

Physical and Economic Viability of Cryptocurrency Mining for Provision of Frequency Regulation: A Real-World Texas Case Study

Rayan El Helou*, Ali Menati* and Le Xie^{†, §}
 {*Student Member, [†]Fellow}, IEEE

Abstract—Demand flexibility plays a pivotal role in modern power systems with high penetration of variable energy resources. In recent years, one of the fastest-growing flexible energy demands has been proof-of-work-based cryptocurrency mining facilities. Due to their competitive ramping capabilities and demonstrated flexibility, such fast-responding loads are capable of participating in frequency regulation services for the grid while simultaneously increasing their own operational revenue. In this paper, we investigate the physical and economic viability of employing cryptocurrency mining facilities to provide frequency regulation in large power systems. We quantify mining facilities' operational profit, and propose a decision-making framework to explore their optimal participation strategy and account for the most influential factors. We employ real-world ERCOT ancillary services data in our case study to investigate the conditions under which provision of frequency regulation in the Texas grid is profitable. We also perform transient level simulations using a synthetic Texas grid to demonstrate the competitiveness of mining facilities at frequency regulation provision.

Index Terms—frequency regulation, cryptocurrency mining, ancillary services, electricity market.

I. INTRODUCTION

ONE of the primary roles of an electric transmission grid operator is to simultaneously balance supply and demand and ensure reliable grid operations. This has become a challenging task in recent years with the proliferation of distributed energy resources, such as renewables and electric vehicles [1], which introduce increasing levels of uncertainty to grid operations. Inevitably, ancillary services are bound to become a cornerstone to the resilience of modern power systems, provided that they can offer fast and affordable flexibility. Among the most pivotal ancillary services is frequency regulation, which is a fast time-scale control service that accounts for electric supply-demand imbalance not covered by the slower time-scale energy market [2]. Large frequency deviations may harm grid infrastructure, trip under-frequency relays, and lead to power system failure [3]–[5].

Traditionally, frequency regulation is done through generation facilities capable of quickly responding either to local frequency measurements or to signals communicated by the grid operator to increase and decrease their power output. In recent years, high elasticity of flexible loads has been utilized to help balance supply and demand and provide

frequency regulation in the grid. Among such flexible loads is the emerging cryptocurrency mining facilities which create a unique opportunity to offer frequency regulation and other ancillary services [6]. In comparison, conventional methods for provision of frequency regulation, such as thermal fuel and hydro generators, are expensive and rather slow [7], but cryptocurrency mining facilities can demonstrably provide significantly faster response [8].

An alternative to cryptocurrency mining facilities is Internet data centers, which are computing-based loads that have also been used to participate in ancillary services. They are faster than conventional generators and are usually paired with battery storage systems in order to provide frequency regulation [9]–[11]. The flexibility of internet data centers is harnessed by managing their delay-tolerant workload through proper hardware design and algorithmic solutions to adapt their processing power. There has been extensive research on data center load management techniques such as speed-scaling [12], power-capping [13], capacity right sizing [14], and geographical load balancing [15], [16]. However, such data centers are limited in their flexibility compared to cryptocurrency mining facilities, due to the former's time-sensitive computational obligations [17]. In practice, since cryptominers are not typically considered critical load, they are fully flexible and could potentially shut down all of their operational capacity within a few seconds [18]. For example, during the winter storm Elliot in December 2022, cryptocurrency mining facilities in Texas provided 1,475 MW of load reduction, which was more than 96% of their total capacity [19].

TABLE I
SUMMARY OF 2022 ANCILLARY SERVICES MARKET AT ERCOT

	Hourly Average	Reg-Up	Reg-Down	RRS	NSRS
Price for Capacity (\$/MW)	21.67	8.46	20.31	22.49	
Capacity Procured (MW)	359	348	2863	3895	
Deployment Rate	16%	25%	≈ 0%	≈ 0%	

In the Electric Reliability Council of Texas (ERCOT), cryptocurrency mining loads are capable of performing all ancillary services and participating in real-time energy market [20]. Some ancillary services like the Responsive Reserve Service (RRS) and the Non-spinning Reserve Service (NSRS) are designed for emergency reserve, and they would be deployed in instances of large generation shortages [21]. Other ancillary

Rayan El Helou, Ali Menati and Le Xie are with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX, USA. Corresponding Author: Le Xie, email: le.xie@tamu.edu

services like Reg-Up/Down are used for frequency regulation, and as shown in Table I, the size of Reg-Up/Down is much smaller than the bulk size of NSRS or RRS. In terms of deployment, it can also be seen that RRS and NSRS are rarely deployed, while Reg-Up/Down is deployed more often. The high flexibility and fast response time of cryptominers qualify them as a suitable candidate for frequency regulation provision, unlike other loads.

This paper investigates cryptominers' participation in frequency regulation from both economic and operational perspectives, and our main contributions are summarized as follows:

- 1) We present a decision-making framework for assessing the economic viability of cryptominers' participation in frequency regulation services considering the most influential factors including market clearing prices and deployment rates of these services, cryptocurrency prices, mining energy, and electricity prices. We utilize a historical record of real-world data to demonstrate under what circumstances participation is profitable in the Texas grid.
- 2) We formulate the real-time frequency regulation dispatch problem from a grid operator's perspective subject to capacity and ramp constraints of resource entities. To solve the problem, we propose alternative real-time regulation dispatch mechanisms that address issues of fairness in allocation and efficiency of resource utilization when dispatching to multiple mining facilities at once.
- 3) We perform transient-level simulations using a 2000-bus synthetic Texas grid to demonstrate the physical viability of cryptominers participation in frequency regulation services, with our results revealing substantial improvement in resiliency compared to existing alternatives.

The remainder of the paper is organized as follows. In Section II, the frequency regulation market is formulated, where we present the parties involved, the day-ahead market, and real-time dispatch. In Section III, profit by mining facilities is formulated, where we present revenue due to mining, profit due to frequency regulation, and a set of proposed strategies on when to participate in the frequency regulation market. In Section IV, real-world data from ERCOT and historical records of Bitcoin prices are utilized to quantify and illustrate the feasibility and viability of participation by Bitcoin mining facilities. Finally, Section V provides concluding remarks.

II. REGULATION MARKET FORMULATION

In this section, we formulate the frequency regulation market, and we use ERCOT as an example for a possible set of rules and requirements adopted in the real-world. In particular, we introduce all the parties involved in this market and what role each plays, we formulate the day-ahead market (DAM) and the clearing process which yields hourly prices and capacities, and we formulate the real-time dispatch of all regulation service providers.

A. Parties Involved

There are four main parties involved in ancillary services markets: the market operator, the buyers, the sellers and the resource entities which provide the actual services, managed

by the sellers. Fig. 1 illustrates the interaction between these parties both in real-time and day-ahead.

Load Serving Entities (LSEs) provide electric service to customers and they serve as the buyers in this market, making them solely responsible for covering the costs of ancillary services. Even though LSEs pay for it, the regulation power demand is not explicitly determined by them, but by the market operator whose role is to ensure grid resiliency, which ultimately serves customers of the LSEs. Reference [2] details methodologies for determining required hourly ancillary services demand in ERCOT. Each electric power transmission system in the United States is operated by a regional transmission organization subject to federal and local regulations. Within its region, such an organization serves as the market operator. In this paper, we focus on the case of Texas, where ERCOT is the market operator for the Texas interconnection. Finally, sellers, who manage the physical mining facilities considered in this paper, are responsible for submitting capacity offers in the day-ahead market, while ensuring that any cleared capacity is deliverable in real-time. In this paper, we use the ERCOT terminology Qualified Scheduling Entities (QSEs) to refer to such sellers who participate in the market on behalf of the mining facilities, as illustrated in Fig. 1.

