l-1)).$$

Then the competitive ratio is

$$cr = \frac{\frac{1}{w+1}\sum_{p=1}^{w+1} cost(FHC^{(p)})}{cost(OPT)} \leqslant \frac{2cost(OPT) + \frac{2}{w+1}\sum\limits_{p=1}^{w+1}\sum\limits_{l \in \Omega_p} s(y^{*(p)}(l-1), y^*(l-1))}{cost(OPT)}$$
$$= 2 + \frac{2\sum_{p=1}^{w+1}\sum_{l \in \Omega_p} s(y^{*(p)}(l-1), y^*(l-1))}{(w+1)cost(OPT)} \leqslant 2 + \frac{2\beta\sum_{t=1}^{T} y^*(t)}{(w+1)cost(OPT)}$$
$$\leqslant 2 + \frac{2\beta\sum_{t=1}^{T} y^*(t)}{(w+1)\alpha\sum_{t=1}^{T} y^*(t)} \leqslant 2 + \frac{2\beta}{(w+1)\alpha}.$$

□

### A.3 Proof of Theorem 4

Proof. According to strong duality theorem, $\sum_{k \in \mathcal{K}} \lambda_k = 1$ implies that the optimal objective value of its dual problem (15) is 1. Note that the solution ($\boldsymbol{\mu} = \mathbf{0}$, $\sigma = 0$, $\tau = 1$) is feasible, the objective value of the dual problem is at most 1. Suppose instead that the optimal objective value $\sum_{i \in \mathcal{I}} \min\{\gamma x_i^*, 1\}\mu_i + \min\{\gamma v^*, Vy\}\sigma + \tau < 1$. What's more, if the dual variable $\mu < 0$, the approximation algorithm for (15) may be inappropriate. We constrain it as:

$$\mu_i^+ = \begin{cases} \mu_i & \text{if } \mu_i \geqslant 0, \text{ and} \gamma x_i^* \leqslant 1 \\ 0 & \text{otherwise} \end{cases}$$

To solve the dual problem, we regard each $\mu_i$ as the cost of bid $i$ and $\sigma$ as the heating cost. The mixed integer solution $(\boldsymbol{x}^k, v^k)$, $k \in \mathcal{K}$ calculated by the $\gamma$-approximation algorithm satisfies:

$$\sigma v^k + \sum_{i \in \mathcal{I}} x_i^k \mu_i^k \leqslant \gamma \tilde{v}^* \sigma + \gamma \sum_{i \in \mathcal{I}} \tilde{x}_i^* \mu_i^+ \leqslant \gamma \sigma v^* + \gamma \sum_{i \in \mathcal{I}} \mu_i x_i^*, \tag{20}$$

where $(\tilde{\boldsymbol{x}}^*, \tilde{v}^*)$ is optimal fractional solution with regarding each $\mu_i$ as the cost of bid $i$ and $\sigma$ as the heating cost respectively.

Note that the DHS can choose more bids to meet the heat demand meanwhile doesn't violate the covering constraint. Denote $\tilde{x_i}^k$ as:

$$\tilde{x_i}^k = \begin{cases} x_i^k & \text{if } \mu_i \geqslant 0, \text{and} \gamma x_i^* \leqslant 1 \\ 1 & \text{otherwise} \end{cases}$$

Then we have:

$$\sum_{i\in\mathcal{I}} \tilde{x_i}^k \mu_i = \sum_{i\in\mathcal{I}} x_i^k \mu_i^+ + \sum_{i:\mu_i<0 \text{ or } \gamma x_i^*>1} \mu_i \leqslant \sum_{i\in\mathcal{I}} x_i^k \mu_i^+ + \sum_{i:\mu_i<0 \text{ or } \gamma x_i^*>1} \min\{\gamma x_i^*, 1\}\mu_i. \tag{21}$$

