

Procuring Spontaneous Session-level Resource Guarantees for Real-time Applications: An Auction Approach

Madhumitha Harishankar, Sireesha Pilaka*, Pragma Sharma*, Nagarjun Srinivasan, Carlee Joe-Wong, and Patrick Tague

Abstract—Real-time multimedia applications such as interactive gaming, live video streaming, and augmented reality have strict latency and bitrate requirements. However, unpredictable network conditions like congestion and link quality can severely degrade the Quality of Experience (QoE). While buffer-based mitigations cannot be applied to real-time applications due to their immediate resource needs, recent innovations in network slicing have demonstrated the feasibility of dedicating specified amounts of network resources to individual sessions in the radio access network. Encouraged by this, we propose to reserve network resources for multimedia sessions *in real time* according to their declared needs, thereby providing *ad hoc session-level performance guarantees*. Through WiFi experiments and trace-driven LTE simulations, we show that such session-level resource provisioning is robust to real-time channel fluctuations and congestion externalities over the lifetime of a session. This approach, however, raises challenges: how can the network ensure that users are honest about their resource needs and optimally allocate its limited resources to users, *under uncertainty in future sessions' resource needs*? We derive a novel Multi-Unit Combinatorial Auction (MUCA) model with a unique structure that can be exploited for fast winner determination and yet incentivize truthful bidding, properties not simultaneously achieved in a generic MUCA but essential to making *real-time* session guarantees. Further, since dynamic bidding in real time is challenging for end-users who are budget-constrained, we develop a Reinforcement Learning based utility-maximizing strategy to distribute their budget across sessions, and show that it yields high user utility.

Index Terms—Quality of service, 5G mobile communication, Resource management, Economics, Machine Learning, Network Slicing, Auction Theory, Reinforcement Learning, Performance guarantees

*These authors contributed equally.

Manuscript received October 19 2018; revised March 16 2019; accepted April 21 2019. This work was supported in part by NSF CNS-1751075 and DARPA contract no. HR001117C0048.

M. Harishankar, C. Joe-Wong and P. Tague are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, CA - 94035, USA (e-mail: {mharisha, cjoewong, tague}@andrew.cmu.edu).

S. Pilaka is with Visa Inc., CA - 94404, USA and was affiliated with the Information Networking Institute, Carnegie Mellon University during her involvement in this work (e-mail: spilaka@alumni.cmu.edu).

P. Sharma is with Amazon Lab126, CA - 94089, USA and was affiliated with the Department of Electrical and Computer Engineering, Carnegie Mellon University during her involvement in this work (e-mail: pragyas@alumni.cmu.edu).

N. Srinivasan is with Mist Systems, CA - 95014, USA and was affiliated with the Department of Electrical and Computer Engineering, Carnegie Mellon University during his involvement in this work (e-mail: nagarjun@alumni.cmu.edu).

I. INTRODUCTION

Cyber-Physical Systems (CPS) and the Internet of Things (IoT) are emerging paradigms for increasingly pervasive and real-time computing environments. Users are coming to rely implicitly on the availability of services like Amazon Alexa and Google Home for their day-to-day tasks, while ambitious next-gen offerings like HoloLens promise to enable new use cases for real-time augmented reality like telepresence. However, the network connectivity that these ubiquitous computing environments rely on is insufficient [1] even for current applications. Widespread mobile multimedia services such as video conferencing and interactive mobile gaming have specific resource needs to provide an acceptable Quality of Experience (QoE) for the end-user. Without any means of conveying these needs to the network, the network may not meet them [2], [3] leading, for instance, to Google deploying its network Espresso.

In this work, we propose to provide resource guarantees on a per-session basis to QoE-sensitive applications by reconciling available resources with their session needs. To see the benefits of such guarantees, consider a Skype user starting a video conferencing session for a job interview who cannot procure any guarantees for the call quality and performance. Mechanisms for QoS-aware allocation [4], [5] typically do not model resource consumption at session-level timescales, hence allowing demand spikes to potentially interrupt or degrade the call. Xu *et al.* [1] and others [6] investigate this problem of high variance in cellular resource availability in the context of real-time applications. With buffer-based reactive measures like HTTP DASH infeasible [7], [8], the authors propose a short-term up-ahead estimation of channel conditions to proactively adjust application behavior, thereby reducing perceived delay. Improved channel estimation, however, cannot help sessions that must be preempted altogether due to congestion or spiky traffic, thereby entirely disrupting the call. Hence, channel estimation itself is insufficient to *guarantee call quality or completion*. We instead propose to proactively allocate resources to real-time sessions to guarantee high QoE over their entire duration. Proactive resource provisioning has been studied [5], [9] at the time-scale of packets or transmission time interval (TTI), using dynamic pricing or auction-based methods to allocate limited resources. However, the resulting allocation at one TTI is entirely independent of the next, limiting their use here. For multimedia applications, while

millisecond level network performance affects user perception, user engagement and ultimately QoE occur at the session level. Therefore, the user’s QoE depends on the resource allocation throughout the duration of the session, which is on the order of minutes or hours. **Our goal in this work is to proactively provide resource guarantees expressed in terms that a user agent or application can understand and negotiate for, abstracting away the lower-level intricacies of network resource allocation as details left to the network operator.** Similar mechanisms have been proposed for the Internet backbone [10] and cloud environments [11], [12], and wireless applications are likely to benefit from them even more due to their best-effort nature. However, this also makes it challenging to provide such guarantees in wireless networks. In fact, offering multiple tiers of service guarantees to wireless users is the goal of the emergent *network slicing* paradigm in 5G. Critical components of network slicing are still in their infancy [13], including *RAN slicing*, i.e., designing the wireless radio access network to enforce per-flow performance guarantees, and *slice admission and management*. This work addresses these research challenges [14] that arise in enabling proactive, session-level resource provisioning for wireless networks.

Feasibility: To proactively provision flows for their anticipated duration, available spectrum resources must be quantified and accurately reconciled with session requirements presumably expressed in terms of bitrate, latency and duration. Is this feasible? Even then, are performance guarantees possible despite uncontrollable wireless influences like fast fading? Recent works in RAN slicing [15], [16] facilitate functional slice isolation and empirically analyze various factors to provide probabilistic performance guarantees in cellular networks. We herein employ an admission control algorithm that allows a flow into the network based on a simple resource forecast and reconciliation model that is generalizable to both scheduled and random-access wireless networks. We conduct extensive WiFi experiments and trace-driven LTE simulations with multimedia applications. In Sections II and III, we find that admitted latency-sensitive as well as bitrate-heavy flows achieve their promised performances and congestion externalities are effectively mitigated. In fact *the network accommodates even more flows by implementing incentivized admission control*. Further, since this may be offered as a *value-added service* that some users may not require, we show reliable guarantees can be made even in the presence of background flows not controlled by our admission algorithm.

Allocation and Incentive Compatibility: Given a forecast of network resource availability and a mechanism to reconcile this with session needs, how should these limited resources be provisioned? The network will likely need to prioritize users with higher resource valuations as it cannot accommodate all session requests. Variation in such valuations could arise from usage context (medium quality for recreational video calls but high for an interview) or device preferences (lower resolution on a smartphone vs a 4K monitor). Further, allocating resources for the duration of a session is particularly difficult as the operator must account for uncertainty in future needs, and users may strategically misrepresent their needs and

valuations. We address these concerns in a novel auction model introduced in Section IV. The operator offers consecutive auctions throughout the day, and users relay their sessions’ resource needs dynamically in a combinatorial bid to the current auction; session durations may span multiple auctions. In Section V, we show that the spontaneous and real-time nature of sessions can be exploited to reduce the search space of the intractable optimization problem of determining winning bids, thereby facilitating *spontaneous guarantees*. In Section VI, we propose multiple ways for the operator to *incentivize truthful user declarations* even under uncertainty of future bid arrivals and analyze *trade-offs in social welfare, incentive compatibility and operator revenue*.

Usability: For users to procure and benefit from performance guarantees in this system, they must engage in routine auctions by bidding. However, studies have shown [17] that dynamic pricing is challenging for end users who are budget constrained and averse to making real-time network consumption decisions. We consequently address the user-facing challenges of *resource-specification overhead, price discovery and budget constraints*. We envision that an automated agent will participate in these session-oriented resource auctions on each user’s behalf, placing bids using a parameterized utility model and enforcing the user’s daily budget. In Section VII, we formulate the distribution of this budget across bids as a dynamic program solved with model-free reinforcement learning, specifically the Monte Carlo policy iteration algorithm [18]. We show via simulation in Section VIII that these agents *maximize user utility* for a given budget within a billing cycle (1 month) *without any loss in revenue to the network operator*.

Overall, *we formulate an end-to-end system for realizing session-level performance guarantees, addressing challenges in the radio access network, incentive mechanisms for resource provisioning, and usability*. We defer discussion of related work to Section X and conclude in Section XI. Additional experimental results and proofs can be found in our extended report [19].

II. FEASIBILITY OF SESSION-LEVEL PERFORMANCE GUARANTEES OVER LTE

Providing session-level performance guarantees requires an Admission Control (AC) procedure that only admits flows with demands that can be fulfilled for the stated session duration. This reconciliation of available and required resources is then expected to result in admitted flows that are robust to externalities and realize their promised performances. While unpredictable real-time channel fluctuations are inevitable (e.g., fast fading), our premise is that the short timescale of these fluctuations affects session-level QoE less than user competition, which occurs at session-level timescales. We validate this via a proof-of-concept trace-driven simulation of users with different traffic types sharing the resources of an LTE eNodeB. We show session performance can be guaranteed by an AC that 1) accounts for resource competition in flow admission and 2) accommodates for unpredictable wireless externalities when provisioning capacity.

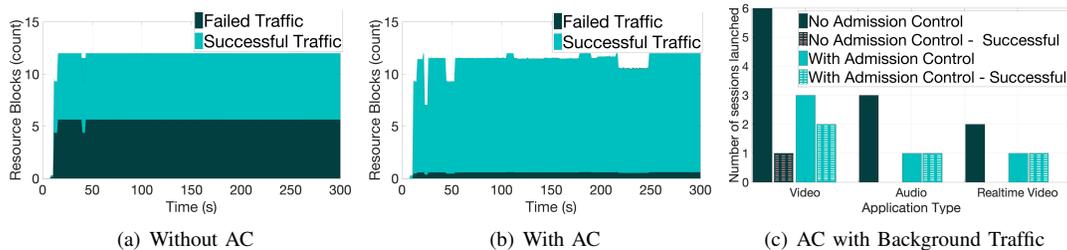


Fig. 1: (a) Without a resource-aware Admission Control (AC) algorithm, almost half of the LTE network’s resources are expended in failed sessions. (b) With the AC in place, this is reduced to $\sim 5\%$ while *preserving high network utilization* and (c) most performance guarantees are met even in the presence of uncontrolled background traffic.