In the day-ahead market, the market operator clears prices and capacities for ancillary services, and in real-time it observes frequency changes and dispatches suppliers of regulation services to increment or decrement power consumption or production in attempt to regulate frequency. Since demand is perfectly inelastic in ancillary service markets, the market operator needs to first determine exactly how much demand is needed for each period of time. In the case of ERCOT, demand for capacity at every hour of the year is determined prior to the beginning of each year and is fixed for one whole year [2]. In Section II-B, we overview how prices and capacities for each resource entity are determined in the market clearing process, and in Section II-C, we overview how resources are dispatched to regulate frequency, and we propose alternative dispatch mechanisms that address fairness and resource utilization.

B. Day-Ahead Market

In this subsection, we detail how prices and quantities for each seller are determined during the market clearing process. The term *day-ahead* comes from the fact that regulation services are purchased as capacity in advance so that they can be deployed in the next day when needed. Here are some key features of the day-ahead market, based on ERCOT's protocols:

- All sellers in the market are paid at the same hourly market clearing price (in \$/MW), i.e. no locational-prices.
- Hourly demand for capacity (in MW) is inelastic, determined year-ahead by forecasting renewables' uncertainty.

Frequency regulation is offered by multiple forms of ancillary services. In this paper, we focus only on two kinds of services, named in ERCOT as Reg-Up and Reg-Down. The former is deployed in real-time when the frequency is lower than desired and more generation (or less consumption) of power is needed, and vice versa for the latter. For each of those two, at the beginning of each day the market clearing process produces

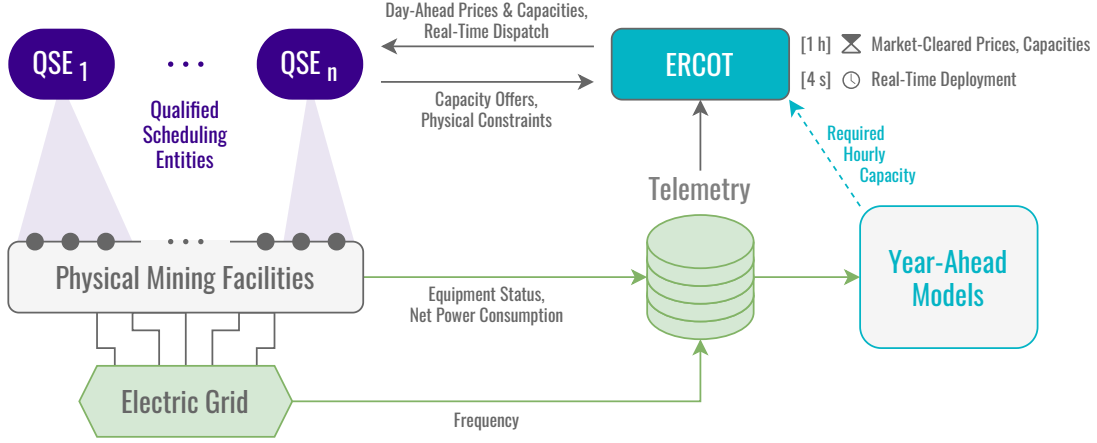


Fig. 1. Physical and economic interaction between ERCOT and the mining facilities participating in frequency regulation. Physical facilities interact with ERCOT via Qualified Scheduling Entities. In the Day-Ahead Market, prices and capacities are cleared hourly, while in Real-Time deployment, dispatch is updated every four seconds. Telemetry is used both for real-time control and for improving year-ahead models that determine future need for capacity.

24 prices (in \$/MW) for Reg-Up and 24 for Reg-Down, one for each hour of the upcoming day. Furthermore, each resource entity which submits an offer (as a seller) for each hour is told for that hour what its individual cleared capacity (in MW) is for both Reg-Up and Reg-Down.

The market clearing process is based on the principle that expected demand for capacity needs to be available at all times. In the year-ahead, as described in reference [2], ERCOT uses historic variations in loads and renewables to forecast frequency fluctuations in the upcoming year, which then translates to expected demand for capacity for each hour of the upcoming year, \widehat{D}^{up} and \widehat{D}^{dn} , corresponding to Reg-Up and Reg-Down respectively. Using this inelastic demand, the market operator clears the market by considering the capacities and prices offered by all n QSEs. Out of the n offers submitted, $m \leq n$ are cleared for non-zero capacity. Let $c_i^{\text{up}}, c_i^{\text{dn}} \forall i \in [1, m]$ be the cleared Reg-Up and Reg-Down capacities respectively. The market operator is expected to ensure that $\sum_{i=1}^m c_i^{\text{up}} = \widehat{D}^{\text{up}}$ and $\sum_{i=1}^m c_i^{\text{dn}} = \widehat{D}^{\text{dn}}$.

C. Real-Time Dispatch

In this subsection, we detail how Reg-Up and Reg-Down power is dispatched for each resource entity during real-time operations, and we investigate alternative approaches for allocation across multiple entities. The term *real-time* comes from the fact that regulation services are communicated and deployed in quick response to changes in frequency. Here are some key features of real-time operations pertaining to this study, based on ERCOT's protocols:

- Regulation dispatch is communicated every 4 seconds. However, some entities might need more than 4 seconds to reach desired value due to ramp constraints.
- The outcome of real-time deployment does not affect financial rewards collected by entities. Sellers are **only rewarded for capacity** in the day-ahead market.

1) Frequency Regulation Problem Formulation

Based on the outcome of the day-ahead clearing process, the market operator knows for each time step in real-time

operations exactly how much capacity is available for deployment by each resource entity. We present a mathematical formulation below for determining how much each entity is dispatched in real-time for any given observed frequency.

Since each hour contains 900 four-second consecutive periods, let $T = 900$ be the length of a discrete-time sequential process, which we'll refer to here as an episode, where each step $t \in [0, T-1]$ corresponds to a step in real-time operations. Since the time step is of length 4 seconds and there are 900 steps, the episode is of length 1 hour. Let $t = 0$ represent the time where each new hour from the day-ahead market begins. That is, for each real-time episode, the market operator may assume a fixed set of available Reg-Up/Down capacities for each resource entity, and at the end of time step $T - 1$, the episode concludes with a new set of capacities in the episode which follows. Without loss of generality, we focus on only one episode moving forward.

For any episode, let m denote the number of resource entities available for dispatch, i.e. with non-zero cleared capacity from the day-ahead market, corresponding to the same one-hour period, as described in Section II-B. Let positive scalars $c_1^{\text{up}}, c_1^{\text{dn}}, c_2^{\text{up}}, c_2^{\text{dn}}, \dots, c_m^{\text{up}}, c_m^{\text{dn}}$ denote the cleared Reg-Up/Down capacities for the m resource entities, and let $d_i[t]$ denote the regulation power dispatched by the market operator at time step $t \in [0, T-1]$ to resource entity $i \in [1, m]$. Here, $d_i[t]$ represents both Reg-Up and Reg-Down simultaneously; the value is positive during Reg-Down and negative during Reg-Up. This convention is chosen since more power consumption is required when frequency is larger than nominal, i.e. when Reg-Down is needed to pull it down. One reasonable assumption to make is that there will always be some exogenous fluctuations in the system at each time step that contribute to changes in the frequency, on top of changes in regulation dispatch. Denote all exogenous perturbations from time step t to $t + 1$ as $\eta[t]$. Note that this is not a scalar, but an arbitrary object that notationally represents a collection of exogenous forces over an entire 4 second window (i.e. one time step t). We first pose the frequency regulation task for each episode

as an ideal optimization problem (Eq. (1)), then we present a set of propositions to simplify the task and yield a practical solution for it.