Moreover, in this case, less heat can be produced to meet the heat demand, i.e., $\tilde{v}^k \leqslant v^k$. According to (20) and (21), we have

$$\begin{aligned}\sigma v^k + \sum_{i\in\mathcal{I}} x_i^k \mu_i^k &\leqslant \sigma \min\{\gamma v^*, Vy\} + \gamma \sum_{i\in\mathcal{I}} x_i^k \mu_i^+ + \sum_{i:\mu_i<0 \text{ or } \gamma x_i^*>1} \min\{\gamma x_i^*, 1\}\mu_i \\ &\leqslant \sigma \min\{\gamma v^*, Vy\} + \sum_{i\in\mathcal{I}} \min\{\gamma x_i^*, 1\}\mu_i \leqslant 1 - \tau.\end{aligned}$$

Note that $(\tilde{\boldsymbol{x}}^k, \tilde{v}^k)$ is a feasible solution to the problem (6) while the constraint $\sum_{i\in\mathcal{I}} x_i^k \mu_i + v^k \sigma + \tau \geqslant 1$ is violated, the assumption $\sum_{i\in\mathcal{I}} \min\{\gamma x_i^*, 1\}\mu_i + \min\{\gamma v^*, Vy\}\sigma + \tau < 1$ is contradictory. Hence we have $\sum_{k\in\mathcal{K}} \lambda_k = 1$.

To solve the dual problem (15) within polynomial time, we utilize the ellipsoid method with regarding Algorithm 1 as the separation oracle to get a polynomial number of separating hyperplanes for the dual problem. After obtaining the optimal solution of (15), the primal problem (14) can be reduced to a linear program with polynomial number of variables corresponding to its dual solutions and solved by efficient linear programming technique such as simplex method. Above all, the convex decomposition can be solve with polynomial time. □

### A.4 Proof of Theorem 6

Proof. *i*) The proposed mechanism achieves individual rationality. Let $U_i$ denote the utility of datacenter $i$ with truthful cost $c_i$. If datacenter $i$ wins in this round, recall that $P_i(c) = \min\{\gamma x_i^*(c, c_{-i}), 1\}$, we have $U_i = r_i - c_i = \frac{\int_{c_i}^{\zeta} P_i(c)dc}{P_i(c_i)} \geqslant 0$. Otherwise, if datacenter $i$ fails, its utility $U_i = 0$.

Above all, all of datacenters can not lose money in the auction, and the proposed mechanism is individually rational.

*ii*) Recall that the Algorithm 1 is used for convex decomposition and achieves 2-approximation ratio in solving (6), *i.e.*, $\gamma = 2$, the expectation objective value of (6) is

$$\begin{aligned}E\Big[\sum_{i=1}^{M} c_i x_i + v p_h\Big] &= \sum_{k\in\mathcal{K}} \lambda_k \Big\{\sum_{i=1}^{M} c_i x_i^k + p_h v^k\Big\} \leqslant \Big\{\sum_{i=1}^{M} c_i \min\{\gamma x_i^*, 1\} + p_h \min\{\gamma v^*, Vy\}\Big\} \\ &\leqslant 2OPT_{LPR} \leqslant 2OPT.\end{aligned}$$

Furthermore, according to Theorem 2, the reverse auction mechanism achieves 2-approximation ratio in expectation.

*iii*) The proposed mechanism achieves truthfulness in expectation. We first prove the $P_i(c_i) = \min\{\gamma x_i^*, 1\}$ is monotonically nonincreasing in $c_i$. Let $C(\boldsymbol{x}, c_i, c_{-i})$ denote the objective value of the problem (6) with bids $(c_i, c_{-i})$ and optimal fractional solution $\boldsymbol{x}$. Fix the $c_{-i}$, and assume $x_i^*$ and $x_i'^*$ are the optimal solution to $i$ with bids $c_i$ and $c_i'$ respectively, where $c_i' \leqslant c_i$. Then we have

$$C(\boldsymbol{x}^*, c_i, c_{-i}) \leqslant C(\boldsymbol{x}'^*, c_i, c_{-i})$$

$$C(\boldsymbol{x}'^*, c_i', c_{-i}) \leqslant C(\boldsymbol{x}^*, c_i', c_{-i})$$

