



HAL
open science

Sequential Resource Access: Theory and Algorithm

Lin Chen, Anastasios Giovanidis, Wei Wang, Lin Shan

► **To cite this version:**

Lin Chen, Anastasios Giovanidis, Wei Wang, Lin Shan. Sequential Resource Access: Theory and Algorithm. IEEE International Conference on Computer Communications (INFOCOM 2021), May 2021, Vancouver (virtual), Canada. hal-03043691

HAL Id: hal-03043691

<https://hal.science/hal-03043691>

Submitted on 23 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sequential Resource Access: Theory and Algorithm

Lin Chen^{*¶}, Anastasios Giovanidis[†], Wei Wang[‡], Lin Shan[§]

^{*}School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China, chenlin69@mail.sysu.edu.cn

[†]Sorbonne University, CNRS-LIP6, Paris, France, Anastasios.Giovanidis@lip6.fr

[‡]College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou, China, wangw@zju.edu.cn

[§]Dept. of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY, USA, shan.x.lin@stonybrook.edu

[¶]Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

Abstract—We formulate and analyze a generic sequential resource access problem arising in a variety of engineering fields, where a user disposes a number of heterogeneous computing, communication, or storage resources, each characterized by the probability of successfully executing the user’s task and the related access delay and cost, and seeks an optimal access strategy to maximize her utility within a given time horizon, defined as the expected reward minus the access cost. We develop an algorithmic framework on the (near-)optimal sequential resource access strategy. We first prove that the problem of finding an optimal strategy is NP-hard in general. Given the hardness result, we present a greedy strategy implementable in linear time, and establish the closed-form sufficient condition for its optimality. We then develop a series of polynomial-time approximation algorithms achieving (ϵ, δ) -optimality, with the key component being a pruning process eliminating dominated strategies and, thus maintaining polynomial time and space overhead.

I. INTRODUCTION

We consider the following generic resource access problem: a user needs to execute a communication or computing task; there are a set of resources she may access; by accessing resource i , she can successfully execute her task with probability p_i and in that case obtains a unit reward; accessing resource i incurs a cost c_i and delay d_i ; the user seeks a *sequential* resource access strategy maximizing her expected utility, defined as the expected reward minus the cost, within a time horizon T corresponding to the maximal delay she disposes to complete her task.

The above resource access problem arises in a variety of engineering fields, where a decision maker disposes a number of heterogeneous computing, communication, or storage resources, and needs to find an optimal strategy to access them to execute her task. Examples fitting into this formulation include communication and computing task offloading, cached data access, opportunistic forwarding, and user-centric network selection (cf. Section II for a detailed description). Theoretically, the intrinsic problem structure we attempt to capture in our formulation is the sequential selection of resources with a hard time constraint. We believe that such formulation, despite being generic, can provide valuable insights in many emerging networking, communication, and computing scenarios.

Furthermore, application examples of our formulation extend well beyond the realm of communication and networking, and include a class of sequential searching problems that can be formulated by the following intuitive example. Given a number of places, each having a certain probability of

containing a prize (e.g., a fugitive actively searched by the police), an agent aims at finding the prize by searching the places. Searching a place incurs certain cost and delay. The agent seeks an optimal sequence of places to search in order to maximize the overall probability of finding the prize while limiting the searching cost within a given time horizon.

Motivated by the above observation, we embark in this paper on a systematic analysis of the sequential resource access problem in its generic form. By generic we mean that no specific problem or context is assumed for the analysis. The only important assumption, made mainly for mathematical tractability, is the independence of success probability among resources. While limiting the applicability of our analysis to some extent, this assumption is justified in many practical situations, where resources are independent one to the other. In case where resources are correlated among them, our model can be extended by taking into the account such inter-resource correlation by e.g., forming super-resources representing correlated ones. We leave a detailed analysis of the correlated case for future research. Despite the generic nature of our analysis, the algorithms we develop are readily implementable once instantiated with system parameters.

Our major technical contribution in this paper is a comprehensive algorithmic framework deriving the optimal or near-optimal sequential resource access strategy. To this end, we first analyze the homogeneous case where the access time d_i is the same among resources. We present a greedy strategy that can be implemented in linear time and mathematically establish sufficient conditions for its optimality. We then develop a polynomial-time algorithm that gives an optimal strategy and design a preprocessing procedure to further reduce its complexity. We then turn to the heterogeneous case where the access time may be different among resources. We prove that the problem of finding an optimal strategy is NP-hard. Given its hardness, we develop a series of polynomial-time approximation algorithms approaching an optimum solution in the sense of (ϵ, δ) -optimality.

The paper is organized as follows. In Section II we formalize the sequential resource access strategy and present some structural properties. In Sections III and IV, we develop a set of greedy and approximation strategies for the homogeneous and heterogeneous cases, respectively. Section V presents the simulation results. We review related work in Section VI. Section VII concludes the paper.

II. MODEL AND PROBLEM FORMULATION

We consider the sequential resource access problem formulated in the Introduction. The user disposes a set \mathcal{N} of N resources; each resource i is characterized by a triple (p_i, c_i, d_i) , known to the user, corresponding to the success probability, access cost, and access delay of the resource. Our goal is to design an optimal resource access strategy, concisely termed as *strategy*, that maximizes the utility of the user within a time horizon T .

A. Assumptions

The formulation of our problem hinges on the following two assumptions.

First, resources are mutually independent. This assumption is justified in many practical situations. For example, resources may map to distinct communication channels where the probability of accomplishing a transmission in one channel is independent to that of the other; resources may also map to a set of hard disks where their failure probabilities are mutually independent. In case of correlated resources, our model can be extended by taking into the account such inter-resource correlation by, e.g., forming super-resources representing correlated ones.

Second, resources are selected without replacement. That is, we focus on the case where the probability p_i of each resource π_i represents the uncertainty about its availability, which is revealed once π_i is accessed. This assumption is valid when the environment is relatively stable during the access process, e.g., in the slow fading case, where each resource maps to a wireless channel, with the channel coherence time longer than the time horizon T .

B. Sequential Resource Access Problem Formulation

Mathematically, a strategy π can be defined as an ordered set of resources $\{\pi_1, \dots, \pi_n\}$, $\pi_i \in \mathcal{N}$ for $1 \leq i \leq n$, that the user accesses sequentially. We denote the number of resources accessed in π by n . As resources are selected without replacement, we have $n \leq N$. We are particularly interested in the case where n is significantly smaller than N , thus requiring a carefully designed access strategy.