By convention, boldface refers to vectors. For example $\mathbf{c}^{\text{up}} \in \mathbb{R}^m$ is the vector of Reg-Up capacities cleared in the day-ahead market for a given episode.

$$\text{minimize}_{\{\mathbf{x}[t]\}_{t=0}^{T-1}} \sum_{t=1}^T \max(\underline{f} - f[t], 0, f[t] - \bar{f}) \quad (1a)$$

$$\text{s.t. } \forall t \in [0, T-1] :$$

$$-\mathbf{c}^{\text{up}} \leq \mathbf{d}[t+1] \leq \mathbf{c}^{\text{dn}} \quad (1b)$$

$$\mathbf{r} \leq \mathbf{x}[t] \leq \bar{\mathbf{r}} \quad (1c)$$

$$\mathbf{d}[t+1] = \mathbf{d}[t] + \mathbf{x}[t] \quad (1d)$$

$$f[t+1] = \text{PhysicalGrid}(f[t], \mathbf{x}[t], \eta[t]) \quad (1e)$$

where \mathbf{x} denotes incremental changes in Reg-Up/Down dispatch and $[\underline{\mathbf{r}}, \bar{\mathbf{r}}]$ denotes ramp constraints due to physical limitations of resource entities. $f[t] \in \mathbb{R}$ denotes real-time system-wide frequency and $[\underline{f}, \bar{f}]$ is the window of acceptable frequencies. That is, the objective function in Eq. (1a) evaluates to zero if $f[t] \in [\underline{f}, \bar{f}]$ and increases linearly as $f[t]$ escapes the window in either direction.

$\mathbf{x}[t] \in \mathbb{R}^m$ is the decision variable for each step t in the episode, and $f[0], \mathbf{d}[0]$ are exogenous to the optimization problem, yet known. Obtaining an exact solution to Eq. (1) is very difficult due to the constraint in Eq. (1e), which is strictly speaking governed by high-order differential equations, inertia and non-linearities in the system. We would require a rough estimate of the impact of regulation dispatch on the frequency to pragmatically solve this problem. In practice, a heuristic two-stage approach is adopted to addressing this issue. In the first stage, the problem is simplified by decoupling the time steps within each episode so that $\mathbf{x}[t]$ is found separately for each step t . It is assumed that all that is needed to find a suitable dispatch is knowledge of system-wide frequency $f[t]$ and operating point $\mathbf{d}[t]$. From this, a required change in dispatch, $\Delta[t]$, is computed proportional to variations in frequency as follows, for some $\beta > 0$:

$$\Delta[t] = \beta [\max(0, f[t] - \bar{f}) + \min(0, f[t] - \underline{f})] \quad (2)$$

Note that if the grid operator desires $\bar{f} = f = 60$ Hz, then the summand in Eq. (1) reduces to $|f[t] - 60 \text{ Hz}|$ and the right hand side of Eq. (2) reduces to $\beta f[t]$.

Finally, the required net dispatch is distributed across $\mathbf{x}[t]$ while respecting both ramp constraints (Eq. (1c)) and market-cleared capacities for Reg-Up/Down (Eq. (1b)). This restricts \mathbf{x} to the following bounds, at each operating point $\mathbf{d}[t]$:

$$\underline{\mathbf{x}}[t] \leftarrow \max(\mathbf{r}, -\mathbf{c}^{\text{up}} - \mathbf{d}[t]) \quad (3a)$$

$$\bar{\mathbf{x}}[t] \leftarrow \min(\bar{\mathbf{r}}, \mathbf{c}^{\text{dn}} - \mathbf{d}[t]) \quad (3b)$$

At each point in time, the market operator computes required change in dispatch $\Delta[t]$ from Eq. (2) along with bounds on incremental dispatch $[\underline{\mathbf{x}}[t], \bar{\mathbf{x}}[t]]$ from Eq. (3). In the second stage, those values for each time step t are used to solve a simplified problem, posed in Eq. (4) below, to yield $\mathbf{x}[t]$.

$$\text{minimize}_{\mathbf{x}[t]} |\Delta[t] - \mathbf{1} \cdot \mathbf{x}[t]| \quad (4a)$$

$$\text{s.t. } \underline{\mathbf{x}}[t] \leq \mathbf{x}[t] \leq \bar{\mathbf{x}}[t] \quad (4b)$$

where $\mathbf{1}$ denotes a vector with m ones and $\mathbf{1} \cdot \mathbf{y}$ is the sum of entries in some $\mathbf{y} \in \mathbb{R}^m$. The solution to Eq. (4) is proposed in Section II-C2 below.

2) Proposed Regulation Dispatch

We propose a pair of possible solutions to Eq. (4) in Propositions 2 and 3 below, both of which adopt the principle assumed in Proposition 1 below.

Proposition 1 (Resource conservation). *At each point in time, if there's a need for regulation, it is either for increasing or decreasing regulation power, but never for both. Thus, to conserve resources, ensure $\forall t \in [0, T-1]$ that no pair of entries in vector $\mathbf{x}[t]$ are of opposite sign.*

For any solution, once action $\mathbf{x}[t]$ is determined, the resulting dispatch follows Eq. (1d). This is effectively integral droop control, wherein regulation dispatch $\mathbf{d}[t]$ continues to increment/decrement while the frequency is not within the desired range. The only time such form of control is incapable of sufficiently regulating the frequency is when there is simply a shortage of regulation resources. In such cases, the market operator resorts to other forms of ancillary services which are otherwise rarely deployed, such as RRS or NSRS in ERCOT.

There are either one or infinitely many solutions to Eq. (4). If $\Delta[t] \geq \mathbf{1} \cdot \bar{\mathbf{x}}[t]$, the optimal solution is $\mathbf{x}[t] \leftarrow \bar{\mathbf{x}}[t]$. If $\Delta[t] \leq \mathbf{1} \cdot \underline{\mathbf{x}}[t]$, the optimal solution is $\mathbf{x}[t] \leftarrow \underline{\mathbf{x}}[t]$. Infinitely many solutions exist in all other cases, i.e. when $\mathbf{1} \cdot \underline{\mathbf{x}}[t] \leq \Delta[t] \leq \mathbf{1} \cdot \bar{\mathbf{x}}[t]$. We use the term *sufficient resources* in both Proposition 2 and 3 below to indicate such cases.

Proposition 2 (Equitable dispatch). *Given sufficient resources, i.e. $\mathbf{1} \cdot \underline{\mathbf{x}}[t] \leq \Delta[t] \leq \mathbf{1} \cdot \bar{\mathbf{x}}[t]$, set $\mathbf{x}[t]$ as follows:*

$$\mathbf{x}[t] \leftarrow \Delta[t] \cdot \frac{\underline{\mathbf{x}}[t]}{\mathbf{1} \cdot \underline{\mathbf{x}}[t]} \quad \text{if } \Delta[t] \leq 0 \quad (5a)$$

$$\mathbf{x}[t] \leftarrow \Delta[t] \cdot \frac{\bar{\mathbf{x}}[t]}{\mathbf{1} \cdot \bar{\mathbf{x}}[t]} \quad \text{if } \Delta[t] > 0 \quad (5b)$$

This ensures that required changes in dispatch are shared in equal proportion among participants, hence the term “equitable dispatch”.

Proposition 3 (Sparse dispatch). *If $\Delta[t] > 0$:*

- 1) *Sort entries in $\bar{\mathbf{x}}[t]$ in descending order. This yields an ordered list of indices \mathcal{I} which is a permutation of $[1, 2, \dots, m]$ in some new order.*
- 2) *Find the smallest integer k such that $\sum_{i=1}^k \bar{x}_{\mathcal{I}_i}[t] \geq \Delta[t]$.*
- 3) *Then, the solution for each i^{th} entry is set as:*

$$x_{\mathcal{I}_i}[t] \leftarrow \begin{cases} \bar{x}_{\mathcal{I}_i}[t] & \text{if } i \in [1, k-1] \\ \Delta[t] - \sum_{j=1}^{k-1} \bar{x}_{\mathcal{I}_j}[t] & \text{if } i = k \\ 0 & \text{otherwise} \end{cases}$$

A similar approach is adopted in the case of $\Delta[t] < 0$, with appropriate changes in signs and inequalities. The term “sparse dispatch” refers to the fact that $(m-k)$ resource entities are allocated zero in their entries in $\mathbf{x}[t]$.