Sum the above two inequations and reformulate it, we have

$$\begin{aligned} C(\boldsymbol{x}^*, c_i, c_{-i}) - C(\boldsymbol{x}^*, c_i', c_{-i}) &\leqslant C(\boldsymbol{x}'^*, c_i', c_{-i}) - C(\boldsymbol{x}'^*, c_i, c_{-i}) \\ \Leftrightarrow (c_i - c_i')x_i^* &\leqslant (c_i - c_i')x_i'^* \\ \Leftrightarrow x_i^* &\leqslant x_i'^* \end{aligned}$$

Therefore, the $P_i(c_i)$ is monotonically nonincreasing in $c_i$.

Consider an extreme case: if the cost of datacenter $i$'s bid is larger than the cost of self-producing $u_i$ plus the cost of turning on enough heating plants for producing it, the datacenter $i$ must fail in the auction. In other words, if $c_i > \zeta = p_h u_i + \lceil \frac{u_i}{V} \rceil(\alpha + \beta)$, then $x_i^* = 0$, and $P_i(c_i) = 0$. Therefore,

$$\int_0^\infty P_i(c)dc = \int_0^{p_h u_i + \lceil \frac{u_i}{V} \rceil(\alpha+\beta)} P_i(c)dc \leqslant p_h u_i + \lceil \frac{u_i}{V} \rceil(\alpha + \beta) < \infty.$$

Further, we have

$$E[r_i] = P_i(c_i)\left(c_i + \frac{\int_{c_i}^{\zeta} P_i(c)dc}{P_i(c_i)}\right) = c_i P_i(c_i) + \int_{c_i}^\infty P_i(c)dc.$$

Above all, the proposed mechanism is truthful in expectation. □

## REFERENCES


[1] Zahra Abbasi, Georgios Varsamopoulos, and Sandeep K. S. Gupta. 2012. TACOMA: Server and Workload Management in Internet Data Centers Considering Cooling-computing Power Trade-off and Energy Proportionality. *ACM Trans. Archit. Code Optim.* 9, 2, Article 11 (June 2012), 37 pages. DOI: http://dx.doi.org/10.1145/2207222.2207227

[2] Aaron Archer and Éva Tardos. 2001. Truthful mechanisms for one-parameter agents. In *Proc. of IEEE FOCS*.

[3] Tim Ashton. 2015. Heat Pumps for Data Centers. http://www.oppenfjarrvarme.se/media/Heat-Recovery-from-Data-Centers-Ashton-20150121.pdf. (2015).

[4] Luiz André Barroso, Jimmy Clidaras, and Urs Hölzle. 2013. *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines, Second Edition.* Morgan & Claypool Publishers.

[5] Josep Ll Berral, Íñigo Goiri, Thu D Nguyen, Ricard Gavalda, Juana Torres, and Ricardo Bianchini. 2014. Building green cloud services at low cost. In *Proc. of IEEE ICDCS*.

[6] Nick Booth. 2015. Fish farmer says data center will kill his fish. http://www.datacenterdynamics.com/design-strategy/fish-farmer-says-data-center-will-kill-his-fish/93448.article. (2015).

[7] Tim Carnes and David Shmoys. 2008. Primal-dual schema for capacitated covering problems. In *Integer Programming and Combinatorial Optimization.* Springer, 288–302.

[8] Robert D Carr, Lisa Fleischer, Vitus J Leung, and Cynthia A Phillips. 2000. Strengthening integrality gaps for capacitated network design and covering problems.. In *Prof. of ACM-SIAM SODA*.