We define the following function $U(\pi)$ to denote the user's utility as a function of her strategy π .

$$U(\pi) \triangleq R(\pi) - C(\pi).$$

Specifically, $U(\pi)$ is defined as the expected reward minus the cost. $R(\pi) \triangleq 1 - \prod_{i=1}^n (1 - p_{\pi_i})$ denotes the expected normalized reward, where the user gets a unit reward (e.g., one euro) if she successfully executes her task before the deadline.

$$C(\pi) \triangleq \sum_{i=1}^n \left[\prod_{j=1}^{i-1} (1 - p_{\pi_j}) \right] p_{\pi_i} \left(\sum_{k=1}^i c_{\pi_k} \right) + \left[\prod_{j=1}^n (1 - p_{\pi_j}) \right] \left(\sum_{k=1}^n c_{\pi_k} \right)$$

denotes the expected normalized cost (e.g., also in euro), where $\left[\prod_{j=1}^{i-1} (1 - p_{\pi_j}) \right] p_{\pi_i}$ is the probability that the user

Section II: model and problem formulation	
\mathcal{N}	Set of accessible resources, $N = \mathcal{N} $
p_i	Success prob. of resource i
q_i	$1 - p_i$
c_i	Access cost of resource i
d_i	Access delay of resource i
T	Maximal delay bound
π	Resource access strategy, $\pi = (\pi_1, \dots, \pi_n)$
n	Number of resources in π , $n = \pi $
$d(\pi)$	Access delay of strategy π , $d(\pi) = \sum_{i=1}^n d_{\pi_i}$
π^*	Optimal feasible resource access strategy
Π	Set of all resource access strategies
Π_f	Set of all feasible resource access strategies
$U(\pi)$	Utility function
$V(\pi)$	Pseudo-cost function, $V(\pi) = 1 - U(\pi)$
Section III: the homogeneous case	
y_i	Rank of q_i among all q 's
z_i	Rank of c_i among all c 's
$\hat{\pi}$	Resource access strategy output by our algorithms
Section IV: the heterogeneous case	
c_{min}	$\min_{i \in \mathcal{N}} c_i$
d_{min}	$\min_{i \in \mathcal{N}} d_i$, normalized to 1
\hat{d}_i	Discretized access delay
\hat{V}	Discretized pseudo-cost
$\hat{d}(\pi)$	Discretized access delay of π , $\hat{d}(\pi) = \sum_{i=1}^{ \pi } \hat{d}_{\pi_i}$
λ	Scaling parameter corresponding to access time
μ	Scaling parameter corresponding to pseudo-cost

TABLE I. Main notations

fails in executing her task with the resources π_1 to π_{i-1} and succeeds with resource i with $\sum_{k=1}^i c_{\pi_k}$ being the related cost, $\prod_{j=1}^n (1 - p_{\pi_j})$ is the probability of failure with all the n resources in π with $\sum_{k=1}^n c_{\pi_k}$ being the related cost.

For ease of presentation, we define a pseudo-cost function

$$V(\pi) \triangleq 1 - U(\pi).$$

The problem of maximizing the utility function $U(\pi)$ maps to minimizing the pseudo-cost function $V(\pi)$. We next give the definition of feasible strategy and formulate the sequential resource access problem.

Definition 1 (Feasible strategy). *For any strategy $\pi \triangleq \{\pi_i\}_{i=1}^n$, let $d(\pi) \triangleq \sum_{i=1}^n d_{\pi_i}$ denote the access delay of π . We call π feasible if its access delay does not exceeds T , i.e., $d(\pi) \leq T$. We denote Π_f the set of all feasible strategies.*

Example 1. *The user disposes 100 resources, half with access delay $d_i = 1$, termed as type-I resources, half with $d_i = 2$, termed as type-II resources, and finite-horizon $T = 3$. The feasible strategy set Π_f consists of the following categories of strategies: (1) accessing 1, 2, or 3 type-I resource(s), (2) accessing 1 type-II resource, (3) accessing 1 resource per type.*

Problem 1 (Sequential resource access). *The sequential resource access problem seeks a feasible strategy π minimizing*

the pseudo-cost function $V(\pi)$, i.e., $\min_{\pi \in \Pi_f} V(\pi)$.

Our problem is by nature a combinatorial optimization problem. Consider a degenerated case with zero access cost. The problem becomes purely combinatorial and can be algebraically formulated as the following Knapsack problem.

$$\min_{\pi} \sum_{i \in \pi} \log(1 - p_i),$$

subject to the Knapsack constraint $\min_{\pi} \sum_{i \in \pi} d_i \leq T$. The order of accessing resources does not matter in this case. On the other hand, in the generic setting, the user may tend to access more reliable resources with high p_i 's; these resources may also incur larger access delay and cost; therefore, she needs to strike a balance among the success probability, access cost, and delay. Moreover, the order of the accessed resources is also important and needs to be optimized.

C. Applicability of Our Problem Formulation

Our generic formulation of the sequential resource access problem is readily applicable in a wide range of resource access problems in emerging communication and computing applications. Below we give three concrete examples.

Data Access in Network Caching. Caching is widely deployed in emerging networking systems such as 5G edge and in-network caches [7], content delivery networks (CDN) [5]. In network caching systems, a cache can be regarded as a data store holding a subset of data that may be accessed by users. Intuitively, knowing which item is stored in each data store can significantly improve user experience. However, maintaining such information may be too expensive. As an alternative solution, it is more practical to maintain an approximate catalog of items at each data store based on compact data structures such as Bloom filters [6]. The price to pay for the space compactness is the well-known false positive, where a Bloom filter returns a positive response while the corresponding data item is not stored at the data store. The false positive rate of a Bloom filter of size m and k hash functions storing n data items is $[1 - (1 - 1/m)^{kn}]^k$, approximately $(1 - e^{-kn/m})^k$.

Now consider the the situation where a user needs to fetch a data item. There are a number of candidate data stores returning positive responses after checking the corresponding Bloom filters. Accessing each data store i incurs a cost c_i and delay d_i . The user disposes time T to fetch the data item. She gets a reward r_1 if successfully fetching it by the deadline, and pays a penalty r_2 otherwise. We can formulate two optimization problems faced by the user. In the first problem, investigated in [11], the user can access multiple data stores and seeks an optimal set of data stores to access simultaneously. The second problem captures the situation where the user is limited to access one cache each time, but can perform multiple queries sequentially as long as the total delay does not exceed T . Compared to accessing multiple caches simultaneously, sequential access can reduce the user's total access cost. The sequential data access problem can be formulated by our sequential resource access problem by

mapping the false positive rate of the Bloom filter of cache i to $1 - p_i$, and normalizing the user's reward, penalty, and cost.

Interest Forwarding in Information Centric Networks (ICN). In ICNs [18], a content is typically divided into chunks. Each chunk is addressed by a unique ID and may have many identical cached copies in the ICN routers across the Internet. A chunk is located and requested by forwarding the so-called *interests*. A user can forward her interest to one or more neighbor ICN routers. If there is no bandwidth or other cost limitation, the user can forward her interest to all available neighbors. However, if there is a bandwidth limitation, or the user has to pay for the interest or delivered content, then she needs to carefully choose which neighbors to forward her interest and in what order, rather than simple flooding [4]. By mapping the neighbor set to \mathcal{N} , the probability that neighbor i can return a copy of the searched chunk to p_i , the average response delay and the access cost to d_i and c_i respectively, the ICN interest forwarding problem can be cast to the sequential resource access problem.

Opportunistic Packet Forwarding. In opportunistic packet forwarding, a user needs to decide the sequence of invoking a set of potential forwarders to transmit a data packet. Each forwarder i has a certain probability p_i of successfully executing the transmission task depending on its channel condition, which is assumed to be stable during the considered time horizon. Invoking forwarder i incurs a cost c_i (e.g., in terms of energy consumption) and delay d_i . The user seeks an optimal sequential strategy to invoke a subset of forwarders to maximize the packet delivery rate within the tolerable delay T by taking into account the related cost. This optimal opportunistic forwarding problem also fits in our formulation.