Setup: We use SimuLTE and INET¹ to simulate LTE with TCP/UDP/IP. Our simulation includes one eNodeB with 12 resource blocks and a noisy channel. Users are randomly dispersed in the coverage area, yielding variation in channel qualities. Data usage is modeled from the multimedia activity found in mobile traffic traces of 20 users collected over 10 days. Multimedia content from the traces includes *video streaming*, *audio streaming* and *real-time video conferencing*, parameterized by bitrates and latencies of known applications^{2,3,4}. Since the AC algorithm compares these required quantities with the available capacity to determine flow feasibility, it must translate granular frequency-time network blocks into bitrate and latency capacity forecasts. While devising an accurate model for this is a challenging task in itself and out of our scope, we presume a naïve model that is sufficient to indicate general feasibility and benefits of this approach. In essence, the AC procedure maps the 12 resource blocks to a conservative estimate of bitrate capacity (e.g., 10Mbps), thereby allowing a buffer of radio capacity that may be consumed, for example, by an admitted flow(s) with poor signal strength or temporary channel degradation. A flow is admitted after verifying that its required capacity can be accommodated, and the capacity forecast reduced according to the requested bitrate for the specified duration.

Network Performance with AC: We measure link-layer utilization both with and without session-level AC. Figure 1(a) shows that without AC, roughly 47% of resource blocks are allocated to sessions that fail (i.e., the stream halts before completion) mainly from excessive resource competition. However, as in Figure 1(b), the AC algorithm drastically reduces this wastage (to below 5%) while *preserving nearly full utilization* of available resources. While the AC allocates a conservative 10 Mbps for flow provisioning to guard against externalities, this would presumably leave the network underutilized. In this case, however, congestion between flows and impact from other externalities were almost entirely eliminated while retaining high utilization. These promised performances were achieved with noise and channel quality variation, indicating that session-level guarantees can be provided in wireless cellular networks without the significant

cost of network under-provisioning. Even the naïve model of capacity forecasting used by our AC procedure proved sufficient to serve most admitted flows of different application types. This validates our premise that wireless radio resource modeling and reconciliation with session-oriented resource requirements are feasible and can likely provide performance guarantees.

Since AC-admitted sessions may co-exist with unregulated background sessions (of users that do not require performance guarantees), we reserve some network capacity for this traffic and apply AC to the remaining capacity. We employ the MAX C/I scheduling at the MAC layer to prioritize flows with better channel quality [20], and modify the scheduler to further prioritize AC-admitted flows before ranking by channel quality. Figure 1(c) compares the resulting performance of sessions belonging to AC and non-AC traffic. An AC session is deemed *successful* if it streams for its entire duration at its guaranteed bitrate on average; real-time sessions additionally require a packet inter-arrival time below 40 ms. A non-AC session is successful if it achieves *any* resolution supported by its traffic type. As in Figure 1(c), the network is highly congested with non-AC traffic, so only one of eleven non-AC flows succeeds, while four of the five admitted AC flows succeed. Figure 1(c) demonstrates that session-level guarantees can be achieved in the presence of non-AC traffic with appropriate reservations in the capacity forecast and MAC prioritization. Our extended report [19] provides additional LTE-based simulation studies.

III. FEASIBILITY OF SESSION-LEVEL PERFORMANCE GUARANTEES OVER WiFi

We now assess whether performance guarantees can be made in a random access medium like WiFi. Unlike LTE, WiFi has a short range and operates in unlicensed spectrum, making the channel more susceptible to interference and externalities. A likely use case for performance guarantees, however, is where multiple users engaged with various apps contend for congested resources of a public WiFi network, for instance, in a café-like scenario. Measurement studies [21], [22] have shown extensive growth in public hotspot traffic and Access Point (AP) deployment, with WiFi traffic doubling every two years ($\sim 35\%$ video). These experiments verify the feasibility of session-level guarantees in such scenarios.

Setup: We launch 50 iPerf⁵ clients in parallel across multiple devices to induce channel quality variations. Clients

¹<http://simulte.com/>, <https://inet.omnetpp.org/>

²<https://support.skype.com/en/faq/FA1417/how-much-bandwidth-does-skype-need>, <http://download.skype.com/share/business/guides/skype-connect-requirements-guide.pdf>

³<https://help.pandora.com/customer/portal/articles/166391-minimum-specifications-to-run-pandora>

⁴<https://support.google.com/youtube/answer/2853702>

⁵<https://iperf.fr/>

connect to an 802.11g AP operating at 2.4GHz and launch five sequential sessions over 50 minutes, each comprising a random duration of video streaming, audio streaming or video conferencing (we continue to use resolution rates and corresponding bandwidth requirements that are widely in practice). We further incorporate non-AC web browsing traffic at 50 Kbps, thereby inducing overall activity variance typical of public WiFi.

Network performance without AC: As a baseline, we first engage the 50 clients in their planned mobile activity over this hotspot without the AC algorithm. Clients request the highest supported resolutions for their multimedia sessions (e.g., 1.5 Mbps and 4.5 Mbps for video conferencing and streaming, respectively) since they have no incentive to request lower bitrates due to free hotspot access. An aggregate data demand of up to 80 Mbps is seen in Figure 4(a). However, although the 802.11g AP has a theoretical capacity of 54 Mbps, *only half of this is realized by the network*, indicating severe performance degradation from congestion. The network also exhibits high latency and jitter, as in Figure 2(b), causing real-time video sessions (requiring ~ 100 ms) to fail or be lag-ridden.

AC Procedure: We now introduce our admission control process. Browsing sessions constitute light-weight traffic and always commence upon launch, modeling the background traffic of regular-access users as in the LTE experiments. The AC algorithm simply uses an estimated bitrate capacity as an abstracted representation of available radio resources and permits a session to start only if requisite session bandwidth is available. We herein refer to this as *Non-Incentivized AC* and introduce a corresponding *Incentivized* version. Since Non-Incentivized AC admits flows (subject to feasibility) on a first-come first-serve basis, users always request high bitrates even when they may be content with lower bitrates (e.g., when using a mobile device with low screen resolution). With Incentivized AC, we presume that an incentive mechanism induces clients to request only their value-maximizing resolutions for multimedia sessions. This incentive mechanism may, for instance, be a payment policy that charges admitted clients to persuade them to state only what they need (developed in Sections IV, V, VI). Using Incentivized AC, users are thus admitted according to their valuation for the appropriate context. We simulate Incentivized AC with clients streaming multimedia sessions at a resolution that is randomly chosen from the supported ones, simulating the distribution of utilities, preferences and budgets in a population.

Performance With AC: While the AP's theoretical capacity is 54 Mbps, this is rarely realized in practice due to time-varying nature of the wireless channel. Since provisioning based on this capacity may result in poor performance of some flows, the AC procedure is initialized with a capacity of 50 Mbps. Figure 3 depicts the number of clients admitted into the network under AC. A few clients are consistently rejected for the first half hour due to lack of capacity as aggregate data demand is highest then, Figure 2(a). However, with Incentivized AC, *the entire pool of 50 clients is admitted into the network* at all times that they initiate multimedia flows. As clients distribute their requesting resolutions in alignment with their true utilities and valuations, aggregate

user demand decreases so much that the network has sufficient capacity to now admit *all of them* (in this case), thereby *increasing the utility of the entire set of users*. Even when the provisionable capacity is reduced to 40 Mbps and number of clients increased to 75, the network admits them all with Incentivized AC, thereby increasing net social welfare of users.

With AC in place, *network throughput increases* to almost 50 Mbps, as in Figure 4(a). Due to random access in WiFi, congestion externalities have a severe impact on the network and are almost entirely mitigated with AC. When the Incentivized AC is deployed, as in Figure 4(a), only around 30 Mbps is typically required of the network now, indicating that incentivizing users for truthfulness may *allow the network to serve more users* overall. With AC, all real-time sessions stream at mean latencies below 50 ms, as in Figure 4(b), while introducing incentives causes lower aggregate demand which further reduces the mean latency to a maximum of 20 ms. In fact, the CDF of per-packet jitter for real-time sessions across all experiments in Figure 4(c) indicates that while more than 80% of packets exhibit jitter above 40 ms without AC, $\sim 80\%$ of packets experience jitter below 40 ms with Non-Incentivized AC and below 20 ms (recommended for real-time video conferencing and gaming) with Incentivized AC.

By controlling flow admittance based on bitrate demand, performance guarantees are delivered to latency-sensitive real-time flows as well as other video and audio streams. We thus validate our premise of session-oriented wireless resource provisioning (more results in our extended report [19]) and shift our focus to the design of the incentive mechanism employed by the AC algorithm. Note that given a process for forecasting network resource availability and reconciling it with session demands, our incentive mechanism is agnostic to the Radio Access Technology (RAT) in use, a point we discuss further in Section IX.

IV. MODELING AUCTIONS FOR SESSION-LEVEL RESOURCE GUARANTEES

We now design the incentive mechanism that determines session admission and cost. We focus on real-time applications that require immediate access, but our model is generalizable to other application types. The network operator discretizes time into a total of T slots per day, and an auction A_t is held in each time slot t . All resource requests are assumed to have a maximum duration of ϕ_{max} time slots; a longer session may simply submit another bid for resources after ϕ_{max} slots. For instance, the network may hold an auction each minute (i.e. $T = 1440$), allowing users to procure resource guarantees almost spontaneously as they launch sessions, and set $\phi_{max} = 20$, for resources to be periodically freed up once every 20 minutes. Similarly, the operator supports a discrete set of resource **modes** m_n for $n = 1, \dots, M$, each corresponding to an operating bandwidth, bitrate, or similar. In Section V, we define a generic mechanism for the network to define these modes, which allows for a wide range of supported bitrates by common applications, while significantly reducing computational overhead of the auction. To characterize the resources being auctioned, the operator computes a forecast

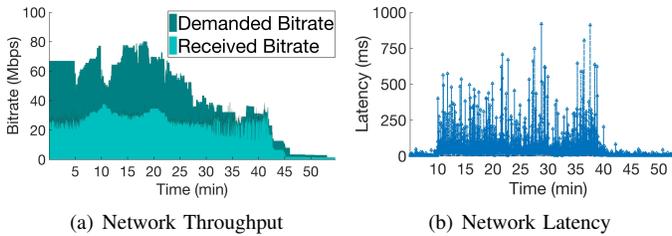


Fig. 2: (a) Due to high data demand, the network is congested and delivers around 30 Mbps throughput despite 54 Mbps capacity and (b) experiences latency spikes between 250 – 750 ms and high jitter.

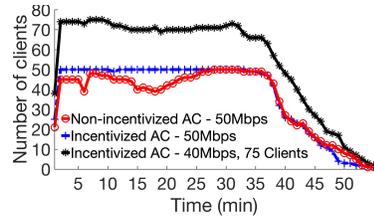


Fig. 3: The incentive mechanism induces heterogeneity in requested resolution rates that allows the network to admit all initiated sessions.