Proposition 2 allocates dispatch proportionally over all resource entities. This ensures that all are sharing the burden of deployment, i.e. sacrificing mining for financial reward. However, this also means that many physical resources are used which could lead to more wear-and-tear. To address this issue, we present Proposition 3 which selects a minimal amount of facilities, those with the fastest response, to collectively meet the required dispatch. While this minimizes the amount of resources used, it has a side effect of punishing larger and/or more nimble facilities, since the market only pays for capacity in the day-ahead market, and not for real-time deployment. This could yield a long term negative effect of less participation by such facilities. As a result, we suggest that the former proposition might be more suitable, but we encourage the reader to consider both.

Later in the paper, we demonstrate how adopting such simple approaches can yield satisfactory results for the physical grid. In the next section below, we explore when it is profitable for mining facilities to adopt such approaches.

III. MINING FACILITIES PROFIT FORMULATION

In this section, we formulate and quantify the economic benefit of mining facilities participating in Reg-Up/Down. In particular, the question we pose in this section is: when is it profitable for mining facilities to participate in the day-ahead Reg-Up/Down market?

A. Mining Rate of Return Pre-Reg-Up/Down

To properly quantify a mining facility's profit during participation in the day-ahead Reg-Up/Down market, we ought to first quantify its profit prior to any participation in this market for a fair comparison post-participation. In particular, we consider the case of Bitcoin mining, since Bitcoin is the leading cryptocurrency among those that rely on proof-of-work, i.e. on electric power consumption [22].

Consider a 24-hour period, and let integer $h \in [0, 23]$ denote each hour of this period. In Section II-B, we let n denote the number of QSEs participating in the day-ahead market, $m \leq n$ of which were cleared with non-zero capacity for dispatch. Without loss of generality, consider that all of those QSEs represent mining facilities, since we are only interested in quantifying the profit for such entities. Let $\lambda_i[h]$ denote average electricity price (in \$/MWh) for QSE $i \in [1, n]$ during hour h . That is, the cost of energy incurred by entity i at the end of hour h is the amount of energy it consumed during that hour multiplied by $\lambda_i[h]$.

Let $u_i^{\text{BTC}}[h]$ denote the long-term utility per Bitcoin mined at hour h as valued by QSE i . Moving forward, we express utility in terms of U.S. dollars to perform like-for-like comparisons. Therefore, the units of $u_i^{\text{BTC}}[h]$ are \$/BTC. Note that we do in fact employ subscript i even though Bitcoin prices are the same for everyone. We do this to distinguish QSEs according to how they perceive the Bitcoins they mine from a long-term perspective. For example, one QSE might elect to immediately sell all Bitcoins it mines, while another might choose to hold on to them for an arbitrarily long time. Hence, although the units of $u_i^{\text{BTC}}[h]$ are \$/BTC, this quantity is not necessarily equal to the currency conversion rate at hour h .

Let $e_i^{\text{BTC}}[h]$ (in MWh/BTC) denote the total amount of energy required to mine one whole Bitcoin in all facilities collectively represented by QSE i for hour h . It is safe to assume that over a 24 hour period, this quantity remains fixed for all hours, but might vary across QSEs, as mining equipment might vary at different facilities. In Section IV, we provide realistic estimates for this quantity.

Define the *mining rate of return* as the value per unit energy returned from mining (in \$/MWh). Let $\rho_i[h]$ denote this amount for QSE i and hour h , defined as follows:

$$\rho_i[h] := \frac{u_i^{\text{BTC}}[h]}{e_i^{\text{BTC}}[h]} - \lambda_i[h] \quad (6)$$

Note that if electricity price $\lambda_i[h]$ exceeds $u_i^{\text{BTC}}[h]/e_i^{\text{BTC}}[h]$ for some hour h , then the mining facility is better off not mining. Therefore, in the absence of any Reg-Up/Down participation, the effective rate of return is actually $\max(0, \rho_i[h])$, not $\rho_i[h]$, during hour h assuming that QSE i is acting rationally.

B. Profit Post-Reg-Up/Down

In Section II-B, the day-ahead market clearing process was introduced. Let $\pi^{\text{up}}[h]$ and $\pi^{\text{dn}}[h]$ (in \$/MW) denote the market-cleared Reg-Up and Reg-Down prices respectively for hour h . Those prices are the same for all participants, but the cleared capacities vary across QSEs. Let $c_i^{\text{up}}[h]$ and $c_i^{\text{dn}}[h]$ denote the market-cleared Reg-Up and Reg-Down capacities respectively for QSE i at hour h , resulting in the following reward collected in the day ahead market:

$$R_i^{\text{DAM}}[h] := \pi^{\text{up}}[h]c_i^{\text{up}}[h] + \pi^{\text{dn}}[h]c_i^{\text{dn}}[h] \quad (7)$$

Let $C_i[h]$ denote the maximum possible capacity that QSE i is able to either mine with or submit offers with at hour h . Indeed it is not necessarily constant over time. For example, the QSE might be involved in some separate demand response program that limits it to some arbitrary amount for that hour. Furthermore, the people behind QSE i might not be willing to offer the full amount each hour. Those are some of the factors that may limit $C_i[h]$.

Now let us say that mining rate of return $\rho_i[h]$ is positive. This means that ideally the mining facilities governed by QSE i should mine as much as possible during hour h , which happens to be $C_i[h]$. However, since the QSE is cleared $c_i^{\text{dn}}[h]$ in day-ahead, it can only mine up to $C_i[h] - c_i^{\text{dn}}[h]$ to ensure that an amount of $c_i^{\text{dn}}[h]$ is readily available for Reg-Down. In contrast, if $\rho_i[h] < 0$, then mining any amount at all yields a loss, but since the QSE is cleared $c_i^{\text{up}}[h]$ in day-ahead, it must mine at least $c_i^{\text{up}}[h]$ to ensure that an amount of $c_i^{\text{up}}[h]$ is readily available for Reg-Up. In either case, we can express the net rewards collected in real-time due to mining **before any real-time deployment** of Reg-Up/Down as follows:

$$R_i^{\text{base}}[h] := \rho_i[h] \begin{cases} C_i[h] - c_i^{\text{dn}}[h] & \text{if } \rho_i[h] \geq 0 \\ c_i^{\text{up}}[h] & \text{if } \rho_i[h] < 0 \end{cases} \quad (8)$$

This amount strictly decreases for any increase in either $c_i^{\text{up}}[h]$ or $c_i^{\text{dn}}[h]$, reflecting a risk being taken during any participation.

Finally, we formulate the rewards collected due to deployment of Reg-Up/Down. In Section II-C, we formulated real-time dispatch, wherein we detail how real-time frequency determines dispatch. Since we do not know a priori how much

dispatch is needed, we treat it as a random variable in this section. In Section IV-A, we explore the distribution of this random variable based on real-world deployment data from ERCOT. Let $\epsilon_i^{\text{up}}[h] \leq 0$ and $\epsilon_i^{\text{dn}}[h] \geq 0$ denote deployment rates for Reg-Up and Reg-Down, with the former being strictly non-positive to reflect the fact that Reg-Up deployment reduces consumption. This results in the following additional rewards collected in real-time due to deployment:

$$R_i^{\text{Reg}}[h] := \rho_i[h] (\epsilon_i^{\text{up}}[h] c_i^{\text{up}}[h] + \epsilon_i^{\text{dn}}[h] c_i^{\text{dn}}[h]) \quad (9)$$

Let $R_i[h]$ denote the net rewards, or profit, for QSE i over hour h . This is defined simply as the sum of rewards collected in day-ahead and real-time, i.e. $R_i^{\text{DAM}}[h] + R_i^{\text{base}}[h] + R_i^{\text{Reg}}[h]$. To distinguish between the impacts of $c_i^{\text{up}}[h]$ and $c_i^{\text{dn}}[h]$ on the net reward, we define two auxiliary variables as follows:

$$w_i^{\text{up}}[h] := \pi^{\text{up}}[h] + \rho_i[h] (\epsilon_i^{\text{up}}[h] + \mathbb{1}_{\{\rho_i[h] < 0\}}) \quad (10a)$$

$$w_i^{\text{dn}}[h] := \pi^{\text{dn}}[h] + \rho_i[h] (\epsilon_i^{\text{dn}}[h] - \mathbb{1}_{\{\rho_i[h] \geq 0\}}) \quad (10b)$$

Auxiliary variables $w_i^{\text{up}}[h]$ and $w_i^{\text{dn}}[h]$ effectively indicate how *worthwhile* it is for QSE i to participate in Reg-Up/Down during hour h . In the next subsection, we address when it is worthwhile to participate, and by how much capacity. Using those variables, we can express net rewards concisely as:

$$\Rightarrow R_i[h] = w_i^{\text{up}}[h] c_i^{\text{up}}[h] + w_i^{\text{dn}}[h] c_i^{\text{dn}}[h] + C_i[h] \max(\rho_i[h], 0) \quad (11)$$

The expression in Eq. (11) contains three distinct terms. The first two suggest that for each of Reg-Up or Reg-Down, it is only worthwhile to have been cleared for capacity if the corresponding auxiliary variables are positive. The last term is completely independent of cleared capacities, and is equal to revenue without participation in Reg-Up/Down whatsoever.