D. Structural Properties of Optimal Feasible Strategy

We conclude this section by showing the following structural properties of $V(\pi)$ and any optimal feasible strategy π^* . To streamline our presentation and due to page limit, readers are referred to the anonymous technical report [1] for the proof of all the lemmas and theorems.

Lemma 1 (Structural properties of $V(\pi)$). *Define $q_i \triangleq 1 - p_i$, $\forall i \in \mathcal{N}$. The following properties hold¹:*

- 1) $V(\pi) = \prod_{i=1}^n q_{\pi_i} + \sum_{i=1}^n \left(\prod_{j=1}^{i-1} q_{\pi_j} \right) c_{\pi_i}$;
- 2) If $q_{\pi_i} + c_{\pi_i} \leq 1$ for $1 \leq i \leq n$, then $V(\pi) \leq 1$;
- 3) Given any strategy π , if there exists a resource $\pi_k \in \pi$ such that $q_{\pi_k} + c_{\pi_k} > 1$, it holds that $V(\pi) > V(\pi_{-k})$, where π_{-k} denotes the strategy by removing π_k from π .

The first property of Lemma 1 demonstrates that $V(\pi)$ consists of two terms: (1) the expected normalized loss due to failing to execute the task $\prod_{i=1}^n q_{\pi_i}$, (2) the expected aggregated access cost $\sum_{i=1}^n \left(\prod_{j=1}^{i-1} q_{\pi_j} \right) c_{\pi_i}$. The second property can be explained intuitively. Upon accessing each resource π_i , the user gets an expected reward p_{π_i} , with an incurred access

¹To make the notation concise, we denote $\prod_{j=1}^i q_{\pi_j} = 1$ for $i < 1$.

cost c_{π_i} . If the expected reward outweighs the cost for each resource π_i , i.e., $p_{\pi_i} \geq c_{\pi_i}$, the global utility $U(\pi)$ is logically non-negative. It follows that $V(\pi) = 1 - U(\pi) \leq 1$. Note that $p_{\pi_i} = 1 - q_{\pi_i}$, a sufficient condition to achieve $V(\pi) \leq 1$ is $q_{\pi_i} + c_{\pi_i} \leq 1$, as stated in the second part of Lemma 1. The third property demonstrates that any rational user never accesses any resource i where $q_i + c_i > 1$, as simply removing it improves the performance. We can thus safely focus on the case where $q_i + c_i \leq 1$ holds for $1 \leq i \leq N$.

Lemma 2 (Structural properties of optimal strategy π^*). *Let π^* denote an optimal strategy, it holds that $p_{\pi_i^*}/c_{\pi_i^*}$ is non-increasing in i , i.e., $p_{\pi_i^*}/c_{\pi_i^*} \geq p_{\pi_{i+1}^*}/c_{\pi_{i+1}^*}$, $1 \leq i \leq |\pi^*| - 1$.*

Lemma 2 demonstrates that, once the set of resources to access at π^* is determined, it suffices to access them in decreasing order of p_i/c_i . In other words, our problem can be transformed to finding an optimal set of resources to access, which is by nature a combinatorial optimization problem.

III. THE HOMOGENEOUS CASE

This section focuses on the homogeneous case where the access delay d_i is identical among resources. To make the notation concise without losing generality, we set $d_i = 1, \forall i \in \mathcal{N}$, i.e., the user can access at most T resources. Our motivation of starting with the homogeneous case is to tackle the problem in a progressive way. As we will demonstrate in this section, analyzing the homogeneous case allows us to obtain more insights on the structure of an optimal strategy, which are useful in the study of the heterogeneous case. We emphasize that the homogeneity only concerns the access delay. Other parameters, such as success probability and access cost, may still be heterogeneous among resources.

A. A Greedy Strategy Based on p_i/c_i

We first investigate a greedy strategy sequentially accessing T resources² in the decreasing order of p_i/c_i , as formalized in Definition 2. Our greedy strategy is motivated by the structural property of an optimal strategy in Lemma 2. From an economic angle, the greedy strategy sequentially chooses T resources by decreasing ratio of success probability (representing profit in certain sense) to access cost. To make our analysis concise, we assume that for any pair of resources i and j , $c_i \neq c_j$.³

Definition 2 (Greedy strategy). *The greedy strategy consists of sequentially accessing T resources by decreasing ratio p_i/c_i , in case of tie choosing the resources by increasing c_i . Mathematically, by sorting the resources such that for each $1 \leq i \leq N - 1$ either (1) $p_i/c_i > p_{i+1}/c_{i+1}$ or (2) $p_i/c_i = p_{i+1}/c_{i+1}$ and $c_i < c_{i+1}$, the greedy strategy sequentially accesses the first T resources.*

²By sequentially accessing a set of resources, we mean by accessing sequentially those resources until a success or the end of time horizon T .

³In case of tie where $c_i = c_j$, we add a small quantity ϵ to either c_i or c_j to break the tie, which leads to at most ϵ in the utility of the optimal strategy. Our analysis can thus be extended in the generic case.

In Theorem 1, we establish the sufficient condition under which the greedy strategy is optimal. In the sequel analysis we assume that the resources are sorted according to Definition 2.

Theorem 1 (Optimality condition of greedy strategy). *If $c_i < c_{i+1}$ and $\frac{p_i - p_{i+1}}{c_i - c_{i+1}} < 1$ hold for $1 \leq i \leq N - 1$, then the greedy strategy is the only optimal strategy.*

By treating p_i as a function of c_i , the condition $\frac{p_i - p_{i+1}}{c_i - c_{i+1}} < 1$ can be essentially regarded as the discrete version of the economic property on the marginal utility $\Delta p / \Delta c$. What we essentially demonstrate is that, if p_i does not increase as much as c_i among resources, then the greedy strategy is optimal.

Despite our efforts in characterizing the optimality of the greedy strategy, the optimality condition established in Theorem 1 may be too stringent in many cases, and the greedy strategy may be far from optimal. To illustrate this, we consider an example where $N = 2$, $p_1 = 0.2$, $c_1 = 0.1$, $p_2 = 0.9$, $c_2 = 0.5$, and $T = 1$; clearly the greedy strategy accesses resource 1 leading to utility 0.1; however, the optimal strategy is to access resource 2 leading to utility 0.4.

Motivated by the above analysis, we proceed to derive the optimal strategy in the generic case.

B. Optimal Strategy based on Dynamic Programming

As in the previous subsection, we sort resources by decreasing p_i/c_i . By Lemma 2, an optimal strategy corresponds to accessing a subset of resources in that order. It remains to find the subset. To this end, we define an auxiliary function $\widehat{V}(r, l)$ ($0 \leq r \leq T$, $r \leq l \leq N$) to denote the minimal expected pseudo-cost by accessing at most r among the first l resources, i.e., $\widehat{V}(r, l) \triangleq \min_{\pi \subseteq \mathcal{N}_l, |\pi|=r} V(\pi)$, where \mathcal{N}_l denotes the subset of \mathcal{N} containing the first l resources.