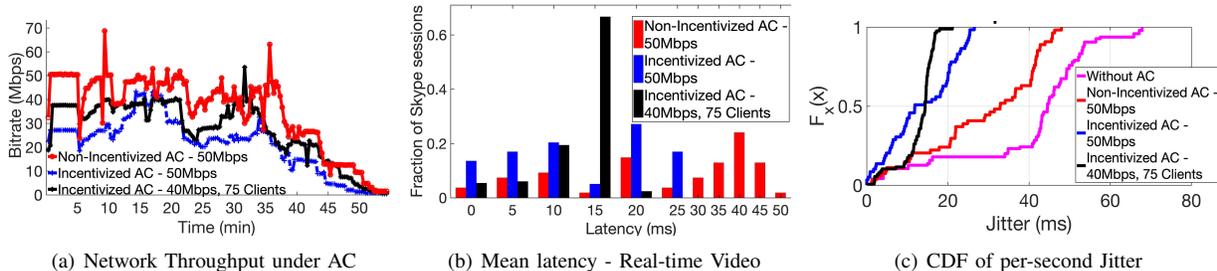


Fig. 4: With AC, (a) network throughput *increases* and real-time video sessions have (b) latencies within 25 – 50 ms and (c) jitter less than 40 ms. Incentivized AC further reduces data demand and improves both jitter and latency.

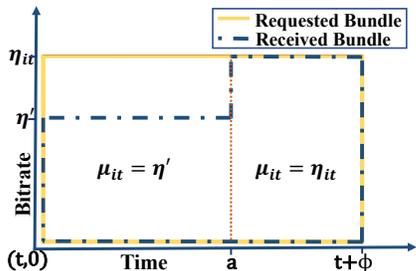


Fig. 5: A bundle B_{it} received by a user i for request R_{it} may provide a subset of the requested resources. For example, R_{it} specified η_{it} for the duration ϕ_{it} but received a bundle providing $\mu_{ik} = \eta' < \eta_{it}$ for $k < a$ and $\mu_{ik} = \eta_{it}$ for $k \geq a$.

$C_t^{(\tau)}$ for auction A_t of the bandwidth resources that will be available in each time slot $t + \tau$, $\tau \in \{0, \dots, \phi_{max}\}$, accounting for resources reserved in earlier winning bids.

The total number of users submitting bids for A_t is denoted by I_t . For users to express a desired combination of resources that the network can actually serve, they must know the resources available to bid on. We develop a two-round interaction mechanism for this resource discovery. In the first round of each auction A_t , user i expresses a request R_{it} that includes the desired duration $\phi_{it} \leq \phi_{max}$ and the desired resource mode $\eta_{it} \in \{m_1, \dots, m_M\}$, along with a corresponding valuation v_{it} . We assume the requested η_{it} is constant over the duration ϕ_{it} , as real-time applications typically have fairly stable resource needs over time. In response to resource requests $R_{it} = (\eta_{it}, \phi_{it})$ from users i , the network operator determines if granting R_{it} is feasible (given projected availability). If not, it generates a set S_{it} of alternate *resource bundles* (based on forecast capacity and adjusting for underlying wireless channel states), where each bundle $B_{it} \in S_{it}$ enumerates the offered resource $\mu_{ik} \leq \eta_{ik}$ for time slots $k = t + 1, \dots, t + \phi_{it}$.

Figure 5 illustrates an example bundle offered in response to a resource request. Given the set S_{it} of available bundles, user i may bid on a bundle $B_{it} \in S_{it}$ by assigning a new bid value v_{it} on it, yielding a bid $b_{it} = (B_{it}, v_{it})$ for the select bundle $B_{it} \in S_{it}$. Once auction A_t is executed, user i learns the result $x_{it} \in \{0, 1\}$ of the bid and starts consumption if $x_{it} = 1$.

Maximizing social welfare to determine bid winners is desirable since Vickrey-Clarke-Groves (VCG) [23] payments can then be charged to incentivize truthful bidding. The resulting computation, however, is an NP-hard problem, the solution time of which is exponential in bid durations. Thus, in the next section we develop novel reductions to the problem by exploiting the *spontaneous* nature of winning sessions, i.e. that they begin consumption immediately. With this, the network can implement the VCG mechanism in real time, and stating true product valuations v_{it} becomes the dominant strategy of users. This allows users to avoid complex estimation of other bidders' strategies to maximize their own utility and allows the network operator to discover the distribution of true valuations across bidders, indicating the perceived value of network resources and the potential revenue. However, bidders are *multi-parameter agents* [24] in this setting; they state not just their valuations but also the desired mode η and duration ϕ . Incentivizing bidders to truthfully report ϕ requires modifications to the mechanism that account for *temporal correlations* between the decisions taken in different auctions. We develop the resulting allocation and payment schemes in Section VI and analyze their auction properties. Proofs can be found in our extended report [19].

V. WINNER DETERMINATION

Since offering performance guarantees is a service in addition to users' normal mobile data plans, it will likely

represent a small portion of overall operator revenue, and the operator may rather wish to maximize users' welfare. Indeed, for "public utility" goods like network resources that are competitively auctioned, Cramton [25] argues for maximizing social welfare rather than network revenue for the sake of long-term user engagement. As such, we study the winner determination problem with the intent of optimizing social welfare and evaluate the achieved revenue via simulation in Section VIII. Thus, in this Multi Unit Combinatorial Auction (MUCA), the network maximizes the declared user valuations v_{it} in auction A_t subject to the resource capacity constraints over time slots, yielding the optimization problem:

$$\begin{aligned} \max_{\{x_{it} \in \{0,1\}\}} & \sum_{i=1}^{I_t} v_{it} x_{it} \\ \text{s.t.} & \sum_{i=1}^I \mu_{i(t+\tau)} x_{it} \leq C_t^{(\tau)}, \tau = 1, \dots, \phi_{\max}. \end{aligned} \quad (1)$$

We recognize this MUCA formulation as the NP-hard multi-dimensional knapsack problem (MKP) [26]. The dimensionality stems from the combinatorial nature of the bids, wherein they span multiple time slots (generalizable to multiple base stations and flows needing uplink/downlink capacity). Solving (1) is thus prohibitive for real-time network use. Many existing algorithms for fast MUCA winner determination [27]–[29] rely on assumptions such as bidder multi-mindedness, submodularity and low number of dimensions in the MKP, which do not apply to our auction model. Using approximation algorithms or other heuristics [30] to find a solution could result in significant loss of network revenue when the number of users or the bid duration increases, especially with frequently repeated auctions. More importantly, an exact solution to the MKP is required to incentivize users to bid truthfully in the auction [31]. We instead exploit the nature of *real-time flow* demands to reduce the complexity of (1) by considering a series of conditions on bid quantities, durations and resource availability. We show the simplification of (1) for each case, gradually leading up to more realistic and less restrictive conditions. The MKP mostly reduces to the knapsack problem solvable in pseudo-polynomial time [32].

A. Bundle Generation Policy

Our first task is to define the network operator's policy for generating the set S_{it} of bundles in response to a resource request R_{it} . The operator constructs a bundle B_{it} with resources at time $t + \tau$ given by $\mu_{i(t+\tau)} = \min(\eta_{it}, C_t^{(\tau)})$ for $\tau = 1, \dots, \phi_{it}$, corresponding to the highest possible resource level not exceeding the request η_{it} , based on projected availability. If $\mu_{i(t+\tau)} = 0$ for any τ , no bundle is offered to that user due to severe lack of resource availability. Hence the bundle a user receives comes closest to what the user requested for the specified duration, given capacity constraints. Given this construction of B_{it} , a bundle submitted to the auction may have different resource demands at different time slots, while the original request R_{it} does not. Note that if R_{it} is feasible, then B_{it} perfectly satisfies it by construction above,

and gets submitted immediately to the auction A_t . Constructed bundles may have the features defined below.

Definition 1 (Upswitch). A bundle B_{it} exhibits an **upswitch** if $\exists \tau \in [1, \phi_{it} - 1]$ s.t. $\mu_{i(t+\tau)} < \mu_{i(t+\tau+1)}$. The number of upswitches in B_{it} is denoted by $U_{B_{it}}$. Bundle B_{it} has an **a/b-upswitch** if $\exists \tau \in [1, \phi_{it} - 1]$ s.t. $a = \mu_{i(t+\tau)} < \mu_{i(t+\tau+1)} = b$.

Definition 2 (Downswitch). A bundle B_{it} exhibits a **downswitch** if $\exists \tau \in [1, \phi_{it} - 1]$ s.t. $\mu_{i(t+\tau)} > \mu_{i(t+\tau+1)}$. The number of downswitches in B_{it} is denoted by $D_{B_{it}}$. Further, bundle B_{it} has an **a/b-downswitch** if $\exists \tau \in [1, \phi_{it} - 1]$ s.t. $a = \mu_{i(t+\tau)} > \mu_{i(t+\tau+1)} = b$.

B. Reduction to Tractable Optimization Problems

We first consider a network capacity projection that increases monotonically over time. That is, if the resource availability projection $C_t^{(\tau)}$ for future timesteps $\tau \in [1, \phi_{\max}]$ shows no decline within that time period, then the auction round A_t satisfies this condition and is said to exhibit property $\mathbb{P}_1(t)$. This would hold, for instance, in any time slot t where no sessions carry over from previous auctions. The entire network's resource capacity is then available equally at t for all future time slots. If A_t exhibits $\mathbb{P}_1(t)$, then there can be no downswitches in any bundles submitted to this auction given our bundle generation policy, i.e., $D_{B_{it}} = 0 \forall i$. We now show several simplifications possible to (1) when $\mathbb{P}_1(t)$ holds, and also determine conditions under which $\mathbb{P}_1(t)$ is guaranteed to hold.

Definition 3 (Uniform quantity bid). A bid b_{it} is a **uniform quantity bid** on a bundle B_{it} if the corresponding resource levels μ_{it} are all equal, i.e., if $\forall \tau \in [1, \phi_{it} - 1], \mu_{i(t+\tau)} = \mu_{i(t+\tau+1)}$.

Theorem 1. If $\mathbb{P}_1(t)$ holds and each bid to auction A_t is a uniform quantity bid, then the outcome $\{x_{it}\}$ that solves (1) with only the capacity constraint at $\tau = 1$ is the solution to (1), reducing it to a knapsack problem. Further, $\mathbb{P}_1(t + 1)$ is guaranteed to hold.