C. When to Participate in Reg-Up/Down

The only time QSE i can make a decision regarding its participation in Reg-Up/Down is once a day in the day-ahead market. Once it submits price-capacity offers, it can no longer control (a) at what prices or capacities it is cleared, or (b) how much it is deployed in real-time. Nonetheless, each QSE i can make an informed decision based on $R_i[h]$ for each hour h since market operators, such as ERCOT, tend to publish historic records of deployment, prices and capacities cleared. ERCOT also publishes year-ahead inelastic Reg-Up/Down demand for each hour. Thus, the decision of whether to participate in the day-ahead market boils down to whether $\mathbb{E}[w_i^{\text{up}}[h] c_i^{\text{up}}[h] + w_i^{\text{dn}}[h] c_i^{\text{dn}}[h]] > 0$ (i.e. profitable) or not. In Section IV-A, we explore real-world hourly data for multiple years to demonstrate when it is profitable.

The amount of potential reward collected is dependent on how much capacity the QSEs are cleared for Reg-Up/Down in any given hour. This can be expressed as the solution to a convex optimization problem as follows. First, for each QSE i and hour h , consider the problem below.

$$\text{maximize}_{c_i^{\text{up}}[h], c_i^{\text{dn}}[h]} \quad \mathbb{E}[w_i^{\text{up}}[h] c_i^{\text{up}}[h] + w_i^{\text{dn}}[h] c_i^{\text{dn}}[h]] \quad (12a)$$

$$\text{s.t.} \quad c_i^{\text{up}}[h] + c_i^{\text{dn}}[h] \leq C_i[h] \quad (12b)$$

$$c_i^{\text{up}}[h] \geq 0 \quad (12c)$$

$$c_i^{\text{dn}}[h] \geq 0 \quad (12d)$$

The constraint in Eq. (12b) is justified as follows. For a given cleared capacity $(c_i^{\text{up}}[h], c_i^{\text{dn}}[h])$, the QSE is obliged to operate at least at $c_i^{\text{up}}[h]$ and at most at $C_i[h] - c_i^{\text{dn}}[h]$ to guarantee delivery of maximum capacity if needed. Thus, $C_i[h] - c_i^{\text{dn}}[h] \geq c_i^{\text{up}}[h]$ is required, which is equivalent to Eq. (12b). If QSE i can determine an estimate for $\mathbb{E}[w_i^{\text{up}}[h]]$ and $\mathbb{E}[w_i^{\text{dn}}[h]]$ for each hour h , as illustrated in Section IV, the objective in Eq. (12a) can be approximated as $\mathbb{E}[w_i^{\text{up}}[h] c_i^{\text{up}}[h]] + \mathbb{E}[w_i^{\text{dn}}[h] c_i^{\text{dn}}[h]]$. Finally, this reduces Eq. (12) into a linear program, whose solution can be shown to be the following:

$$(c_i^{\text{up}}[h], c_i^{\text{dn}}[h])^* \leftarrow \begin{cases} (C_i[h], 0) & \text{if } \mathbb{E}[w_i^{\text{up}}[h]] \geq \mathbb{E}[w_i^{\text{dn}}[h]] \\ & \text{and } \mathbb{E}[w_i^{\text{up}}[h]] > 0 \\ (0, C_i[h]) & \text{if } \mathbb{E}[w_i^{\text{dn}}[h]] > \mathbb{E}[w_i^{\text{up}}[h]] \\ & \text{and } \mathbb{E}[w_i^{\text{dn}}[h]] > 0 \\ (0, 0) & \text{otherwise} \end{cases} \quad (13)$$

That is, for a fixed estimate for $\mathbb{E}[w_i^{\text{up}}[h]]$ and $\mathbb{E}[w_i^{\text{dn}}[h]]$, we can conclude two things:

- (a) It is never optimal to be cleared for both Reg-Up and Reg-Down simultaneously.
- (b) The optimal net reward for QSE i at hour h is:

$$R_i^*[h] = C_i[h] \max(\mathbb{E}[w_i^{\text{up}}[h]], \mathbb{E}[w_i^{\text{dn}}[h]], 0) \quad (14)$$

The intuition behind the three possible solutions in Eq. (13) is the following. $\mathbb{E}[w_i^{\text{up}}[h]]$ and $\mathbb{E}[w_i^{\text{dn}}[h]]$ are the expected per-MW-capacity profit due to participation in Reg-Up/Down for some hour h . Therefore, if neither are positive, then it is simply best not to participate, and if any is positive, then the greater of the two indicates whether to participate in Reg-Up or Reg-Down. In this paper, we do not further explore a risk-controlled optimization formulation, but in such a case, one could argue that the solution will be a convex combination of the three possible solutions representing a less risky participation in the market. In the next section, we explore a realistic distribution for w^{up} and w^{dn} , and analyze when it is profitable for mining facilities to realistically participate in the Reg-Up/Down day-ahead market.

IV. VIABILITY ANALYSIS: A TEXAS CASE STUDY

In this section, we present a Texas-based case study which quantifies and illustrates the feasibility of participation by Bitcoin mining facilities in ERCOT's Reg-Up/Down market. First, we rely on historical records of the quantities formulated in Section III to determine under what circumstances it is ideal for such facilities to participate in the market. Second, we simulate, using a synthetic model of the Texas grid, the impact of the control strategies proposed in Section II-C on both the electric grid and the participants involved in the market.

A. Real-World Historical Data

In Section II-C, a formulation for real-time regulation dispatch was presented, wherein a claim was made that simple linear integral control could closely approximate real-world dispatch. In Fig. 2, we display a time-synchronized pair of frequency and deployed regulation power, as observed by ERCOT, at a 4-second resolution over a two-hour window sampled from 2022. Here is what we observe from the data:

- Change in deployed regulation power ($\Delta[t]$ from Section II-C) is strongly correlated with frequency deviation from the nominal 60 Hz.
- Regulation is deployed even for minuscule frequency deviations (i.e. \underline{f} and \bar{f} from Section II-C would be set as very close to 60 Hz).
- For extended periods of time, regulation power saturates causing the frequency to remain away from 60 Hz.
- There is a visible ramp rate constraint which seems to limit how fast the frequency recovers. The amounts are approximately 1.50 MW/s and 1.15 MW/s for system-wide Reg-Up and Reg-Down respectively.

The mining facilities considered in this paper are a prime candidate at addressing the issue of limited ramping since they are able to ramp up/down consumption for an arbitrary amount of mining machines in parallel in a matter of a few seconds [8]. In contrast, as seen in Fig. 2, a similar task with existing service providers would take minutes, not seconds, to complete. If we are able to demonstrate that mining facilities can offer capacity with overall profit, then this would address the issue of saturation. We argue this case below.

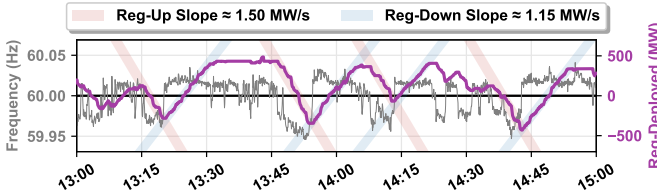


Fig. 2. Time-synchronized frequency and deployed regulation power, as observed by ERCOT, over a 2-hour window, showing a strong correlation between frequency deviation and change in regulation power.