The optimal algorithm we develop is based on dynamic programming, hinging on the following recursive property:

$$\widehat{V}(r, l) = \min\{c_l + q_l \cdot \widehat{V}(r - 1, l - 1), \widehat{V}(r, l - 1)\},$$

$$1 \leq r \leq T, r \leq l \leq N. \quad (1)$$

The above property follows from the observation that $\widehat{V}(r, l)$ is the minimum of the following two strategies: (1) an optimal strategy that accesses $(r - 1)$ resources among the first $(l - 1)$ resources followed by accessing resource l , (2) an optimal strategy that accesses r resources among the first $(l - 1)$ resources without accessing resource l . The border values are given by $\widehat{V}(0, l) = 1, 0 \leq l \leq N$.

The pseudo-code of our optimal algorithm is described in Algorithm 1, essentially consisting of two iterations (except the first **for** iteration for initialization). The first iteration computes the values of $\widehat{V}(r, l)$. Note that we only compute $\widehat{V}(r, l)$ for $r = 1$ to T and $l = r$ to $(N - T + r)$ for a given r , because these are sufficient to calculate the optimal cost $\widehat{V}(T, N)$. This can be physically implemented by a $T \times (N - T + 1)$ array. The second iteration starts from (T, N) and traces back to $(1, 1)$ to derive the optimal strategy π^* . During this process, r traces the current index of the resource accessed in π^* . Both the time and space complexity of Algorithm 1 is $O(NT)$, or more precisely, $O((N - T)T)$.

Algorithm 1 Finding optimal strategy: homogeneous case

```
1: Input:  $T, \{q_i, c_i\}_{1 \leq i \leq N}$ 
2: Output: an optimal strategy  $\pi^* = \{\pi_r^*\}_{1 \leq r \leq T}$ 
3: Sort resources by decreasing  $p_i/c_i$ 
4: for  $l = 1$  to  $N$  do
5:    $\widehat{V}(0, l) \leftarrow 1$ 
6: end for
7: for  $r = 1$  to  $T, l = r$  to  $N - T + r$  do
8:   if  $c_l + q_l \cdot \widehat{V}(r-1, l-1) \leq \widehat{V}(r, l-1)$  then
9:      $\widehat{V}(r, l) \leftarrow c_l + q_l \cdot \widehat{V}(r-1, l-1)$ 
10:  else
11:     $\widehat{V}(r, l) \leftarrow \widehat{V}(r, l-1)$ 
12:  end if
13: end for
14:  $r \leftarrow T$ 
15: for  $l = N$  to  $1$  do
16:   if  $c_l + q_l \cdot \widehat{V}(r-1, l-1) \leq \widehat{V}(r, l-1)$  then
17:      $\pi_r^* \leftarrow l$ 
18:      $r \leftarrow r - 1$ 
19:   end if
20: end for
```

Finding an optimum strategy can be cast to finding a path between the root of a tree to a leaf. This formulation helps us gain more insights on the problem. Specifically, we build a graph $G \triangleq (\mathcal{V}, \mathcal{E})$, where the set of vertexes $\mathcal{V} \triangleq \{(r, l)\}$, $0 \leq r \leq T, r \leq l \leq N$. We add an edge between the vertexes (r, l) and $(r-1, l-1)$ if $c_l + q_l \cdot \widehat{V}(r-1, l-1) \leq \widehat{V}(r, l-1)$, and between vertexes (r, l) and $(r, l-1)$ otherwise, where $\widehat{V}(r, l)$ is derived recursively by (1). We can check that G is a tree rooted at $\widehat{V}(0, 0)$, and that $\widehat{V}(T, N)$ is a leaf of G . If we can trace the path between $\widehat{V}(0, 0)$ and $\widehat{V}(T, N)$, we can establish the optimal strategy. The problem of finding an optimal strategy thus maps to the problem of finding a path between $\widehat{V}(0, 0)$ and $\widehat{V}(T, N)$. The last iteration of Algorithm 1 can be regarded as the procedure of finding such path by tracing from the leaf $\widehat{V}(T, N)$ back to the root $\widehat{V}(0, 0)$.

C. Complexity Reduction via Preprocessing

As an optimization to further reduce the complexity of Algorithm 1, we add a preprocessing phase identifying the resources that are guaranteed to be included (excluded, respectively) in any optimal strategy π^* . The preprocessing allows to reduce the size of the problem to be solved by Algorithm 1. Our preprocessing phase hinges on the structural properties of π^* given in Lemma 4, which hinges on another structural property given in Lemma 3 and Definition 3.

Definition 3 (Dominance). *Given a pair of resources i and j , we say that i dominates j in q (in c , respectively) if $q_i \leq q_j$ ($c_i \leq c_j$). We say that i dominates j if i dominates j in both q and c . The dominance is said to be strict if at least one inequality holds strictly.*

It can be noted that dominance defined above is a partial order, and that dominance in q and in c are total orders.

Lemma 3. *Given any strategy π , if there exists a resource $\pi_k \in \pi$ dominated by another resource $\pi'_k \notin \pi$, it holds that $V(\pi) \geq V(\pi')$, where π' denotes the strategy by replacing π_k by π'_k and keeping the other resources and their order as in π , i.e., $\pi'_j = \pi_j$ for $j \neq k$.*

Lemma 3 is intuitive to understand in the sense that if resource π'_k is superior to π_k in both reward and cost, the user should choose π'_k over π_k . Hence, π^* never contains any resource dominated by another resource outside π^* . This intuition is further formalized and generalized in Lemma 4.

Lemma 4. *We sort the resources increasingly by q_i and c_i , respectively. Let y_i and z_i denote the rank of resource i based on the sorting of q_i and c_i , respectively. The following two properties hold.*

- 1) *If $y_i + z_i \leq T + 1$, then $i \in \pi^*$.*
- 2) *If $y_i + z_i \geq N + T + 1$, then $i \notin \pi^*$.*

Armed with Lemma 4, we can develop a preprocessing procedure identifying all the resources satisfying the properties in Lemma 4. We can then safely remove all resources identified by the preprocessing procedure, and only solve the remaining problem by Algorithm 1 with reduced input size before reconstructing the optimal strategy by Lemma 2. The preprocessing is straightforward to implement, whose pseudo-code is omitted here. The complexity of preprocessing is dominated by the sorting operation, which sums up to $O(N \log N)$ using heap-based implementation [2].