While Theorem 1 simplifies the winner determination significantly by reducing the MKP to a single knapsack problem in one time slot, it only applies when all bids are of uniform quantity. Consider the following instance. The network has no active sessions at time t and projects $C_t^\tau = 10, \forall \tau$. At t , it admits two flows that consume 3 Mbps for 5 minutes and 5 Mbps for 7 minutes, respectively. At $t + 1$, the projected availability is $C_{t+1}^\tau = 2, \tau = 1, \dots, 5$ and $C_{t+1}^\tau = 5, \tau = 6, 7$, and $C_{t+1}^\tau = 10, \tau = 8, 9$, and so on. A user requesting 4 Mbps for 10 minutes at $t + 1$ would thus receive a bundle granting 2 Mbps for the first 4 minutes and 4 Mbps thereafter. This bundle exhibits an upswitch and is therefore not uniform. The network may be able to force the construction of suboptimal uniform bundles if this is a reasonable restriction for some applications or use cases. However, upswitches are likely due to varying availability constraints and data consumption patterns.

Lemma 1. *If a bundle B_{it} exhibits an a/b -upswitch at time τ , then any bundle in A_t corresponding to a request $R_{i't}$ with $\eta_{i't} \geq b$ exhibits an a/b' -upswitch with $b' \geq b$ at time τ .*

Proof. Given a bundle B_{it} with $a = \mu_{i(t+\tau)} < \mu_{i(t+\tau+1)} = b$, the bundle generation policy implies $C_t^{(\tau)} = a$, $\eta_{it} \geq b$, $C_t^{(\tau+1)} \geq b$. Hence, for any other bid $b_{i't}$ with $\eta_{i't} \geq b$, we have $\mu_{i'(t+\tau+1)} = \min(\eta_{i't}, C_t^{(\tau+1)}) \geq b$, $\mu_{i'(t+\tau)} = a$. \square

We now derive results that guide the network operator in defining its operating bitrate modes such that the MKP can be simplified even with upswitches. First, if the quantity expressed by each supported mode is equally spaced, we called these *evenly dispersed modes*.

Definition 4 (Evenly dispersed). *A set of modes $\{m_1, \dots, m_M\}$ is **evenly dispersed** if $\exists y, z \in \mathbb{N}$ s.t. $m_n = z(n + y)$, $\forall n \in \{1, M\}$.*

Theorem 2. *If the auction modes are evenly dispersed, then the outcome $\{x_{it}\}$ that solves (1) with only the capacity constraint at $\tau = 1$ is the solution to (1), reducing it to a knapsack problem. Further, $\mathbb{P}_1(t)$ holds for all t if auction modes are evenly dispersed.*

Theorem 2 allows the network operator to support different operating bitrates while solving (1) with a single knapsack. Specifically, if the supported modes were evenly dispersed, e.g., 2, 4 and 6 Mbps, then it is sufficient in every auction to solve (1) in the first time slot. The operator can choose the exact operating modes by examining those required by target applications and its ability to reserve resources. However, certain real-time applications may not lend themselves to this, e.g., Skype has discrete modes with unevenly dispersed bitrate requirements⁶, and the network may therefore offer arbitrary modes to serve these applications. We first note that even in this case, the upswitch count $U_{B_{it}}$ cannot exceed $M - 1$ as long as $\mathbb{P}_1(t)$ holds.

Theorem 3. *If $\mathbb{P}_1(t)$ holds, modes are not evenly dispersed, and upswitches occur at time slots $\tau_1, \tau_2, \dots, \tau_k$ (as in Lemma 1), where $k \leq M - 1$, then restricting the capacity constraint in (1) to these time slots, along with the first time slot, yields the overall optimal solution. However, $\mathbb{P}_1(t + 1)$ need not hold.*

Theorem 3 shows that even with arbitrarily defined modes, the complexity of the winner determination scales only with the number of modes supported and not the number of time slots, as long as the availability projection at t increases monotonically. However, there is no guarantee that capacity projections will continue to be monotonically increasing for future auctions. We can simplify the MKP without \mathbb{P}_1 if bids are uniform quantity.

Theorem 4. *If $\mathbb{P}_1(t)$ does not hold, but bids are of uniform quantity, solving (1) in only the time slots given by Algorithm 1 yields the optimal solution. Further, if all bids are of equal duration, then Algorithm 1 reduces to a single knapsack*

⁶<http://download.skype.com/share/business/guides/skype-connect-requirements-guide.pdf>

Algorithm 1 Pruning time slots under uniform quantity bids ($\mu_{it} = \mu_i$) when $\mathbb{P}_1(t)$ does not hold. \mathbf{W} is set of bids

```

1: procedure COMPUTECONSTRAINEDTIMESLOTS( $\mathbf{W}, \mathbf{C}, t, \phi_{\max}$ )
2:    $slots[0, :] \leftarrow [0]$ 
3:    $consSlotInInterval \leftarrow 0$ 
4:    $consValue \leftarrow 0$ 
5:   for  $\tau \leftarrow 1, \dots, \phi_{\max}$  do
6:      $sumAllAsks \leftarrow 0$ 
7:      $endInterval \leftarrow 0$ 
8:     for  $i \in \mathbf{W}$  do
9:        $sumAllAsks \leftarrow sumAllAsks + \mu_i$ 
10:      if  $\phi_i == (t + \tau)$  then
11:         $endRegion \leftarrow 1$ 
12:       $currSlotConstraint \leftarrow \frac{sumAllAsks}{C_t^{(\tau)}}$ 
13:      if  $currSlotConstraint > consValue$  then
14:         $consValue \leftarrow currSlotConstraint$ 
15:         $consSlotInInterval \leftarrow \tau$ 
16:      if  $endInterval == 1$  then
17:         $slots = [slots, \tau]$ 
18:         $consSlotInInterval \leftarrow 0$ 
19:         $consValue \leftarrow 0$ 
20:   return  $slots$ 

```

problem, solved for the time slot with the largest ratio of requested to available capacity (i.e., $\sum_{i: b_{it} \neq \emptyset} \mu_{it} / C_t^{(1)}$).

Algorithm 1 iterates over each time-slot $\tau \leq t + \phi_{max}$. If a submitted bid(s) is scheduled to finish consumption at τ , it finds the time-slot with the largest ratio of requested to available capacity between τ and the last time-slot when a submitted bid ended. These timeslots are used to solve (1). Hence, if bids are uniform quantity, the dimensionality of (1) scales with the variance in bid durations, not the number of time slots, leading to relatively fast solutions even for large ϕ_{max} .

VI. INCENTIVE COMPATIBILITY

We have shown several ways for the network to simplify the winner determination in (1), making it feasible to optimize for social welfare in real time. This allocation objective, in conjunction with carefully designed payment schemes, can induce strong properties. We first consider a single auction A_t in isolation, and induce a *myopic* notion of truthfulness using the VCG mechanism. We then frame A_t in the context of repeated auctions, where we account for the impact of decisions made in A_t on subsequent auctions. In both cases, we *ensure dominant strategy incentive compatibility while inducing desirable properties that are often challenging to achieve simultaneously, such as revenue monotonicity and ex post individual rationality*.

A. Myopic Truthfulness

The VCG mechanism has gained wide popularity in its ability to guarantee socially optimal results through dominant strategy incentive compatibility (DSIC); i.e., every bidder's best interest is to bid truthfully, regardless of the strategies of other bidders [23]. Since the network maximizes social welfare, it can implement the VCG mechanism by charging auction winners their *social cost*. The social cost of each bidder i

is computed as the difference between the maximum feasible welfare without i and the welfare to others given i 's presence, i.e., $\max_{x_{it} \in [0,1]} \sum_{k=1, k \neq i}^{k=I_t} v_{kt} x_{kt} - \sum_{k=1, k \neq i}^{k=I_t} v_{kt} x_{kt}^*$, where x_{kt}^* represents the optimal solution with i present.

When applied in combinatorial auctions, however, the VCG mechanism is known to exhibit undesirable failures in *bidder revenue monotonicity* [33], meaning the network's revenue from VCG payments may in fact *decrease* when some bids enter the system. An auction is said to be robust for a set of bidders Δ under VCG payments p_{it} if

$$\forall j \in \Delta \sum_{i \in \Delta} p_{it}(v_{it}, \Delta) \geq \sum_{i \in \Delta \setminus \{j\}} p_{it}(v_{it}, \Delta \setminus \{j\}). \quad (2)$$

See Rastegari *et al.* [33] for a more formal treatment of revenue monotonicity. Another type of VCG failure in combinatorial auctions is *goods revenue monotonicity failure* [34], when the operator could increase revenue by not auctioning certain goods (in our case, resource quantities and time slots), hence acquiring an incentive to hide goods from bidders. Most prior work on combinatorial auction frameworks does not address the issue of VCG-induced monotonicity failures, which are especially challenging to manage in MUCA settings such as ours. We show, however, that under certain conditions applying VCG payments is guaranteed to result in revenue monotonicity. To do this, we rely on the property of *bidder submodularity* [33], [35] which builds on the maximum social welfare $V(\Delta)$ of a set of bidders Δ , corresponding to the objective in (1) restricted to Δ . Bidder submodularity holds for bidder sets Δ and Δ' with $\Delta \subseteq \Delta'$ if and only if $\forall i V(\Delta \cup \{i\}) - V(\Delta) \geq V(\Delta' \cup \{i\}) - V(\Delta')$.

Theorem 5. *If (1) can be solved in a single time slot t' (e.g., as in Thm 1), and $\mu_{i't'} = \mu_{i't} \forall i, i' \in [1, I_t]$, then A_t is guaranteed to be revenue monotonic in bidders under VCG payments.*

In this scenario, winner determination is a knapsack problem where bids request the same resources but potentially different valuations. Then, all bids compete equally for capacity, and therefore, removing a bid cannot increase another's social cost, resulting in revenue monotonicity.

Lemma 2. *In our auction, bidder revenue monotonicity implies goods revenue monotonicity.*

In our auction, bidders desire and bid on exactly one bundle, a property referred to as single mindedness. When single-minded bidders exhibit bidder revenue monotonicity, they are goods revenue monotonic as well [34]. Thus, by exploiting the structure of users' real-time resource requests, we have shown that under reasonable conditions, users have an incentive to bid truthfully and the network operator has no incentive to discourage bids from users or hide resources, as doing so will not increase its revenue. However, we also note the following limitation.

Lemma 3. *As long as the auctioneer solves (1) for winner determination of A_t and charges winners their social cost, bidders may have an incentive to submit a false session duration ϕ_{it} .*

Consider a case where $\mathbb{P}_1(t)$ holds and all bids in A_t are uniform quantity. The auctioneer then only solves the knapsack problem in the first time-slot (Thm. 1), and users' choice of ϕ_{it} has no impact on their bid allocations or payments. Indeed, as seen in Section V, even when $\mathbb{P}_1(t)$ does not hold and results in arbitrary up/downswitches, the winner determination and hence payments depend only on the time slots of these switches. Hence, maximizing the social welfare at A_t only with respect to A_t does not directly incentivize truthfulness in declaration of ϕ_{it} .