[9] Edward H Clarke. 1971. Multipart pricing of public goods. *Public choice* 11, 1 (1971), 17–33.

[10] Daikin. 2017. The COP of heat recovery unit. http://www.daikinac.com/content/assets/DOC/Product%20Brochures/PCVUSE13-05C-VRVIII-Brochure-Daikin-AC.PDF. (2017).

[11] Data Center Map. 2017. http://www.datacentermap.com. (2017).

[12] DECC. 2013. Research on the costs and performance of heating and cooling technologies. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/204275/Research_on_the_costs_and_performance_of_heating_and_cooling_technologies__Sweett_Group_.pdf. (2013).

[13] Wei Deng, Fangming Liu, Hai Jin, Bo Li, and Dan Li. 2014. Harnessing renewable energy in cloud datacenters: opportunities and challenges. *IEEE Network* 28, 1 (2014), 48–55.

[14] Wei Deng, Fangming Liu, Hai Jin, and Xiaofei Liao. 2013a. Online control of datacenter power supply under uncertain demand and renewable energy. In *Proc. of IEEE ICC*.

[15] Wei Deng, Fangming Liu, Hai Jin, and Chuan Wu. 2013b. SmartDPSS: Cost-minimizing multi-source power supply for datacenters with arbitrary demand. In *Proc. of IEEE ICDCS*.

[16] Wei Deng, Fangming Liu, Hai Jin, Chuan Wu, and Xue Liu. 2013. Multigreen: Cost-minimizing multi-source datacenter power supply with online control. In *Proc. of ACM e-Energy*.

[17] District Heating. 2017. https://en.wikipedia.org/wiki/District_heating. (2017).