To gain more quantitative insight on the benefit of preprocessing, we consider a system setting where q_i and c_i for each resource i are independently and randomly ranked in $[1, N]$. We can derive the percentage of resources removed by preprocessing. To this end, for any $i \in \mathcal{N}$, we have

$$\begin{aligned} \Pr[y_i + z_i \leq T + 1] &= \sum_{y=1}^T \Pr[z_i \leq T + 1 - y] \cdot \Pr[y_i = y] \\ &= \sum_{y=1}^T \frac{T + 1 - y}{N} \cdot \frac{1}{N} = \frac{T(T + 1)}{2N^2}. \end{aligned}$$

Symmetrically, we have

$$\Pr[y_i + z_i \geq N + T + 1] = \frac{(N - T)(N - T + 1)}{2N^2}.$$

Denote η the percentage of resources removed by preprocessing in average and let $\beta \triangleq T/N$, asymptotically we have

$$\begin{aligned} \eta &= \Pr[y_i + z_i \leq T + 1] + \Pr[y_i + z_i \geq N + T + 1] \\ &\simeq \frac{\beta^2 + (1 - \beta)^2}{2}. \end{aligned}$$

Algebraically we have $0.25 \leq \eta \leq 0.5$. The preprocessing procedure can thus reduce up to half of the total resources in the best case, thus halving the problem size passed to Algorithm 1. Even in the worst case where $T = N/2$ (corresponding to $\beta = 0.5$), it can still filter out 25% resources.

IV. THE HETEROGENEOUS CASE

We proceed to the generic case where the access time is heterogeneous among resources. We develop our analysis

by first demonstrating the hardness of the problem and then investigating the greedy strategy, followed by the development of a set of approximation algorithms approaching the minimal pseudo-cost with polynomial time and space complexity.

Theorem 2 (NP-hardness of sequential resource access). *The sequential resource access problem in Definition 1 is NP-hard.*

The proof of Theorem 2, detailed in [1], consists of relating the sequential resource access problem to the 0–1 Knapsack problem which is known to be NP-hard [17]. Given the NP-hardness of our problem, we explore two directions in the sequel analysis. The first is to develop specific strategies that perform optimally under certain conditions. The second is to design efficient approximation algorithms with bounded efficiency loss to the optimal utility. By efficient, we mean the algorithm has polynomial complexity in both time and space.

A. The Greedy Strategy

We consider the greedy strategy defined in Definition 2, i.e., accessing the resources in the decreasing order of p_i/c_i until success, or reaching time T . Theorem 3 establishes the sufficient conditions for the optimality of the greedy strategy.

Theorem 3 (Optimality condition of greedy strategy). *If $c_i < c_{i+1}$, $d_i \leq d_{i+1}$ and $\frac{p_i - p_{i+1}}{c_i - c_{i+1}} < 1$ hold for $1 \leq i \leq N - 1$, the greedy strategy is the only optimal strategy.*

Theorem 3 essentially demonstrates that the greedy strategy is optimal if (1) resources with higher access cost also incur higher delay, e.g., in the scenarios where the access cost is positively correlated or even proportional to the access delay, and (2) p_i does not increase as much as c_i among resources.

More generically, if the access delay and cost are not positively correlated, the greedy strategy may not be optimal, motivating our following analysis on the most generic case.

B. Approximation Algorithm Based on Dynamic Programming

In contrast to the homogeneous case, our sequential resource access problem in the heterogeneous case is NP-hard. We thus concentrate on developing approximation algorithms achieving near-optimal performance. We first formalize in Definition 4 the way we approximate an optimal strategy.

Definition 4 ((ϵ, δ) -optimality). *A strategy π is called (ϵ, δ) -optimal if $V(\pi) \leq (1 + \epsilon)V(\pi^*)$ and $d(\pi) \leq (1 + \delta)T$.*

By Definition 4, we allow the total access time to slightly exceed the given time constraint T . Our formulation makes sense if the delay constraint is not strictly inviolable so as to tolerate certain “overflow”. The quantity of such overflow can be controlled by the parameter δ . When $\epsilon = 0$, $(0, \delta)$ -optimality degenerates to the standard optimality with a δ relaxed constraint. When $\delta = 0$, $(\epsilon, 0)$ -optimality degenerates to the classic ϵ -optimality. The focus of our work is the development and analysis of two polynomial algorithms outputting an (ϵ, δ) -optimal strategy.

Our first approximation algorithm, presented in this subsection, extends from the dynamic programming approach in the homogeneous case. The core idea is to discretize the access delay to $O(1/\delta)$ values by scaling and rounding each d_i . This discretization step allows us to adapt the dynamic programming approach to the generic heterogeneous case. Specifically, our algorithm runs in two steps. To make the analysis concise without losing generality, we normalize $d_{min} \triangleq \min_{i \in \mathcal{N}} d_i = 1$.

Step 1: discretization. Given $\delta > 0$, we set $\lambda \triangleq \lceil 1/\delta \rceil$ as a scaling parameter⁴, further replace each d_i by $\hat{d}_i \triangleq \lfloor d_i \lambda \rfloor / \lambda$, i.e., we round down the fractional part of d_i , $d_i - \lfloor d_i \rfloor$, to the closest fraction of the form a/λ with $a < \lambda$ being an integer.

Step 2: dynamic programming. Let $\hat{T} \triangleq \lceil T/\delta \rceil$. For any pair of integers (t, l) , $0 \leq t \leq \hat{T}$, $1 \leq l \leq N$, let $\hat{V}(t, l)$ denote the minimal expected cost achievable by accessing a subset of resources among the first l resources with a total delay at most $\delta \hat{T}$, i.e., $\hat{V}(t, l) \triangleq \min_{\pi \in \Pi, \hat{d}(\pi) \leq t\delta} V(\pi)$ where $\hat{d}(\pi) \triangleq \sum_{i=1}^{|\pi|} \hat{d}_{\pi_i}$. By extending the results in the homogeneous case, we can establish the recursive property concerning $\hat{V}(t, l)$.

$$\hat{V}(t, l) = \min\{c_l + q_l \cdot \hat{V}(t - \hat{d}_l \lambda, l - 1), \hat{V}(t, l - 1)\}, \\ 1 \leq t \leq \hat{T}, 1 \leq l \leq N. \quad (2)$$

The pseudo-code of our approximation algorithm is described formally in Algorithm 2. The only notable difference compared to the homogeneous case is the way how $\hat{\pi}$ is established. Specifically, since we do not know the number of resources accessed in $\hat{\pi}$, we reconstruct $\hat{\pi}$ from the last element backward to the first and then reverse $\hat{\pi}$, as depicted in the last line of Algorithm 2, where REVERSE($\hat{\pi}$) denotes the operation of reversing $\hat{\pi}$. Both the time and space complexity of Algorithm 2 is $O(N\hat{T})$, i.e., $O(NT/\delta)$. Theorem 4 establishes the $(0, 2\delta)$ -optimality of Algorithm 2.

Theorem 4. *Algorithm 2 outputs a $(0, 2\delta)$ -optimal strategy.*

To further reduce the complexity of Algorithm 2, we can extend the preprocessing phase developed in the homogeneous case. It suffices to modify the definition of dominance by taking into account the heterogeneous access delay as below. Given a pair of resources i and j , we say that i dominates j in q (in c, d , respectively) if $q_i \leq q_j$ ($c_i \leq c_j, d_i \leq d_j$). We say that i dominates j if i dominates j in q, c and d .

C. Improved Approximation Algorithm

By examining Algorithm 2, we observe that the 2-dimensional table \hat{V} is not an efficient data structure for our problem, as some of the entries are not needed to determine an optimal strategy. Motivated by this observation, we develop an improved algorithm, whose key idea is exposed as follows.