In Section III, we concluded that whether or not it is profitable on average for QSE i to participate in Reg-Up/Down during hour h is determined by estimates for $\mathbb{E}[w_i^{\text{up}}[h]]$ and $\mathbb{E}[w_i^{\text{dn}}[h]]$. For the remainder of this subsection, we quantify such estimates based on historical records of the components that determine these estimates. Such components include **market cleared prices for capacity**, **Bitcoin prices**, **energy to mine a Bitcoin**, **electricity prices**, and **deployment rates**. Accordingly we reveal under what circumstances it is profitable to participate in each of Reg-Up and Reg-Down.

First, we consider **deployment rates**, denoted by ϵ^{up} and ϵ^{dn} in the previous section. In Fig. 3, we show the average capacity procured and average deployed power for each integer hour of the day $h \in [0, 23]$. We use negative values for Reg-Up in reference to the fact that Reg-Up corresponds to reduction in consumption, and for visual clarity. From these results, we draw the following conclusions:

- There are critical hours of the day (e.g. 6-7 AM) where the Reg-Up capacity is much larger than the deployed amount, i.e. ϵ^{up} is very small. Being cleared for Reg-Up capacity during these hours is desirable since it implies that very little opportunity cost is incurred (due to not mining) in exchange for large rewards from the market.
- In contrast, as explained in Section III, a large ϵ^{dn} is desired. Unfortunately, the results reveal ϵ^{dn} is mostly small since Reg-Down deployment (i.e. request to mine more) tends to be much smaller on average than capacity, except for very few hours (e.g. 2-5 AM).
- Demand for capacity tends to be quite modest, two orders of magnitude less than demand in the energy market. Consequently, if a few dozen mining facilities were to meet all the required capacity, $c_i^{\text{up}}[h]$ would be on the order of 10 MW or less on average for each i .

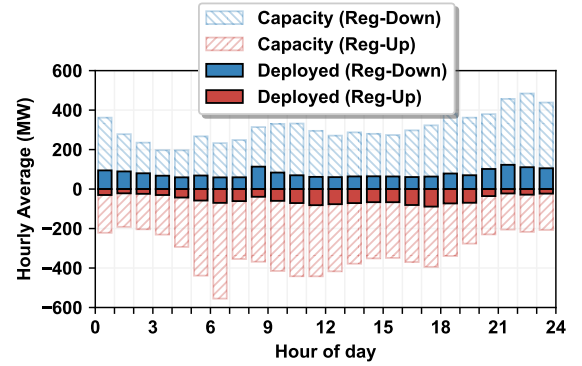


Fig. 3. Hourly average procured capacity and deployed amounts. Negative values refer to Reg-Up.

Second, we consider **market cleared prices for capacity**, denoted by π^{up} and π^{dn} in the previous section. In Fig. 4, we reveal that even in non-extreme cases, the prices, especially for Reg-Up, can reach sufficiently large amounts that it becomes more profitable to forgo mining in exchange for regulation service provision. For example, as shown in the bottom part of the figure, about 2.3% of the time, Reg-Up prices for capacity exceeded \$100/MW. This number by itself holds no meaning; rather, it's strictly about whether it is greater than the opportunity cost due to not mining or not. We illustrate this in the following example.

In Fig. 5, a 4-day example from the summer of 2022 is shown to illustrate that during certain hours of the day, even when the prices are much less than \$100/MW, it is net profitable to take a risk and to offer Reg-Up capacity. The quantity labelled *Expected Loss of Mining Profit* is effectively the non-price term in Eq. (10a) and is based on deployment rate ϵ^{up} . To use the same example, it is evident that deployment rates were sufficiently low during the afternoons such that it was profitable to be cleared for Reg-Up even at low prices. The example shown in Fig. 5 is for less than 1 week of data. Moving forward in this section, we examine the entire historical record spanning four years (Jan 2019 to Dec 2022) to perform a much more exhaustive analysis and determine under what circumstances it is more profitable to have been cleared

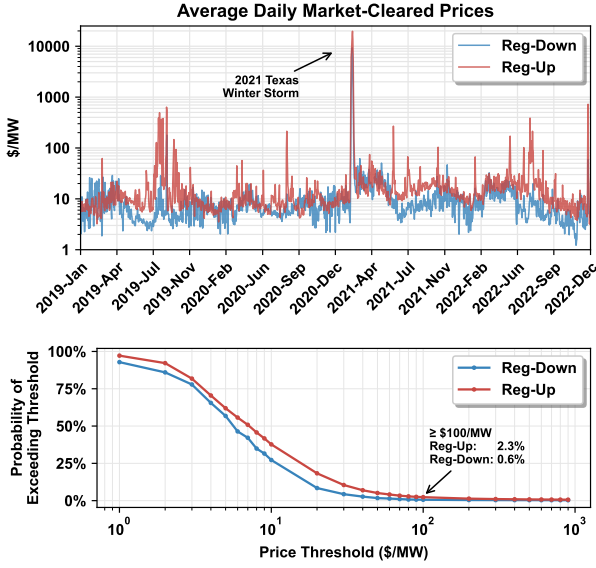


Fig. 4. Historical market-cleared prices for Reg-Up/Down. Timeline of prices (top), and probability of exceeding different price thresholds (bottom).

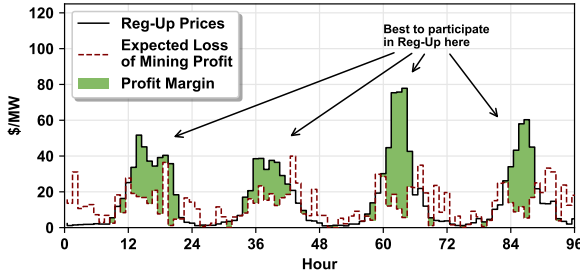


Fig. 5. An example from June 28 to July 2, 2022 illustrating a potential for profit due to participation in Reg-Up, even at low Reg-Up prices.

for capacity than not. Having considered the deployment rates and the market-cleared prices for capacity, what is left is to address is the energy required to mine Bitcoin as well as this cryptocurrency's prices over the years.

Finally, we consider the **energy required to mine Bitcoin** as well as **Bitcoin prices** and **electricity prices**. We rely on assumptions summarized in Table II to conclude the following. Consider two state-of-the-art Bitcoin mining machines, released in the summers of 2022 and 2021 respectively: Bitmain Antminer S19 XP and S19J Pro. According to the assumptions made in the table, based on real-world data and blockchain theory, the expected energy required to be awarded one Bitcoin from those two machines amounts to 147 MWh and 202 MWh respectively.

For example, at an exchange rate of \$22,050/Bitcoin, the S19 XP's revenue is roughly \$150/MWh. If we consider electricity prices ranging from 5¢/kWh to 10¢/kWh, the net opportunity cost for 1 hour of cleared capacity ranges from \$50/MW to \$100/MW. This is assuming 100% deployment for the entire hour. Fortunately, as shown in Fig. 3, deployment rates are well below 100%. As a result, the effective oppor-

TABLE II
ASSUMPTIONS ON BITCOIN MINING CAPABILITY (EARLY 2023)

	Quantity
Bitmain Antminer S19 XP (or S19J Pro)	21.5 (or 29.5) J/TH
Network Difficulty D (unitless)	39.35×10^{12}
Expected Hashes Per Block	$D/2^8$ TH/block
Bitcoin Reward (next halving: Apr. 2024)	6.25 Bitcoin/block

tunity cost reduces to as low as \$10/MW or less. However, if you consider Bitcoin prices greater than \$22,050/Bitcoin, the effective opportunity cost obviously increases. A historical record of average Bitcoin prices is shown in Fig. 7.