[18] Khosrow Ebrahimi, Gerard F Jones, and Amy S Fleischer. 2014. A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities. *Renewable and Sustainable Energy Reviews* 31 (2014), 622–638.
[19] EIA. 2014. Industrial Sector Energy Price Estimates, 2013. http://www.eia.gov/state/seds/data.cfm?incfile=/state/seds/sep_sum/html/sum_pr_ind.html&sid=US. (2014).
[20] EnergyPlus. 2017. Weather data. https://energyplus.net/weather. (2017).
[21] Enwave. 2017. Heating Toronto's core for over 40 years. http://enwavetoronto.com/district_heating_system.html. (2017).
[22] Fortum. 2015. Open District Heating for Data centers. https://www.oppenfjarrvarme.se/media/Heat-Recovery-from-Data-Centers-Rylander-20150121.pdf. (2015).
[23] Fortum. 2017. Profitable recovery with Open District Heating. http://www.oppenfjarrvarme.se/media/open-district-heating-bahnhof-thule.pdf.pdf. (2017).
[24] Wendelin Friedel and Werner Neumann. 2002. Systematic implementation of CHP plants in Frankfurt am Main. https://www.frankfurt.de/sixcms/media.php/738/ieecb_chp.pdf. (2002).
[25] Henrik Gadd and Sven Werner. 2013. Daily heat load variations in Swedish district heating systems. *Applied Energy* 106 (2013), 47–55.
[26] Anshul Gandhi, Yuan Chen, Daniel Gmach, Martin Arlitt, and Manish Marwah. 2011. Minimizing data center SLA violations and power consumption via hybrid resource provisioning. In *Proc. of IEEE IGCC*.
[27] Íñigo Goiri, Thu D Nguyen, and Ricardo Bianchini. 2015. CoolAir: Temperature-and Variation-Aware Management for Free-Cooled Datacenters. In *Proc. of ACM ASPLOS*.
[28] Theodore Groves. 1973. Incentives in teams. *Econometrica: Journal of the Econometric Society* (1973), 617–631.
[29] Lena Gunnarsson. 2012. Fortum — Open District Heating. http://rhc-platform.org/fileadmin/2012_RHC_Conference/Day_1_-_26.04/P_120426_1530_SessionE_Lena_Gunnarsson.pdf. (2012).
[30] Heating Degree Day. 2017. https://en.wikipedia.org/wiki/Heating_degree_day. (2017).
[31] Mohammad A. Islam, A. Hasan Mahmud, Shaolei Ren, and Xiaorui Wang. 2015. Paying to save: Reducing cost of colocation data center via rewards. In *Proc. of IEEE HPCA*.
[32] Mohammad A Islam, Xiaoqi Ren, Shaolei Ren, Adam Wierman, and Xiaorui Wang. 2016. A Market Approach for Handling Power Emergencies in Multi-Tenant Data Center. In *Proc. of IEEE HPCA*.
[33] Soochan Lee, Patrick E Phelan, and Carole-Jean Wu. 2015. Hot Spot Cooling and Harvesting Central Processing Unit Waste Heat Using Thermoelectric Modules. *Journal of Electronic Packaging* 137, 3 (2015), 031010.
[34] Li Li, Wenli Zheng, Xiaodong Wang, and Xiaorui Wang. 2014. Placement optimization of liquid-cooled servers for power minimization in data centers. In *Proc. of IEEE IGCC*.
[35] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. 2013. Dynamic Right-Sizing for Power-Proportional Data Centers. *IEEE/ACM Transactions on Networking* 21, 5 (Oct 2013), 1378–1391. DOI: http://dx.doi.org/10.1109/TNET.2012.2226216
[36] Jie Liu, Michel Goraczko, Sean James, Christian Belady, Jiakang Lu, and Kamin Whitehouse. 2011. The data furnace: heating up with cloud computing. In *Proc. of USENIX HotCloud*.
[37] Longjun Liu, Chao Li, Hongbin Sun, Yang Hu, Juncheng Gu, Tao Li, Jingmin Xin, and Nanning Zheng. 2015. HEB: deploying and managing hybrid energy buffers for improving datacenter efficiency and economy. In *Proc. of ISCA*.
[38] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. 2012. Renewable and cooling aware workload management for sustainable data centers. In *Proc. of ACM SIGMETRICS*.
[39] L. Lu, J. Tu, C. Chau, M. Chen, and X. Lin. 2013. Online Energy Generation Scheduling for Microgrids with Intermittent Energy Sources and Co-Generation. In *Proc. of ACM SIGMETRICS*.
[40] Lauren Macready. 2015. Center for Biological Sequence Analysis at the Technical University of Denmark Installs Liquid-Cooled Solution from CoolIT Systems. http://www.prweb.com/releases/2015/06/prweb12795601.htm. (2015).
[41] Iyswarya Narayanan, Di Wang, Abdullah-Al Mamun, Anand Sivasubramaniam, and Hosam K Fathy. 2014. Should we dual-purpose energy storage in datacenters for power backup and demand response?. In *Prof. of USENIX HotPower*.
[42] Haga Niklas, Kortela Veikko, and Ahnger Anders. 2012. White paper: Smart Power Generation — District heating solutions. (2012).
[43] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Vol. 1. Cambridge University Press Cambridge.
[44] NYC. 2017. Important Safety Issues — Heat and Hot Water. https://www1.nyc.gov/site/hpd/renters/important-safety-issues-heat-hot-water.page. (2017).
[45] Open District Heating. 2017. http://www.opendistrictheating.com/. (2017).
[46] NYC Global Partners. 2011. Best Practice: District Heating System. (2011).
[47] PJM. 2017a. Hourly Real-Time & Day-Ahead LMP. http://www.pjm.com/markets-and-operations/energy/real-time/monthlylmp.aspx. (2017).

[48] PJM. 2017b. Locational Marginal Pricing. http://www.pjm.com/~/media/about-pjm/newsroom/fact-sheets/locational-marginal-pricing-fact-sheet.ashx. (2017).