In addition to discretize d_i , we also discretize the cost $V(\pi)$ for any strategy π by setting $\mu \triangleq \lceil n^*/\epsilon c_{min} \rceil$ and replacing $V(\pi)$ by $\hat{V}(\pi) \triangleq \lfloor V(\pi)\mu \rfloor / \mu$, where $n^* \triangleq \lfloor \pi^* \rfloor$ denotes the

⁴In practice, δ is usually very small. We can thus conveniently approximate λ as $1/\delta$ and treat it as a large integer.

Algorithm 2 Approximation algorithm: heterogeneous case

```
1: Input:  $T, \{q_i, c_i\}_{1 \leq i \leq N}$ 
2: Output: a  $(0, 2\delta)$ -optimal strategy  $\hat{\pi} = \{\hat{\pi}_i\}_{1 \leq i \leq |\hat{\pi}|}$ 
3: Sort resources by decreasing  $p_i/c_i$ 
4:  $\lambda \leftarrow \lceil 1/\delta \rceil, \hat{T} \leftarrow \lceil T/\delta \rceil$ 
5: for  $i = 1$  to  $N$  do
6:    $\hat{d}_i \leftarrow \lfloor d_i \lambda \rfloor / \lambda$ 
7: end for
8: for  $l = 0$  to  $N, t = 0$  to  $\hat{T}$  do
9:    $\hat{V}(t, l) \leftarrow 1$ 
10: end for
11: for  $l = 1$  to  $N, t = \hat{d}_l \lambda$  to  $\hat{T},$  do
12:   if  $c_l + q_l \hat{V}(t - \hat{d}_l \lambda, l - 1) \leq \hat{V}(t, l - 1)$  then
13:      $\hat{V}(t, l) \leftarrow c_l + q_l \hat{V}(t - \hat{d}_l \lambda, l - 1)$ 
14:   else
15:      $\hat{V}(t, l) \leftarrow \hat{V}(t, l - 1)$ 
16:   end if
17: end for
18:  $i \leftarrow 1$ 
19: for  $l = N$  to  $1$  do
20:   if  $c_l + q_l \hat{V}(t - \hat{d}_l \lambda, l - 1) \leq \hat{V}(t, l - 1)$  then
21:      $\hat{\pi}_i \leftarrow l, i \leftarrow i + 1, t \leftarrow t - \hat{d}_l \lambda$ 
22:   end if
23: end for
24: REVERSE( $\hat{\pi}$ )
```

number of resources accessed in π^* .⁵ Recall that $d_{min} = 1$, we can loosely upper-bound n^* by T and set $\mu = T/\epsilon c_{min}$.

We then generate a list Γ of all feasible pairs of $(\hat{d}(\pi), \hat{V}(\pi))$ ($\pi \in \Pi$) where $\hat{d}(\pi) \leq T$. Technically, Γ can be generated in N iterations. Initially we put $(0, 1)$ in Γ . At iteration l , from each pair $(\hat{d}(\pi), \hat{V}(\pi))$, we generate another pair $(\hat{d}(\pi) + \hat{d}_l, \lfloor (c_l + q_l \hat{V}(\pi)) \mu \rfloor / \mu)$ if $\hat{d}(\pi) + \hat{d}_l \leq T$ based on the recursive property of $\hat{d}(\pi)$ established in (2), and add the new pair to Γ if it does not duplicate any existing pair. By doing so, at the end of iteration l , each pair in Γ represents a strategy of accessing a subset of resources among the first l resources, whose total access time is upper-bounded by T , and inversely, each such strategy is represented by a pair. Once Γ is established, we return the strategy $\hat{\pi}$ corresponding to the pair with minimal discretized cost $\hat{V}(\hat{\pi})$ as an approximate optimal strategy.

The above algorithm can be further improved by noting that not all pairs in Γ are needed to derive an approximate optimal strategy. In fact we can safely remove a pair $(\hat{d}(\pi), \hat{V}(\pi))$ if there exists another pair $(\hat{d}(\pi'), \hat{V}(\pi'))$ in Γ such that the former is dominated by the latter in the sense $\hat{d}(\pi) \geq \hat{d}(\pi')$ and $\hat{V}(\pi) \geq \hat{V}(\pi')$. After eliminating all dominated pairs, each remaining pair $(\hat{d}(\pi), \hat{V}(\pi))$ is Pareto-optimal by satisfying the following conditions at the end of iteration l : (1) $\hat{d}(\pi)$ is the smallest access delay with access cost $\hat{V}(\pi)$ if only the

⁵Without introducing much ambiguity, we use the same notation \hat{V} as in Algorithm 2, as they essentially play similar roles in both algorithms.

first l resources are allowed to be accessed, (2) $\hat{V}(\pi)$ is the smallest access cost with access delay $\hat{d}(\pi)$ if only the first l resources are allowed to be accessed.

Based on the above idea, we can now modify our improved algorithm by adding a pruning phase at the end of each iteration. We sort the pairs in Γ in strictly increasing order of \hat{d} and in strictly decreasing order of \hat{V} due to elimination of dominated pairs. In iteration l , we produce a new list Γ' as follows: for each pair $(\hat{d}(\pi), \hat{V}(\pi)) \in \Gamma$, we add a pair $(\hat{d}(\pi) + \hat{d}_l, \lfloor (c_l + q_l \hat{V}(\pi)) \mu \rfloor / \mu)$ to Γ' if $\hat{d}(\pi) + \hat{d}_l \leq T$. Since the pairs in Γ is in increasing order of $\hat{d}(\pi)$, the process of establishing Γ' can be terminated whenever a pair $(\hat{d}(\pi), \hat{V}(\pi))$ in Γ is reached, for which $\hat{d}(\pi) + \hat{d}_l(\pi) > T$. Once Γ' is established, we merge Γ' to Γ by removing dominated pairs. This is easily accomplished given the strict ordering of $\hat{d}(\pi)$ and $\hat{V}(\pi)$ in the list. At the end of the last iteration (iteration N), the last pair in Γ gives the cost and the corresponding access delay for an approximated optimal solution we look for. To reconstruct the set of resources to access, we can simply add a pointer to each pair, pointing to the parent pair from which the current pair is generated. Via these pointers we can trace back from the last pair in Γ to the first pair $(0, 1)$, and reconstruct the set of resources to access.

The pseudo-code of the improved approximation algorithm is described in Algorithm 3. The algorithm is mainly composed of two iterations, except that of initialization. The first iteration (line 9 to line 19) establishes Γ as described above. The second iteration (the **while** loop) reconstructs the entire set of resources to access. In our algorithm, the following elementary functions are used. They are straightforward to implement and are omitted here for brevity.

- INSERT(e, γ) inserts the element e in the list γ .
- MERGE(Γ, Γ') merges Γ' into Γ , removing dominated pairs and sorting the resulting list. Given the structure of Γ , where the pairs are in increasing order of $\hat{d}(\pi)$ and decreasing order of $\hat{V}(\pi)$, MERGE(Γ, Γ') can be implemented in linear time to the size of Γ and Γ' .
- LAST(Γ) returns the last element in the list Γ .
- FIND(γ, ptr) returns the resource such that by accessing it, the pair (access delay, cost) changes from the values in the pair pointed by ptr to the values in γ . This can be achieved via a standard search in $O(\log N)$ time if we provide a sorted list of resources based on access delay.