As a reminder, to determine whether it is profitable for QSE i to have been cleared for capacity during hour h , we need to find when $R_i^*[h] > 0$, as defined in Eq. (14). We perform an exhaustive search over the entire 4-year historical record and we reach the conclusion shown in Fig. 6. We purposely split the data into two: one portion before February 2021 and one after. The reason is twofold: (a) as shown in Fig. 4, prices were unusually high during the infamous 2021 Texas winter storm, and (b) as shown in Fig. 7, Bitcoin prices were significantly lower before 2021. As a result, when prices were higher, there were much less occurrences where $R_i^*[h] > 0$ after Feb. 2021. Surprisingly, the records reveal that even though there were less occurrences, the average profit was greater during that period, due to increased demand for ancillary services.

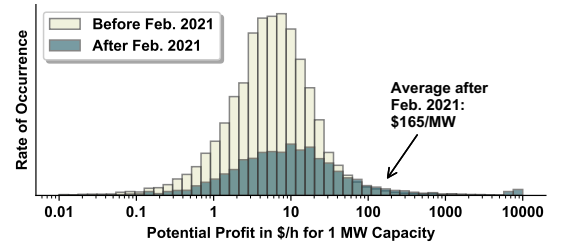


Fig. 6. Potential profit (Eq. 14) under proposed strategy (Eq. 13). Note that the horizontal scale is logarithmic.

Note that the scale in Fig. 6 is logarithmic, which means that the center of mass of the distribution is much farther to the right than where the peak appears to be on the figure. For example, the average is \$165/MW profit after Feb. 2021. Nonetheless, the key takeaway here is that even though profit is much smaller for ancillary services markets than it is for energy markets, it is still worthwhile for mining facilities to consider participation for the following reasons: (a) by providing much faster regulation service and more flexible capacity, they can significantly boost grid resiliency in the face of renewables, which better positions them in the eye of the public, and (b) they do not incur a net loss on average, and the overhead is low to participate since they are already in close communication with local grid operators with minimal communication infrastructure required to set up.

To get a closer look at when it is profitable to participate, the results in Fig. 6 do not suffice. We examine each month-

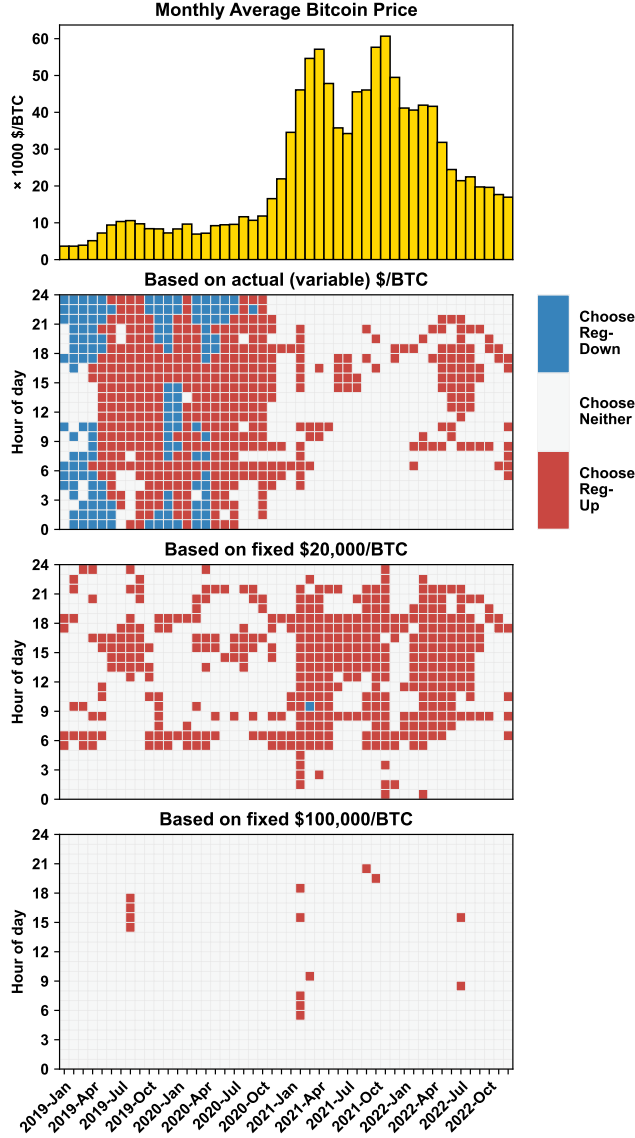


Fig. 7. Optimal choice (Reg-Up, Down or neither) for each month-hour pair, under different Bitcoin price scenarios, from Jan 2019 to Dec 2022.

hour pair in the historical record and correspondingly adopt Eq. (13) to choose either to participate in Reg-Up, Reg-Down or neither. Furthermore, we repeat this experiment under three different scenarios for Bitcoin prices: actual historical records, \$20,000/Bitcoin, and \$100,000/Bitcoin. The results are shown in Fig. 7, and we draw the following conclusions:

- There is a clear correlation between increase in Bitcoin prices and increase in choosing not to participate.
- The key separating factor between optimally choosing Reg-Up over Reg-Down is time of day.
- Even at \$100,000/Bitcoin (never yet witnessed in the historical record), it is still profitable to participate in Reg-Up, but only when the grid is under extreme stress, e.g. during summer heat waves or the 2021 winter storm.

The reason we consider month-hour pairs is based on a protocol followed by ERCOT, wherein, demand for capacity is fixed for each month-hour pair. For example, the demand

for Reg-Up capacity during the month of June 2022 at 10 AM is 629 MW for any day of that month, and this was calculated in a year-ahead study [2]. Accordingly, we claim that grouping the data in Fig. 7 by month-hour pairs is appropriate.

The result in Fig. 6 reflects a distribution over potential profit assuming that Bitcoin prices followed historical records. In that experiment, we also assumed a fixed 150 MWh/Bitcoin and a retail electricity price of 5¢/kWh. To further relax those assumptions, we conduct a series of experiment where we vary all of (a) Bitcoin prices, (b) MWh per Bitcoin and (c) electricity prices, and the results are shown in Fig. 8 and Fig. 9. In either figure, we sweep over Bitcoin prices from \$20,000/Bitcoin to \$100,000/Bitcoin and plot the average profit per hour under different scenarios. In Fig. 8, scenarios are defined by energy required to mine one Bitcoin, while in Fig. 9, scenarios are defined by electricity prices. Based on these results, we conclude that even under extreme situations, it continues to be profitable on average to participate in Reg-Up/Down. For example, for an average cleared Reg-Up capacity of 10 MW, a mining facility can net over \$6M/yr in profit. Note that this profit is calculated relative to the base case of no participation.

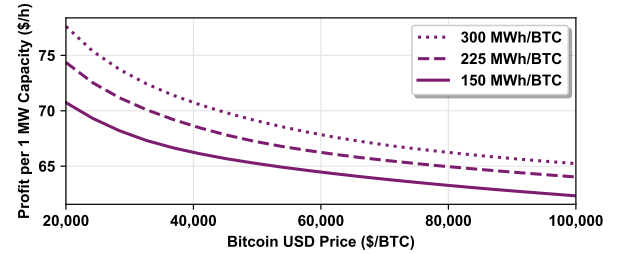


Fig. 8. Average profits under proposed strategy (Eq. 13) for variable Bitcoin prices and multiple scenarios of energy required to mine one Bitcoin.

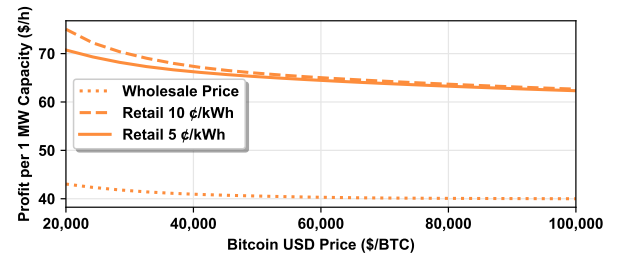


Fig. 9. Average profits under proposed strategy (Eq. 13) for variable Bitcoin prices and multiple scenarios of electricity prices.