B. Truthfulness Amidst Temporal Correlations

In our setting, the true social cost of a bit b_{it} is not only a function of other bids submitted to the current auction (as discussed earlier). Selecting a bid b_{it} as a winner of auction $A(t)$ directly reduces available capacity in the next ϕ_{it} time slots, which impacts the bids that can be accepted in subsequent auctions A_j , $j \in [t+1, t+\phi_{it}]$. To account for this *temporal correlation* between these periodic auctions, we are faced with the challenge of factoring in the uncertainty in future bids in admitting the present bids. *We thus develop mechanisms to push this uncertainty either to the user or the network, inducing different properties accordingly.*

The *temporal trickle effect* of a winning bid may in fact extend beyond its duration; for instance, allocating resources for a bid b_{it} might preclude allocation for a bid $b_{j(t+\phi_{it}-1)}$, which might in turn allow for the allocation of a bid at $t+\phi_{it}+1$ that would have been infeasible had j 's bid been allocated. We argue, however, that it is unreasonable to charge users their social cost beyond the duration ϕ_{it} (unlike the treatment by Parkes *et al.* [36]). First, it is extremely challenging to predict and model the trickle effects starting from allocation of a bid until the last auction in the system, leading to significant computational overhead and possible infeasibility. Second, since mobile network use is dense and diversified, the extended effects of a single bid in the system would arguably be too little to cause a significant impact in the overall social welfare and hence not worth accounting for. Hence we propose to hold user i accountable for their "first-order" social cost with respect to arriving bids during $[t, t+\phi_{\max}]$, hence capturing direct impact during i 's consumption and any immediate ripples until $t+\phi_{\max}$.

We now formulate strategies that induce desirable properties despite this temporal correlation amongst auctions and future uncertainty. We first provide definitions of these properties (see [23] for a thorough treatment). **Individual Rationality** is achieved when no bidder receives a negative utility from participating in the auction, i.e., no winning bid is charged more than its reported value v_{it} and no losing bid is charged. The winner determination is **Allocatively Efficient** when social welfare is maximized in the allocation outcomes. If the sum of all payments charged by the auctioneer is non-negative, i.e., the auctioneer does not suffer a net loss, then the mechanism is (weakly) **Budget Balanced**. A property is said to hold **ex ante** if it holds in expectation over the private and unknown information of all bidders, **ex interim** if it holds when a bidder knows their private information but

others only in expectation, and **ex post** if it is guaranteed to hold even when all bidder parameters are revealed. We now develop allocation and payment schemes that navigate trade-offs in these properties by factoring in future bid uncertainty differently.

1) *Maximize Expected Social Welfare*: Let us consider an allocation strategy alternate to (1) to determine winners of auction A_t . Let \mathbb{O}_t be the set of all feasible allocations of A_t . Then

$$o_t^* = \operatorname{argmax}_{\mathbb{O}_t} \sum_{i=1}^{I_t} v_{it}(o_t) + \sum_{j=t+1}^{t+\phi_{\max}} \mathbb{E}_{D(j)}^t \sum_{k=1}^{I_j} [v_{kj}(o_{tj})], \quad (3)$$

where I_j is the number of bidders placing bids in the system at time j and $\mathbb{E}_{D(j)}^t$ denotes the expectation at time t taken over the distribution $D(j)$ of bids at j (consisting of bid arrivals, requested mode in the bid, duration and valuation). With this allocation rule, the network maximizes the expected social welfare of the next ϕ_{\max} time steps in deciding the allocation, rather than maximizing only for the welfare of bidders at A_t . This is implicit in the dependence between o_t and o_{tj} , wherein the latter captures the allocation decision taken at time t for the timestep j in the estimated look-ahead model. Point estimates derived from, for example, Monte Carlo sampling of outcomes starting from t can be used for unknown parameters of future timesteps [37]. The approach in (3) presumes that the network has learned this distribution of bids spanning the next ϕ_{\max} time steps. Indeed, computing a deterministic ϕ_{\max} -step look-ahead model is far more feasible than computing the optimal solution for the multi-stage stochastic programming problem of expected welfare maximization for all remaining auctions [37], [38]. Models requiring computation of the optimal value function at every time-step [36] pose severe feasibility challenges. Further, the time period ϕ_{\max} intuitively lends itself as a reasonable look-ahead period since all user allocations starting at t must end by then, providing a standard and relatively short time window for computing prices.

2) *Charge Expected Social Cost*: We now design payment rules which operate in conjunction with the allocation rule in (3) to induce desirable auction properties. First, we consider a rule similar to above wherein a winning bidder i is charged its expected social welfare cost p_{it} as the difference between maximum welfare without i and welfare to others given i 's presence, i.e.,

$$p_{it} = W_t^{-i} - \left(\sum_{l=1, l \neq i}^{I_t} v_{lt}(o_t^*) + \sum_{j=t+1}^{t+\phi_{\max}} \sum_{k=1}^{I_j} \mathbb{E}_{D(j)}^t [v_{kj}(o_{tj}^*)] \right), \quad (4)$$

where W_t^{-i} represents the welfare without i as

$$W_t^{-i} = \max_{\mathbb{O}_t} \left(\sum_{l=1, l \neq i}^{I_t} v_{lt}(o_t^{-i}) + \sum_{j=t+1}^{t+\phi_{\max}} \sum_{k=1}^{I_j} \mathbb{E}_{D(j)}^t [v_{kj}(o_{tj}^{-i})] \right), \quad (5)$$

o_t^{-i} represents an allocation outcome at time t without i in the system, and o_t^* is the optimal solution with i present, with look-ahead model decisions o_{tj}^* . Note that we assume users have quasi-linear utility functions, as ubiquitously done [23].

Theorem 6. *The mechanism implementing the allocation rule in (3) and the payment rule in (4) is DSIC in all bid parameters. Further, it is ex post individually rational and budget balanced.*

By simply maximizing social welfare in expectation of $[t, t + \phi_{\max}]$ and charging winning bids their expected social cost, the network can not only incentivize dominant strategy truthfulness in η_{it} , ϕ_{it} and v_{it} , but also ensure that no bidders are charged more than what they bid for. We now introduce a new property to evaluate this payment scheme.

Definition 5 (Payment Efficient). *A mechanism is said to exhibit **Payment Efficiency** if it charges winners their social cost. For instance, the VCG mechanism is payment efficient, since winning bids are charged the difference in social welfare to others due to their presence.*

Lemma 4. *The mechanism implementing the allocation rule in (3) and the payment rule in (4) is ex ante allocatively efficient and ex ante payment efficient at t for the time-period $[t, t + \phi_{\max}]$.*

While this mechanism maximizes expected social welfare, it may well be the case that the auctioneer under-predicts demand between $[t, t + \phi_{\max}]$ in retrospect which makes the allocation and payment decisions made earlier suboptimal. The auctioneer's revenue would then be higher if winning bids at t were charged their social cost at $t + \phi_{\max}$ after observing the actual demand.

3) *Charge Realized Social Cost*: The allocation decision for A_t must be made in real-time for immediate session needs. However, we realize that *payments need not be computed or charged in real time for users to start their sessions*. Consider the payment rule given by the actual after-the-fact difference between welfare without and with i after $t + \phi_{\max}$ time steps have elapsed:

$$p_{it}^{t+\phi_{\max}} = \max_{\mathbb{O}_t} \sum_{j=t}^{t+\phi_{\max}} \sum_{k=1, k \neq i}^{I_j} (v_{kj}(o^{-i}) - v_{kj}(o^*)), \quad (6)$$

where o^* is the optimal allocation of users other than i given by

$$o^* = \operatorname{argmax}_{\mathbb{O}_{it}} \sum_{j=t}^{t+\phi_{\max}} \sum_{k=1, k \neq i}^{I_j} v_{kj}(o), \quad (7)$$

and \mathbb{O}_{it} is the set of all feasible allocations given i is allocated. Admitted users are now charged at time $t + \phi_{\max}$, by which time all bids starting consumption at t are guaranteed to end. The auctioneer may now use its retrospective knowledge of bids that came in from t to $t + \phi_{\max}$ to calculate the exact first-order social cost for each bid at t . Before analyzing the unique properties that this rule yields, we first define a modified notion of DSIC.

Definition 6 (DSICE). *Consider bidders that maximize expected utility. If truthful revelation maximizes the expected utility of bidders, regardless of the strategy of other bidders, then the mechanism is said to be **Dominant Strategy Incentive Compatible in Expectation**.*

Theorem 7. *The mechanism implementing the allocation rule in (3) and the payment rule in (6) is DSICE in all bid parameters, ex interim individually rational and ex post budget balanced.*

By charging bidders their true social cost at $t + \phi_{\max}$ based on actual bids that arrived after t , the mechanism essentially *shifts the risk of demand under-prediction* to the bidder. However, winning bidders now bear the risk of being charged more than their bid.

Lemma 5. *The mechanism implementing the allocation rule in (3) and payment rule in (6) is ex ante allocatively efficient and ex post payment efficient with respect to bids at t for $[t, t + \phi_{\max}]$.*

As we allocate using (3), ex ante allocative efficiency holds. By design of (6), winning bids pay their true social cost at $t + \phi_{\max}$ and hence the network is also payment efficient. However, if the auctioneer over-predicts resource demand between $[t + 1, t + \phi_{\max}]$, charging users their actual social cost at $t + \phi_{\max}$ yields less revenue than charging them their expected cost at t .

We have proposed three distinct payment mechanisms for our auction model: traditional VCG (Section VI-A), paying the expected social cost, and paying the realized social cost (Section VI-B). Since the latter two mechanisms require knowledge of future bids, the network would likely introduce VCG payments first. By evaluating each A_t in isolation (i.e., using (1) and VCG payment), the network can learn users' true bid valuations and required bitrates. By estimating bid durations with historical usage, the network may use this distribution of future bid parameters it has learned to offer the latter mechanisms. This would likely increase social welfare in the system, since the network now accounts for the impact of allocation decisions between multiple rounds of auction as in (3). In choosing between the two payment rules as in Table I, the network must make a design choice. It may either guarantee individual rationality by choosing (4) and assume the risk of under-predicting resource demand, or it may choose (6) and ensure payment efficiency while allowing winners to be charged higher than their bids. However, in the latter case, the network now assumes the risk of over-predicting resource demand.

VII. USABILITY CONSTRAINTS

In Sections V and VI, we developed practical allocation and payment strategies for our auctions that achieve spontaneous resource guarantees for real-time applications. We now turn to challenges faced by end users of this system. Most data plans in the US provide known and fixed up-front pricing for the month [39], so engaging spontaneously in auctions may add uncomfortable *expense uncertainty* for the average mobile user. Further, explicitly conveying an app's resource needs and bid parameters every time the user desires guaranteed data access can be a significant deterrent. To address these usability issues, we propose to have automated agents on users' devices that act on their behalf, abstracting them away from resource specification and bidding overhead. We first formulate

a user-parameterized utility framework using which agents can discover user valuations transparently for resource guarantees of specific sessions. We then propose a reinforcement learning strategy to enforce users' daily budgets.