We next analyze the time and space complexity of Algorithm 3. In this regard, we first derive the maximal size of Γ and Γ' . Note that (1) the finest granularity of the discretized access delay is $1/\lambda$, and the total delay is upper bounded by T , (2) the finest granularity of the discretized cost $V(\pi)$ is $1/\mu$, and $V(\pi) \leq 1$ from Lemma 1, and (3) all the dominated pairs are removed from Γ , which implicates that (3.1) for any discretized access delay \hat{d} there is at most one pair with delay \hat{d} , (3.2) for any discretized cost \hat{v} there is at most one pair with cost \hat{v} . Therefore, the maximal size of Γ is $\min\{T/\lambda, 1/\delta\}$, which, by injecting λ and μ , sums up to $\min\{|\pi^*|T/\epsilon c_{min}, 1/\delta\}$, where $|\pi^*|$ can be upper-bounded by

Algorithm 3 Improved approximation algorithm

```
1: Input:  $T, \{q_i, c_i\}_{1 \leq i \leq N}$ 
2: Output: an  $(\epsilon, \delta)$ -optimal strategy  $\hat{\pi} = \{\hat{\pi}_i\}_{1 \leq i \leq |\hat{\pi}|}$ 
3: Sort resources by decreasing  $p_i/c_i$ 
4:  $\lambda \leftarrow \lceil 1/\delta \rceil, \mu \leftarrow \lceil 1/\epsilon c_{\min} T \rceil$ 
5: for  $i = 1$  to  $N$  do
6:    $\hat{d}_i \leftarrow \lfloor d_i \lambda \rfloor / \lambda$ 
7: end for
8:  $\Gamma \leftarrow (0, 1, \emptyset)$ 
9: for  $l = 1$  to  $N$  do
10:   $\Gamma' \leftarrow \emptyset$ 
11:  for each  $\gamma = (\hat{d}(\pi), \hat{V}(\pi), ptr(\pi)) \in \Gamma$  do
12:    if  $\hat{d}(\pi) + \hat{d}_l > T$  then
13:      break
14:    else
15:       $ptr \leftarrow \&\gamma$ 
16:      INSERT( $(\hat{d}(\pi) + \hat{d}_l, [(c_i + q_l \hat{V}(\pi))\mu] / \mu,$ 
17:         $ptr), \Gamma')$ 
18:    end if
19:  end for
20:  MERGE( $\Gamma, \Gamma'$ )
21: end for
22:  $i \leftarrow 1, \gamma \triangleq (\hat{d}(\pi), \hat{V}(\pi), ptr(\pi)) \leftarrow \text{LAST}(\Gamma)$ 
23: while  $ptr(\pi) \neq \emptyset$  do
24:    $\hat{\pi}_i \leftarrow \text{FIND}(\gamma, ptr(\pi))$ 
25:    $\gamma \leftarrow ptr(\pi), i \leftarrow i + 1$ 
26: end while
27: REVERSE( $\hat{\pi}$ )
```

T . By the construction of Γ' in Algorithm 3, where each pair in Γ generates at most one pair in Γ' , we can upper-bound the size of Γ' also by $\min\{|\pi^*|T/\epsilon c_{\min}, 1/\delta\}$. The overall time complexity of Algorithm 3, dominated by the second **for** loop, sums up to $O(N \min\{|\pi^*|T/\epsilon c_{\min}, 1/\delta\})$. Algorithm 3 also needs to stock Γ and Γ' . Hence the space complexity sums up to $O(\min\{T^2/\epsilon c_{\min}, 1/\delta\})$, if we do not count the space required to store the parameters. We remark that this is an order-of-magnitude space gain compared to Algorithm 2.

Theorem 5 establishes the (ϵ, δ) -optimality of Algorithm 3.

Theorem 5. *Algorithm 3 gives an (ϵ, δ) -optimal strategy.*

We conclude this section by summarizing and comparing the different algorithms we develop in Table II. For the time and space complexity, the time and space required to sort and store the parameters of each resource (e.g., p_i, c_i, d_i) are not taken into account. The minimal access delay d_{\min} is normalized to 1.

V. NUMERICAL ANALYSIS

In this section, we conduct numerical analysis to evaluate the performance of our approximation algorithm. We evaluate our algorithm against the following three strategies:

- **Randomized strategy:** the user accesses a randomly chosen resource each time until success or reaching T ;

- **Greedy- p :** the user accesses the resources in decreasing order of p_i until success or the end of T ;
- **Greedy- c :** the user accesses the resources in decreasing order of c_i until success or the end of T ;
- **Greedy- p/c :** the user accesses the resources in decreasing order of p_i/c_i until success or the end of T ; this strategy is the greedy strategy analyzed in Section III and IV.

Specifically, we simulate a system consisting of $N = 20$ resources, where each p_i is randomly chosen from $[0, 1]$, c_i randomly from $[0, p_i]$, $d_i = 1$ for the homogeneous case and randomly chosen from $[0, 2]$ in the heterogeneous case, T varies from 4 to 16. We take the randomized strategy as the baseline and trace the ratio between the utility of the other simulated strategies and the utility of the randomized strategy, denoted by Ψ . In other words, Ψ quantifies the performance gain of the simulated strategy over the randomized baseline strategy. For our algorithm, we run Algorithm 3 with both δ and ϵ set to 0.01. Figure 1(a) and 1(b) traces the average of Ψ over 10000 runs for the homogeneous and heterogeneous cases, respectively. We make the following observations.

Our algorithm performs constantly and significantly better than all the other strategies. This is coherent to the theoretical results as our algorithm is proved to approach the system optimum.

Among the other strategies, none outperforms the others in all the simulated scenarios. We can only observe that, in the average sense, greedy- p and greedy- p/c performs better with small T , while greedy- c performs better with large T . However, this performance trade-off depends on the system parameters.

The performance gain of our algorithm is more significant where approximately half of the resources can be selected (i.e., $T = 8$ and 12). This is also the least tractable cases, since when the number of selected resources approaches $N/2$, the number of choices are maximized, hence the optimum strategy may significantly outperforms a heuristic one, as demonstrated by our results.