[49] Neil Rasmussen. 2011. Calculating Total Cooling Requirements for Data Centers. http://apcmedia.com/salestools/NRAN-5TE6HE/NRAN-5TE6HE_R3_EN.pdf. (2011).

[50] Shaolei Ren and Mohammad A Islam. 2014. Colocation demand response: Why do I turn off my servers. *Proc. of USENIX ICAC* (2014).

[51] McMillan Robert. 2013. Servers Too Hot? Intel Recommends a Luxurious Oil Bath. https://www.wired.com/2012/08/servers-too-hot-intel-recommends-a-luxurious-oil-bath/. (2013).

[52] Alexander Schrijver. 1998. *Theory of linear and integer programming*. John Wiley & Sons.

[53] Ren Shaolei. 2017. Power Management in Multi-Tenant Colocation Data Centers. (2017). http://www.ece.ucr.edu/~sren/doc/tech/colo_summary.pdf

[54] Shinryo Corporation. 2017. District heating system in Tokyo. https://www.shinryo.com/en/tech/bf_dhc.html. (2017).

[55] TeleGeography. 2014. Telegeography Colocation Database. http://www.telegeography.com/research-services/colocation-database/. (2014).

[56] Uptime Institute. 2014. 2014 Data Center Industry Survey. https://journal.uptimeinstitute.com/2014-data-center-industry-survey/. (2014).

[57] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.

[58] E.M. Ward, M.L. Goedke, P. Brenner, and D.B. Go. 2012. A simplified thermodynamic model for waste heat utilization from a containerized data center experimental platform. In *Proc. of IEEE ITherm*. DOI:http://dx.doi.org/10.1109/ITHERM.2012.6231475

[59] Hong Xu, Chen Feng, and Baochun Li. 2015. Temperature Aware Workload Managementin Geo-Distributed Data Centers. *Parallel and Distributed Systems, IEEE Transactions on* 26, 6 (2015), 1743–1753.

[60] Linquan Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2015. A truthful incentive mechanism for emergency demand response in colocation data centers. In *Proc. of IEEE INFOCOM*. DOI:http://dx.doi.org/10.1109/INFOCOM.2015.7218654

[61] Linquan Zhang, Shaolei Ren, Chuan Wu, and Zongpeng Li. 2016. A Truthful Incentive Mechanism for Emergency Demand Response in Geo-Distributed Colocation Data Centers. *ACM Trans. Model. Perform. Eval. Comput. Syst.* 1, 4, Article 18 (Sept. 2016), 23 pages. DOI:http://dx.doi.org/10.1145/2950046

[62] Linquan Zhang, Chuan Wu, Zongpeng Li, Chuanxiong Guo, Minghua Chen, and Francis Lau. 2013. Moving big data to the cloud: an online cost-minimizing approach. *Selected Areas in Comm., IEEE Journal on* 31, 12 (2013), 2710–2721.

[63] Ruiting Zhou, Zongpeng Li, and Chuan Wu. 2015. An online procurement auction for power demand response in storage-assisted smart grids. In *Proc. of IEEE INFOCOM*.

[64] Zhi Zhou, Fangming Liu, Bo Li, Baochun Li, Hai Jin, Ruolan Zou, and Zhiyong Liu. 2014. Fuel cell generation in geo-distributed cloud services: A quantitative study. In *Proc. of IEEE ICDCS*.

[65] Zhi Zhou, Fangming Liu, and Zongpeng Li. 2016. Bilateral Electricity Trade Between Smart Grids and Green Datacenters: Pricing Models and Performance Evaluation. *IEEE Journal on Selected Areas in Communications* 34, 12 (2016), 3993–4007.

[66] Z. Zhou, F. Liu, Z. Li, and H. Jin. 2015. When Smart Grid Meets Geo-distributed Cloud: An Auction Approach to Datacenter Demand Response. In *Proc. of IEEE INFOCOM*.


Received April 2017; revised November 2017; accepted March 2018