A. Bundle Utility

We determine user i 's valuation of a resource bundle $R_{it} = (\eta_{it}, \phi_{it})$. Let ψ_{it} denote the utility per unit of time consumption for the mode η_{it} requested, normalized between $[0, 1]$ across applications and pre-configured by the user. We use an α -fair model [40] to capture diminishing returns in utility over longer session durations ϕ_{it} . The user's utility associated with request R_{it} is thus given by $U_{it}(R_{it}) = \frac{\phi_{it}^{1-\alpha}}{1-\alpha} \psi_{it}$. If the network responds with an alternate bundle B_{it} that returns a mode below that requested in R_{it} , we impose a penalty to represent the dissatisfaction in receiving a lower mode. We model this penalty for each affected time slot $t + \tau$ as a multiplicative factor $\xi_{i\tau}$, such that a higher penalty corresponds to a smaller mode. We denote $\xi_{i\tau} = 1/(1 + \psi_{it} - \psi_{i(t+\tau)}^*)$, where $\psi_{i(t+\tau)}^*$ denotes user i 's valuation of the mode corresponding to the offered $\mu_{i(t+\tau)}$. Since users may be especially dissatisfied if their session experiences a downswitch, we further apply a downswitch penalty that grows with the magnitude $(a - b)$ of an a/b -downswitch. We model this penalty for each τ as a multiplicative factor $\zeta_{i\tau} = \rho_i/(\mu_{i(t+\tau)} - \mu_{i(t+\tau+1)})$ when $\mu_{i(t+\tau)} > \mu_{i(t+\tau+1)}$, where $\rho_i \in [0, 1]$ is user-specific. The utility of B_{it} relative to that of R_{it} is then given by

$$U_i(B_{it}|R_{it}) = \frac{\phi_{it}^{1-\alpha}}{1-\alpha} \sum_{\tau=1}^{\phi_{it}} \xi_{i\tau} \zeta_{i\tau} \psi_{i(t+\tau)}^*. \quad (8)$$

B. Budget Constraints

Building on (8), we develop an algorithm for the user agent to satisfy a daily budget constraint while placing bids that are proportional to the user's true utility U_i . If agents distribute budgets poorly (as some of the naïve algorithms demonstrated in Section VIII do), users consistently lose in their auctions of interest, hence forming the false impression that the market rate is prohibitively high and exiting the system. The budget distribution problem relies on a policy to select a valuation v_{it} to declare on a given bundle B_{it} (interchangeably R_{it}) that maximizes the user's total expected future utility, subject to the budget constraint. Without loss of generality, we collapse the distinction between B_{it} and R_{it} by considering a virtual round where the network offers $B_{it} = R_{it}$ if R_{it} is perfectly available. To discover the optimal policy, we model the user environment as a Markov Decision Process (MDP) wherein actions correspond to placing bids, and the user receives a reward equal to the utility $U_i(B_{it}|R_{it})$ if the bid wins and zero otherwise. As seen in Figure 6, the state of user i in time slot t is defined as $\sigma_{it} = (t, \beta_{it}, B_{it}, R_{it})$, where β_{it} is the remaining budget at time t (with β_{i1} as the total daily budget). The overall system state determines the probability of winning the bid $P_{\text{win}}(\sigma_{it}, v_{it})$, which also represents the state

Payment Scheme	DSIC	Payment Efficiency	Individual Rationality	Allocative Efficiency	Budget Balance
Expected Social Cost	✓	Ex Ante	Ex Post	Ex Ante	Ex Post
Realized Social Cost	In Expectation	Ex Post	Ex Interim	Ex Ante	Ex Post

TABLE I: We summarize trade-offs between charging expected social cost at t and realized social cost at $t + \phi_{\max}$.

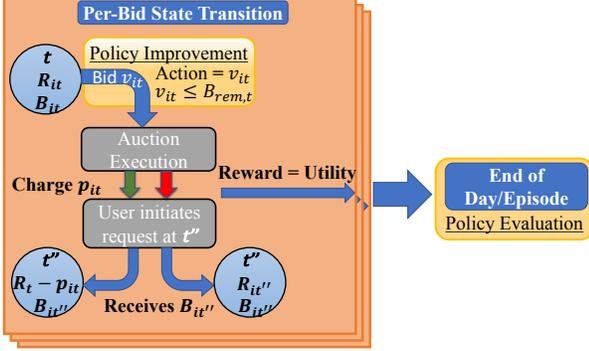


Fig. 6: The MCPI agent bids based on the current policy. Rewards from the states encountered and the actions taken during the day are used to update the policy end of day.

transition probabilities. The optimal budget distribution policy is then given by:

$$\begin{aligned}
& \max_{b_{it}} U_i(B_{it}|R_{it})P_{\text{win}}(\sigma_{it}, v_{it}) + \\
& \sum_{k=t+1}^T U_i(B_{ik}|R_{ik})P_{\text{win}}(\sigma_{ik}, v_{ik}) \\
& \text{s.t.} \sum_{k=t}^T p_{ik}P_{\text{win}}(\sigma_{ik}, v_{ik}) \leq \beta_{it}.
\end{aligned} \tag{9}$$

However, users cannot solve (9) as the environment is only Partially Observable (yielding a POMDP); they only observe their own actions and rewards and therefore cannot compute the transition probabilities $P_{\text{win}}(\sigma_{it}, v_{it})$. We hence employ a model-free reinforcement learning mechanism to determine the optimal user actions under uncertainty. The offline and episodic Monte Carlo Policy Iteration (MCPI) technique [18] is particularly suitable here as users typically exhibit periodicity in daily mobile activities and resource needs, allowing us to consider a day as an episode. MCPI seeks the optimal bidding policy $\pi_{it}^* = v_{it}^*(\sigma_{it})$ by iteratively evaluating a candidate policy π and updating the action value function $q_{\pi}(\sigma_{it}, v_{it})$ from episodes sampled from the POMDP. We define $q_{\pi}(\sigma_{it}, v_{it})$ as the return obtained by placing bid v_{it} in state σ_{it} and then following policy π , averaged over all future states and actions. The return $G_{it} = \sum_{j=t}^T \lambda^j U_i(B_{ij}|R_{ij})$ for a given series of states is the total future discounted reward, where λ is a discount factor representing how much present value a user assigns to future rewards. This captures a degree of uncertainty about the future that stems from the environment as well as the user's estimate of their future session desires. At the end of each episode (e.g., day), the action-value function is updated using

$$q_{\pi}(\sigma_{it}, v_{it}) \leftarrow q_{\pi}(\sigma_{it}, v_{it}) + \chi(G_{it} - q_{\pi}(\sigma_{it}, v_{it})), \tag{10}$$

where χ is the learning rate. We follow the well-known ϵ -greedy approach [18] to balance the trade-off between exploring the environment further to know it better (i.e., choose v_{it} randomly with probability ϵ) and exploiting current knowledge of the environment to maximize current returns (i.e., chooses v_{it} to maximize q_{π}). A bid v_{it} cannot exceed β_{it} , the current available budget. We *anneal* epsilon to eventually always exploit after the environment has been explored sufficiently, which yields the optimal policy if the environment is periodic [18]. We set $\epsilon = 1/N(s)$, where $N(s)$ is the number of times state s is visited, resulting in continuous reduction of exploration from a state as it is visited further and guaranteeing that π_{it} approaches the optimal $v_{it}^*(\sigma_{it})$ as $N(s) \rightarrow \infty$. In Section VIII, we show that this exploration leads to convergence in relatively few iterations and study its performance under increasing complexity.

VIII. EVALUATION OF BUDGET DISTRIBUTION STRATEGY

Setup: We simulate a network of 100 users that participate in auctions for performance guarantees over the course of 80 days and have predetermined schedules of resource requests R_{it} to place during the day. To allow all users to have a reasonable chance of acquiring at least some of their desired resources, we set uniform daily budgets as $\beta_{i1} = \beta$. Note that the MCPI strategy can only help each user achieve the maximum utilities for their given budget and is of little help if market rates are prohibitively high (see Section IX).

Baseline strategies: To assess the performance of MCPI-based budget distribution, we introduce two alternative strategies to compare against. Bidders of the *greedy bidding* strategy bid as much of their remaining budget as needed to achieve the desired utility in each auction as $v_{it} = \min(U_i(B_{it}|R_{it}), \beta_{it})$. Bidders of the *rationed bidding* strategy spread their daily budgets evenly over their (known) daily resource requests, setting v_{it} to the minimum of the bid utility $U_i(B_{it}|R_{it})$ and the session budget. Any residual session budget rolls over to the next session. The maximum realizable daily utility subject to a user's daily budget is

$$U_{i,\max} = \max_{\{y_{it} \in \{0,1\}\}} \sum_{t=1}^T U_i(B_{it}|R_{it})y_{it} \quad \text{s.t.} \sum_{t=1}^T \theta_{it}y_{it} \leq \beta_{i1}, \tag{11}$$

where θ_{it} is the *critical price* of the auction A_t , i.e., the social cost of admitting the bid in this allocation. Since our auctions incentivize truthfulness, θ_{it} is also the payment charged to user i if their bid wins. Note that users themselves cannot compute $U_{i,\max}$ due to partial observability. We use this value as the

theoretical maximum to evaluate the MCPI generated utility against.

Performance of MCPI-based bidding: We first study a deterministic setting where we ensure resources are available for every auction, meaning every request elicits a viable bundle. Four scenarios are simulated, with the first two using the greedy and rationed strategies for all users, respectively. In these cases, we measure the fraction of $U_{i,\max}$ that bidders achieves at the end of each day and show the mean and standard deviation across bidders in Figure 7(a). In the third and fourth scenarios, we introduce one bidder using MCPI amongst the greedy and rationed users, respectively, then measure the MCPI bidder’s U_i . As Figure 7(a) shows, the MCPI bidder succeeds at exploring various actions and reaching 100% $U_{i,\max}$ by day 40 (even sooner against the rationed bidders), while neither of the naïve strategies achieves more than 40%. In the presence of MCPI bidders, greedy and rationed bidders continue to have poor mean performance but marginally higher deviation until the MCPI agent converges (not shown). *Hence, a bidder that previously realizes no more than 40% of the maximum utility achievable with their budget (potentially believing that the market rate is prohibitively high), now wins more by bidding per the MCPI-based budget distribution algorithm.* To increase realism, we next consider a congested setting wherein some resource requests may not elicit any bundles. This is done by increasing session durations, which also increases the likelihood of a resource request being turned away due to ongoing consumption of previously admitted flows. Hence, in addition to budget constraints, MCPI learning must implicitly account for resource availability. For instance, if resources are typically unavailable at 6:00PM, then the optimal strategy might be to distribute the budget to other times of day, since the agent will likely not get a chance to express a bid for resource needs at 6:00PM. Figure 7(b) shows the resulting increase in the time for the strategy to stabilize. The MCPI bidder still outperforms the naïve bidders and converges, but only 85% of $U_{i,\max}$ is reached by day 30. In this case, exploration has a ripple effect on returned bundles in subsequent auctions, which affects budget changes and slows convergence.