VI. RELATED WORK

The problem we tackle in this paper is related to the sequential decision making problem, among which perhaps the most pertinent example to our context is the classic optimal stopping problem [16]. Stemmed from the famous secretary problem dating back to the late 1950's, the theory of optimal stopping has received a lot of research attention by constituting today a field of study in applied mathematics and statistics [12]. Briefly speaking, the theory of optimal stopping is concerned with the problem of deciding when to stop based on sequentially observed random variables in order to maximize an expected payoff. Interested readers may refer to [15], [16] for a detailed treatment. Compared to the standard optimal stopping problem where only a binary action, stopping or proceeding to observe, needs to be taken upon each observation, our problem has a more combinatorial flavor as the decision maker needs to choose which resources to probe and in what sequence. In contrast, our problem does not have

Algorithm	Optimality	Time complexity	Space complexity
Greedy algorithm	optimal under certain condition	$O(T)$	$O(T)$
DP-based algorithm (Algo. 1): homogeneous case	optimal	$O(NT)$	$O(NT)$
Approximation algorithm (Algo. 2)	$(0, 2\delta)$ -optimal	$O(NT/\delta)$	$O(NT/\delta)$
Improved approximation algorithm (Algo. 3)	(ϵ, δ) -optimal	$O\left(N \min\left\{\frac{T^2}{\epsilon c_{min}}, \frac{1}{\delta}\right\}\right)$	$O\left(\min\left\{\frac{T^2}{\epsilon c_{min}}, \frac{1}{\delta}\right\}\right)$

TABLE II. Comparison of different algorithms developed in our work

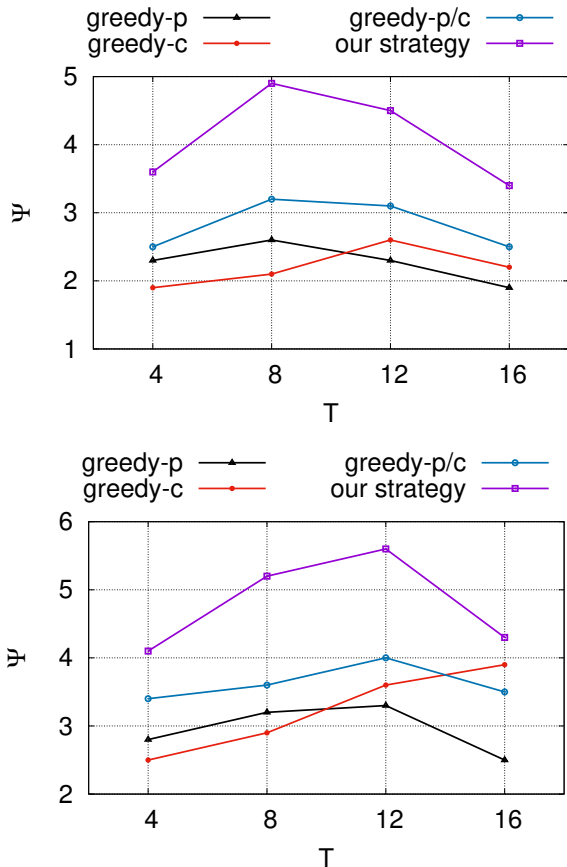


Fig. 1: Performance of different access strategies: heterogeneous (top) and homogeneous (bottom) cases.

the stochastic component as in the optimal stopping problem since the system parameters are known *a priori*.

From an application point of view, since the last decade, there has been a surge of interest in using optimal stopping theory in emerging communication and networking paradigms [3], [8]–[10], [13], [14]. These studies consider a variety of scenarios where a user seeks a sequential strategy that probes a subset of resources with potentially unknown distribution of states so as to maximize her utility. The considered problem can be cast to the optimal stopping problem by regarding resource states as the observable variable. Two variants have been investigated depending on whether the user is allowed to use a previously probed resource or not. Under certain system setting, the optimal probing policy has a threshold structure and can be derived in a subset of special

cases (e.g., with two resources) [9]. Compared to these works addressing concrete application examples with particular constraints and specificity therein, our work take a more generic view without sacrificing the complexity and intrinsic structure of the problem. Therefore, we believe that our analysis, despite being generic, can serve as a complementary line of research, and provide valuable insights in a variety of resource access problems in emerging networking and computing systems.

VII. CONCLUSION AND PERSPECTIVE

We have formulated and analyzed a generic resource access problem, where a user, having access to a set of resources with heterogeneous success probability, access cost and latency, seeks an optimal strategy to achieve the best trade-off between limiting the access cost and maximizing the overall success probability. We proved that the problem of finding an optimal strategy is NP-hard. Given its NP-hardness, we presented a greedy strategy that can be implemented in linear time and mathematically establish sufficient conditions for its optimality. We then developed a series of polynomial-time approximation algorithms approaching an optimum solution in the sense of (ϵ, δ) -optimality.

A natural generalization of our current work, as mentioned in the Introduction, is to extend our algorithmic framework to investigate the case with correlated resources. This is an important facet of the problem, and will provide valuable insight on how to exploit the correlation among resources to improve the performance. Another direction is to investigate the scenario, where the user is allowed to access up to C resources simultaneously each time, and can perform multiple rounds of queries, as long as the total delay does not exceed T . This direction adds more combinatorial flavor to the problem.

REFERENCES

- [1] Sequential resource access: theory and algorithm, technical report, drive.google.com/file/d/1ok5ftYtl7WwyXYY52QS-QD2Ps13GIFSR.
- [2] Algorithms. *Communications of the ACM*, 7(6):347–349, June 1964.
- [3] I. Alghamdi, C. Anagnostopoulos, and D. P. Pezaros. On the optimality of task offloading in mobile edge computing environments. In *Proc. IEEE GLOBECOM*, 2019.
- [4] K. Avrachenkov and P. Jacko. Ccn interest forwarding strategy as multi-armed bandit model with delays. In *Proc. Intl. Conf. Network Games, Control and Optimization (NetGCoop)*, 2012.
- [5] D. S. Berger, R. K. Sitaraman, and M. Harchol-Balter. Adaptsize: Orchestrating the hot object memory cache in a content delivery network. In *Proc. NSDI*, 2017.
- [6] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [7] F. Boccardi, R. W. Heath, A. Lozano, T. L. Marzetta, and P. Popovski. Five disruptive technology directions for 5g. *IEEE Communications Magazine*, 52(2):74–80, February 2014.

- [8] C. Borrego, M. Amadeo, A. Molinaro, P. M. Mendes, R. C. Sofia, N. Magaia, and J. Borrel. Forwarding in opportunistic information-centric networks: an optimal stopping approach. *IEEE Communications Magazine*, 2020.
- [9] N. B. Chang and M. Liu. Optimal channel probing and transmission scheduling for opportunistic spectrum access. *IEEE/ACM Transactions on Networking*, 17(6):1805–1818, 2009.
- [10] H. T. Cheng and W. Zhuang. Simple channel sensing order in cognitive radio networks. *IEEE Journal on Selected Areas in Communications*, 29(4):676–688, 2011.
- [11] I. Cohen, G. Einziger, R. Friedman, and G. Scalosub. Access strategies for network caching. In *Proc. INFOCOM*, 2019.
- [12] T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- [13] T. Ouyang, X. Chen, L. Zeng, and Z. Zhou. Cost-aware edge resource probing for infrastructure-free edge computing: From optimal stopping to layered learning. In *Proc. IEEE RTSS*, 2019.
- [14] Y. Pei, Y. Liang, K. C. Teh, and K. H. Li. Energy-efficient design of sequential channel sensing in cognitive radio networks: Optimal sensing strategy, power allocation, and sensing order. *IEEE Journal on Selected Areas in Communications*, 29(8):1648–1659, 2011.
- [15] G. Peskir and A. Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Birkhuser Basel, 2006.
- [16] A. N. Shiryaev. *Optimal Stopping Rules*. Springer, 2011.
- [17] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 1st edition, 2011.
- [18] G. Zhang, Y. Li, and T. Lin. Caching in information centric networking: A survey. *Computer Networks*, 57(16):3128–3141, 2013.