B. Texas-Based Simulation

In Section IV-A, it was shown that existing ancillary service providers in ERCOT collectively provide Reg-Up support with a speed on the order of 1 MW/s. In this subsection, we present simulation results which reveal that mining facilities could significantly outperform existing providers at the same regulation task. We rely on a synthetic Texas-based 2000-bus transmission network model [23], [24] to simulate system-wide average frequency at the transient level, and the results

are shown in Fig. 10. In this simulation, there is loss in generation starting at time 5 seconds, and three different strategies for recovering the frequency are compared: (a) using same regulation speed from Fig. 2 (collectively 1.5 MW/s), (b) using ten times faster regulation, and (c) using a collection of large-scale mining facilities.

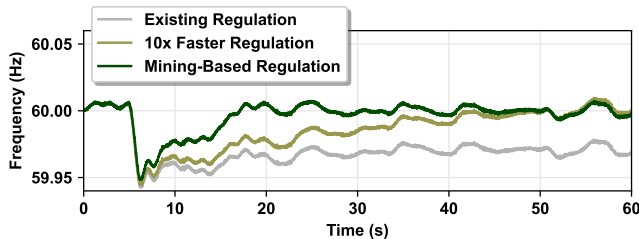


Fig. 10. Simulated system-wide average frequency in response to loss in generation in a Texas synthetic grid under different scenarios. It takes almost ten times the existing regulation speed to recover the frequency in under a minute. With a total of 500 MW from a collection of mining facilities, the frequency is recovered in much less time.

For the mining facility, we assume that each individual machine can be ramped down in 10 seconds, which is a conservative estimate according to reference [8]. Considering a collection of large-scale mining facilities which collectively offer 500 MW Reg-Up capacity, this yields an effective ramp rate of 50 MW/s. The results in Fig. 10 reveal that it takes only a few large-scale mining facilities to outperform even ten times more powerful regulation than what is currently used.

V. CONCLUDING REMARKS

This paper investigates the potential of utilizing cryptocurrency mining facilities for frequency regulation in electric energy systems. The frequency regulation problem is formulated considering various market participants with different physical constraints, and load dispatch strategies are presented to accommodate these differences. We present market participation strategies that help improve the quality of frequency regulation by exploiting the fast-responding potential of the mining loads while maximizing their operational profit. Our real-world case study and our synthetic transient-level simulations corroborate our theoretical analysis and highlight the advantages of our proposed methods from both economical and physical perspectives. Future work will examine proper ancillary service market designs to incentivize maximum participation of mining facilities in providing grid-level services.

REFERENCES

- [1] B. Wang, P. Dehghanian, and D. Zhao, "Chance-constrained energy management system for power grids with high proliferation of renewables and electric vehicles," *IEEE Transactions on Smart Grid*, vol. 11, no. 3, pp. 2324–2336, 2019.
- [2] Nitika Mago. (2021) Item 18: 2022 ERCOT Methodologies for Determining Minimum Ancillary Service Requirements. ERCOT. [Online]. Available: https://www.ercot.com/files/docs/2021/12/02/18_2022_ERCOT_Methodologies_for_Determining_Minimum_AS_Requirements.pdf
- [3] Y.-K. Wu, S. M. Chang, and Y.-L. Hu, "Literature review of power system blackouts," *Energy Procedia*, vol. 141, pp. 428–431, 2017, power and Energy Systems Engineering. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1876610217354619>

- [4] R. Yan, N.-A. Masood, T. Kumar Saha, F. Bai, and H. Gu, "The anatomy of the 2016 south australia blackout: A catastrophic event in a high renewable network," *IEEE Transactions on Power Systems*, vol. 33, no. 5, pp. 5374–5388, 2018.
- [5] J. W. Busby, K. Baker, M. D. Bazilian, A. Q. Gilbert, E. Grubert, V. Rai, J. D. Rhodes, S. Shidore, C. A. Smith, and M. E. Webber, "Cascading risks: Understanding the 2021 winter blackout in texas," *Energy Research & Social Science*, vol. 77, p. 102106, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2214629621001997>
- [6] A. Menati, Y. Cai, R. E. Helou, C. Tian, and L. Xie, "Optimization of cryptocurrency miners' participation in ancillary service markets," *arXiv preprint arXiv:2303.07276*, 2023.
- [7] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Potentials and economics of residential thermal loads providing regulation reserve," *Energy Policy*, vol. 79, pp. 115–126, 2015.
- [8] Steve Kinard. (2022) LFLTF Collaboration Deck, ERCOT Large Flexible Load Task Force (LFLTF) Meeting. Texas Blockchain Council. [Online]. Available: <https://www.ercot.com/calendar/04262022-Large-Flexible-Load-Task>
- [9] Y. Shi, B. Xu, D. Wang, and B. Zhang, "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2882–2894, 2017.
- [10] B. Aksanli and T. Rosing, "Providing regulation services and managing data center peak power budgets," in *2014 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2014, pp. 1–4.
- [11] S. Li, M. Brocanelli, W. Zhang, and X. Wang, "Integrated power management of data centers and electric vehicles for energy and regulation market participation," *IEEE Transactions on Smart Grid*, vol. 5, no. 5, pp. 2283–2294, 2014.
- [12] L. Chen and N. Li, "On the interaction between load balancing and speed scaling," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 12, pp. 2567–2578, 2015.
- [13] H. Chen, C. Hankendi, M. C. Caramanis, and A. K. Coskun, "Dynamic server power capping for enabling data center participation in power markets," in *2013 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2013, pp. 122–129.
- [14] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska, "Dynamic right-sizing for power-proportional data centers," in *2011 Proceedings IEEE INFOCOM*, 2011, pp. 1098–1106.
- [15] Y. Yao, L. Huang, A. Sharma, L. Golubchik, and M. Neely, "Data centers power reduction: A two time scale approach for delay tolerant workloads," in *2012 Proceedings IEEE INFOCOM*, 2012, pp. 1431–1439.
- [16] L. Zhang, S. Ren, C. Wu, and Z. Li, "A truthful incentive mechanism for emergency demand response in colocation data centers," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 2632–2640.
- [17] D. Quirk and M. Stabinski, "Cryptocurrency drives data center innovation," *ASHRAE Journal*, vol. 63, no. 10, pp. 44–49, 2021.
- [18] A. Menati, K. Lee, and L. Xie, "Modeling and analysis of utilizing cryptocurrency mining for demand flexibility in electric energy systems: A synthetic texas grid case study," *IEEE Transactions on Energy Markets, Policy and Regulation*, pp. 1–11, 2022.
- [19] Electric Reliability Council of Texas, "Large flexible load analysis," <https://www.ercot.com/calendar/02172023-LFLTF-Meeting> [Online], 2023.
- [20] Kenan Ögelman, "Overview of demand response in ertcot," https://www.ercot.com/files/docs/2021/09/15/ERCOT_Demand_Response_Summary_PUCT_Working_Session_091621.pdf [Online], 2021.
- [21] Electric Reliability Council of Texas, "Large flexible load resource participation in the ertcot region," https://www.ercot.com/files/docs/2022/04/25/Load_Resource_Participation_for_LFLs.pptx [Online], 2022.
- [22] White House, "Climate and energy implications of crypto-assets in the United States," <https://www.whitehouse.gov/wp-content/uploads/2022/09/09-2022-Crypto-Assets-and-Climate-Report.pdf> [Online], 2022.
- [23] A. B. Birchfield, T. Xu, K. M. Gegner, K. S. Shetye, and T. J. Overbye, "Grid structural characteristics as validation criteria for synthetic networks," *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 3258–3265, Jul. 2017. [Online]. Available: <https://doi.org/10.1109/tpwrs.2016.2616385>
- [24] T. Xu, A. B. Birchfield, K. S. Shetye, and T. J. Overbye, "Creation of synthetic electric grid models for transient stability studies," in *The 10th Bulk Power Systems Dynamics and Control Symposium (IREP 2017)*, 2017, pp. 1–6.