We increase uncertainty in the environment by offsetting resource request schedules by a random time period ω_{it} for fraction f of the user population. Hence a user’s resource request times are no longer perfectly periodic. Resource availability as well as critical prices during $[t, t + \phi_{it} + \omega_{it}]$ are affected by this variance, making the MCPI learning more challenging. However, as Figure 7(c) shows, the MCPI bidder’s utility does not decrease as a function of ω_{\max} or f , even as these vary from 1-3 hours and 1-100% respectively, indicating that MCPI is beneficial to deploy in realistic network scenarios. We also study the impact of MCPI bidding on network revenue. Figure 7(d) shows the network revenue *increasing* marginally with the number of MCPI bidders (never decreasing). *This is a direct effect of combining our budget distribution strategy with an auction payment and allocation scheme that incentivize truthfulness.* Users, in their best interest, request resources only when needed and have no incentive to misrepresent their valuation (and no value for leftover budget end of the day).

MCPI bidding then serves only to spend the budget in ways that simultaneously best represents users’ utilities and the chances of winning. More results from our simulation study are provided in our extended report [19].

IX. DISCUSSION

We now discuss practical considerations around deploying our system.

Network Requirements: Our choice of WiFi and LTE for feasibility experiments is motivated by the widespread proliferation of these RATs. The auction model can, in-fact, be executed on any topology where resource availability can be forecasted and reconciled with session needs. Hence, emerging RATs like mmWave (5G) and other WiFi versions are candidate network topologies. The resource modeling mechanism, however, depends on the RAT in use. For instance, in a WiFi network using the MAC-layer point coordination function instead of the distributed coordination function, the access points has more centralized control of flows and therefore may better eliminate wireless externalities. This would likely be factored in the forecast model. In addition, the underlying network must support real-time flow control, e.g., with software-defined networking, to ensure flows do not consume more than their stated bandwidth.

Discovering Session Needs: The app-specific resource needs of a session are best determined by the app itself. Hence, wrapper libraries for network access protocols such as TCP/IP can be used by apps to state the required resources (e.g., bitrate or latency) as they open a new socket connection. The user agent may then be a background process that receives this information. User-specific factors such as intended session duration or daily budget can be explicitly set by the user or estimated by the agent based on historical user activity.

Diversity in Budgets: In Section VI, we propose multiple methods for the network to maximize social welfare and incentivize truthfulness in users’ resource and valuation specifications. This ensures a certain notion of *competitive fairness* [25]. Since it is not in users’ best interest to lie about their needs, only users who truly have the highest value for resources win allocations. However, since this value is expressed in monetary terms, their budgets play a limiting role in their winning chances. In this case, the RL strategy formulated in Section VII will be especially valuable to help more constrained users place bids in auctions with lower critical prices and achieve the best utility possible given their affordability constraints.

X. RELATED WORK

Auctions in wireless networks have been mainly studied in three contexts: spectrum license allocation, secondary cognitive radio allocation, and QoS-aware resource allocation. Auctions for long-term spectrum licenses are held over hours or days with multiple bidding rounds before the auction ends and winners are determined, like the popular simultaneous ascending and combinatorial clock auctions [25], [41]. They do not account for the faster time scales of session-level allocations for spontaneous application sessions. Further, the combinatorial nature of our auctions presents significant challenges

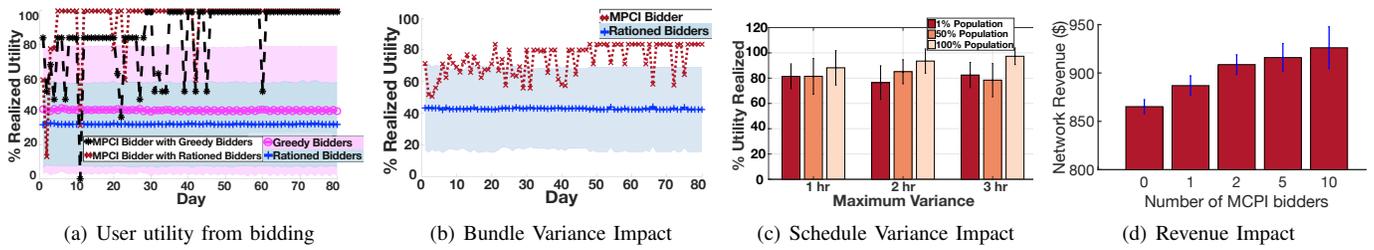


Fig. 7: (a) MCPI far outperforms naïve strategies, achieving 100% of the maximum utility. (b) MCPI realizes 85% utility despite high uncertainty in resource availability. (c) Temporal variance in resource requests does not significantly degrade performance. (d) MCPI bidders increase revenue by driving up critical prices in auctions.

in the context of this prior work. Cognitive radio auctions [42] do not consider session-level app performance, instead availing opportunistic spectrum for much shorter time scales. Auctions have also been employed for QoS-aware real-time channel allocation to primary users in mobile networks. The goals of such approaches [5], [9] differ from ours in their focus on sub-carrier allocation with millisecond granularity and interference mitigation. Using auctions for short-term resource allocation does not guarantee session-level performance, which introduces new combinatorial characteristics that we address.

Our work furthers 5G’s envisioned network slicing capabilities [13], [14]. We verify the premise of RAN slicing for both LTE and WiFi networks, recently also studied by Foukas *et al.* [15] and others [16], [43] in the cellular context, and provide incentive-compatible mechanisms for modeling slices and admitting users to them. We also address budget optimization in the context of repeated auctions, which has been studied in limited settings and even fewer of them combinatorial. Gummadi *et al.* [44] study budget-constrained bidding for sponsored search auctions, but with strong assumptions about the system that guarantee equilibrium. Janssen *et al.* [45] study the combinatorial setting, but their work is limited to the combinatorial clock auction. Almost no work has considered whether reinforcement learning can inform auction bidding strategies as we do here.

XI. CONCLUSION

We design an auction model that captures the market for session-level performance guarantees, decoupling the user-facing auction from network-facing wireless resource management. Our model moves away from radio resource auctions and focuses on application-level provisioning, which is especially useful for emerging real-time multimedia applications. Through trace-driven LTE simulations and extensive WiFi experiments, we verify that not only can wireless externalities be minimized with resource-aware admission control of flows, but more flows can be accommodated by implementing incentive-compatible auction-based admission control. We further analyze the winner determination of our proposed auction model with regard to real-time multimedia applications and show that there are several realistic conditions that render the multi-dimensional knapsack problem solvable in pseudo-polynomial time. These reductions make it feasible to implement the incentive-compatible VCG mechanism and even lead to revenue monotonicity in certain cases. We also analyze the impact of temporal correlations between auctions

on incentive compatibility and define novel payment and allocation schemes to handle future bid uncertainty and navigate trade-offs in desirable properties. Finally, we use the Monte Carlo Policy Iteration technique to show that even budget-constrained users can achieve high utility from these auctions.

XII. ACKNOWLEDGEMENTS

We thank Sriram V. Iyer (Head - Innovation Labs, Myntara-Jabong-Flipkart) and Phani Ramamurthy for their valuable contributions to our experimental setup and insights.

REFERENCES

- [1] Q. Xu, S. Mehrotra, Z. Mao, and J. Li, “Proteus: Network Performance Forecast for Real-time, Interactive Mobile Applications,” in *Proceeding of the 11th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys ’13. Taipei, Taiwan: ACM, 2013, pp. 347–360.
- [2] M. Jarschel, D. Schlosser, S. Scheuring, and T. Hoßfeld, “Gaming in the clouds: QoE and the users perspective,” *Mathematical and Computer Modelling*, vol. 57, no. 11-12, pp. 2883–2894, 2013.
- [3] S. Baraković and L. Skorin-Kapov, “Survey and Challenges of QoE Management Issues in Wireless Networks,” *Journal of Computer Networks and Communications*, vol. 2013, 2013.
- [4] W. Miao, G. Min, Y. Jiang, X. Jin, and H. Wang, “QoS-aware resource allocation for LTE-A systems with carrier aggregation,” in *2014 IEEE Wireless Communications and Networking Conference (WCNC)*, April 2014, pp. 1403–1408.
- [5] M. S. ElBamby and K. M. F. Elsayed, “An auction approach to resource allocation with interference coordination in LTE-A systems,” in *2014 IEEE Wireless Communications and Networking Conference (WCNC)*, Istanbul, Turkey, April 2014, pp. 1885–1890.
- [6] X. K. Zou, J. Erman, V. Gopalakrishnan, E. Halepovic, R. Jana, X. Jin, J. Rexford, and R. K. Sinha, “Can accurate predictions improve video streaming in cellular networks?” in *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*. ACM, 2015, pp. 57–62.
- [7] T. Stockhammer, “Dynamic Adaptive Streaming over HTTP –: Standards and Design Principles,” in *Proceedings of the Second Annual ACM Conference on Multimedia Systems*, ser. MMSys ’11. San Jose, CA, USA: ACM, 2011, pp. 133–144.
- [8] M. Xiao, C. Zhou, V. Swaminathan, Y. Liu, and S. Chen, “Bas-360: Exploring spatial and temporal adaptability in 360-degree videos over http/2,” *IEEE*, 2018.
- [9] K. Yang, N. Prasad, and X. Wang, “An Auction Approach to Resource Allocation in Uplink OFDMA Systems,” *IEEE Transactions on Signal Processing*, vol. 57, no. 11, pp. 4482–4496, Nov 2009.
- [10] T. Wolf, J. Griffioen, K. L. Calvert, R. Dutta, G. N. Rouskas, I. Baldin, and A. Nagurney, “ChoiceNet: Toward an Economy Plane for the Internet,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 58–65, jul 2014.
- [11] Y. Peng, D.-K. Kang, F. Al-Hazemi, and C.-H. Youn, “Energy and QoS Aware Resource Allocation for Heterogeneous Sustainable Cloud Datacenters,” *Opt. Switch. Netw.*, vol. 23, no. P3, pp. 225–240, Jan. 2017.
- [12] A. Karamoozian, A. Hafid, M. Boushaba, and M. Afzali, “QoS-aware resource allocation for mobile media services in cloud environment,” in *2016 13th IEEE Annual Consumer Communications Networking Conference (CCNC)*, Jan 2016, pp. 732–737.

- [13] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, "Network Slicing in 5G: Survey and Challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94–100, May 2017.
- [14] M. Harishankar, P. Tague, and C. Joe-Wong, "Network Slicing as an Ad-Hoc Service: Opportunities and Challenges in Enabling User-Driven Resource Management in 5G," in *Proceedings of 1st International Workshop on Trustworthy and Real-time Edge Computing for Cyber-Physical Systems (TREC4CPS)*. Nashville, TN, USA: Institute for Software Integrated Systems, Vanderbilt University, Dec 2018.
- [15] X. Foukas, M. K. Marina, and K. Kontovasilis, "Orion: RAN Slicing for a Flexible and Cost-Effective Multi-Service Mobile Network Architecture," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '17. Snowbird, Utah, USA: ACM, 2017, pp. 127–140.
- [16] C. Marquez, M. Gramaglia, M. Fiore, A. Banchs, and X. Costa-Perez, "How Should I Slice My Network?: A Multi-Service Empirical Evaluation of Resource Sharing Efficiency," in *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, ser. MobiCom '18. New Delhi, India: ACM, 2018, pp. 191–206.
- [17] S. Ha, S. Sen, C. Joe-Wong, Y. Im, and M. Chiang, "Tube: Time-dependent pricing for mobile data," in *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, ser. SIGCOMM '12. Helsinki, Finland: ACM, 2012, pp. 247–258.
- [18] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.
- [19] M. Harishankar, S. Pilaka, P. Sharma, N. Srinivasan, C. Joe-Wong, and P. Tague, "Procuring Spontaneous Session-level Resource Guarantees or Real-time Applications: An Auction Approach," *tech. rep.*, 2018. [Online]. Available: <https://www.dropbox.com/s/coib0cgvjgw9e2/JSACExtendedPaper.pdf>.
- [20] F. Capozzi, G. Piro, L. A. Grieco, G. Boggia, and P. Camarda, "Downlink Packet Scheduling in LTE Cellular Networks: Key Design Issues and a Survey," *IEEE Communications Surveys Tutorials*, vol. 15, no. 2, pp. 678–700, Second 2013.
- [21] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren, "Usage Patterns in an Urban WiFi Network," *IEEE/ACM Transactions on Networking*, vol. 18, no. 5, pp. 1359–1372, Oct 2010.
- [22] K. Fukuda, H. Asai, and K. Nagami, "Tracking the Evolution and Diversity in Network Usage of Smartphones," in *Proceedings of the 2015 Internet Measurement Conference*, ser. IMC '15. Tokyo, Japan: ACM, 2015, pp. 253–266.
- [23] V. Krishna, *Auction Theory*. Academic press, 2009.
- [24] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter Mechanism Design and Sequential Posted Pricing," in *Proceedings of the Forty-second ACM Symposium on Theory of Computing*, ser. STOC '10. Cambridge, Massachusetts, USA: ACM, 2010, pp. 311–320.
- [25] P. Cramton, "Spectrum Auction Design," *Review of Industrial Organization*, vol. 42, no. 2, pp. 161–190, 2013.
- [26] A. Freville, "The multidimensional 0-1 knapsack problem: An overview," *European Journal of Operational Research*, vol. 155, no. 1, pp. 1–21, 2004.
- [27] P. Krysta, O. Telelis, and C. Ventre, "Mechanisms for Multi-unit Combinatorial Auctions with a Few Distinct Goods," in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, ser. AAMAS '13. St. Paul, MN, USA: International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 691–698.
- [28] T. Kelly, "Generalized Knapsack Solvers for Multi-unit Combinatorial Auctions: Analysis and Application to Computational Resource Allocation," in *Proceedings of the 6th AAMAS International Conference on Agent-Mediated Electronic Commerce: Theories for and Engineering of Distributed Mechanisms and Systems*, ser. AAMAS'04. New York, NY: Springer-Verlag, 2005, pp. 73–86.
- [29] A. Giovannucci, J. A. Rodriguez-Aguilar, J. Cerquides, and U. Endriss, "Winner Determination for Mixed Multi-unit Combinatorial Auctions via Petri Nets," in *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '07. Honolulu, Hawaii: ACM, 2007, pp. 104:1–104:8.
- [30] M. J. Varnamkhandi, "Overview of the algorithms for solving the multidimensional knapsack problems," *Advanced Studies in Biology*, vol. 4, no. 1, pp. 37–47, 2012.
- [31] V. M. Copping, V. Smith, and J. A. Titus, "Incentives and Behavior in English, Dutch and Sealed-Bid Auctions," *Economic Inquiry*, vol. 18, no. 1, pp. 1–22, 1980.
- [32] D. Pisinger, "Where Are the Hard Knapsack Problems?" *Comput. Oper. Res.*, vol. 32, no. 9, pp. 2271–2284, Sep. 2005.
- [33] B. Rastegari, A. Condon, and K. Leyton-Brown, "Revenue Monotonicity in Combinatorial Auctions," *SIGecom Exch.*, vol. 7, no. 1, pp. 45–47, dec 2007.
- [34] N. Muto and S. Yasuhiro, "Goods Revenue Monotonicity in Combinatorial Auctions," Graduate School of Economics, Hitotsubashi University, Discussion Papers 2013-13, Oct. 2013.
- [35] L. M. Ausubel and P. R. Milgrom, "Ascending auctions with package bidding," *Advances in Theoretical Economics*, vol. 1, no. 1, 2002.
- [36] D. C. Parkes and S. Singh, "An MDP-based Approach to Online Mechanism Design," in *Proceedings of the 16th International Conference on Neural Information Processing Systems*, ser. NIPS'03. Whistler, British Columbia, Canada: MIT Press, 2003, pp. 791–798.
- [37] W. B. Powell, "Clearing the Jungle of Stochastic Optimization," in *Bridging data and decisions*. Informs, 2014, pp. 109–137.
- [38] M. Römer and T. Mellouli, "Future Demand Uncertainty In Personnel Scheduling: Investigating Deterministic Lookahead Policies Using Optimization And Simulation," in *ECMS*, Regensburg, Germany, 2016, pp. 502–507.
- [39] L. Zheng, C. Joe-Wong, M. Andrews, and M. Chiang, "Optimizing data plans: Usage dynamics in mobile data networks," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. Atlanta, GA, USA: IEEE, 2018, pp. 2474–2482.
- [40] L. Zheng, C. Joe-Wong, J. Chen, C. G. Brinton, C. W. Tan, and M. Chiang, "Economic viability of a virtual ISP," in *IEEE INFOCOM 2017 - IEEE Conference on Computer Communications*, Atlanta, GA, USA, May 2017, pp. 1–9.
- [41] P. Cramton, "Simultaneous Ascending Auctions," *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [42] Y. Zhang, D. Niyato, P. Wang, and E. Hossain, "Auction-based Resource Allocation in Cognitive Radio Systems," *IEEE Communications Magazine*, vol. 50, no. 11, pp. 108–120, November 2012.
- [43] V. Petrov, M. A. Lema, M. Gapeyenko, K. Antonakoglou, D. Moltchanov, F. Sardis, A. Samuylov, S. Andreev, Y. Koucheryavy, and M. Dohler, "Achieving End-to-End Reliability of Mission-Critical Traffic in Softwarized 5G Networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 3, pp. 485–501, March 2018.
- [44] R. Gummadi, P. Key, and A. Proutiere, "Repeated Auctions under Budget Constraints: Optimal Bidding Strategies and Equilibria," in *the Eighth Ad Auction Workshop*, 2012.
- [45] M. Janssen, V. Karamychev, and B. Kasberger, "Budget Constraints in Combinatorial Clock Auctions."



Madhumitha Harishankar is a PhD Candidate at the Department of Electrical and Computer Engineering, Carnegie Mellon University, Silicon Valley. She is advised by Prof. Patrick Tague as a member of the MEWS research group. She received her B.S. in Electrical and Computer Engineering at Rutgers University (summa cum laude, James Slade Scholars, Honors program) and was a software engineer at Barclays Investment Bank and then at Amazon Web Services, before joining CMU. Her research interests include wireless networking and network economics, specifically in applying pricing and learning techniques for better resource allocation in data networks. She received an NSF GRFP Honorable Mention in 2017, as well as CIT Dean's 1-year fellowship in 2016. Her more recent research interests are in blockchain economics.



Sireesha Pilaka Sireesha Pilaka has received the MS degree from Information Networking Institute, Carnegie Mellon University in 2017 and B.Tech. degree from Indian Institute of Technology, Kharagpur in 2014. She was a research assistant at Carnegie Mellon University in 2017. She is currently a Sr. Network Engineer at Visa's network architecture team. Her research interests include wide area networks, data-center networks, SDN/NFV, network automation and machine learning.



Patrick Tague is an Associate Research Professor at Carnegie Mellon University with appointments in the Electrical and Computer Engineering Department and the Information Networking Institute, and he is also the Associate Director of the INI. Patrick leads the Mobile, Embedded, and Wireless Security group at the Silicon Valley Campus of CMU. His research interests include wireless communications and networking; wireless/mobile security and privacy; robust and resilient networked systems; and analysis and sense-making of sensor network data.

He received PhD and MS degrees in Electrical Engineering from the University of Washington as a member of the Network Security Lab and BS degrees in Mathematics and Computer Engineering from the University of Minnesota. Patrick received the Yang Research Award for outstanding graduate research in the UW Electrical Engineering Department, the Outstanding Graduate Research Award from the UW Center for Information Assurance and Cybersecurity, and the NSF CAREER award. He is a member of IEEE, including the Communication, Computer, and Signal Processing Societies, and ACM, including SIGCOMM, SIGMobile, and SIGSAC.



Pragya Sharma received her Bachelors in Technology (B.Tech) from the Department of Electronics and Communication, Vellore Institute of Technology, India in 2017 and her M.S. in Electrical and Computer Engineering from Carnegie Mellon University in 2018. She is currently an Embedded Software Engineer at Amazon Lab126. Her research interests include wireless communications, low-power sensor networks and embedded systems.



Nagarjun Srinivasan received his Bachelor in Technology (B.Tech) in Electrical Engineering from the Indian Institute of Technology (IIT), Hyderabad in 2016 and Master of Science (MS) in Electrical and Computer Engineering from Carnegie Mellon University in 2017. He currently works as a Developer at Mist Systems, a Juniper Company, on building real-time, reliable and scalable Indoor Location Based Services. His research interests include Signal Processing and Filtering, Indoor Location Technologies, Optimization and Reinforcement Learning.



Carlee Joe-Wong is an Assistant Professor of Electrical and Computer Engineering at Carnegie Mellon University. She received her A.B. degree (magna cum laude) in Mathematics, and M.A. and Ph.D. degrees in Applied and Computational Mathematics, from Princeton University in 2011, 2013, and 2016, respectively. Dr. Joe-Wong's research is in optimizing networked systems, particularly on applying machine learning and pricing to resource allocation in data and computing networks. From 2013 to 2014, she was the Director of Advanced Research

at DataMi, a startup she co-founded from her Ph.D. research on mobile data pricing. Her research has received several awards, including the ARO Young Investigator Award in 2019, the NSF CAREER Award in 2018, the INFORMS ISS Design Science Award in 2014 and the Best Paper Award at IEEE INFOCOM 2012.