Secure Outsourcing of Multi-Armed Bandits

Radu Ciucanu
INSA Centre Val de Loire, LIFO
radu.ciucanu@insa-cvl.fr

Pascal Lafourcade
Univ. Clermont Auvergne, LIMOS
pascal.lafourcade@uca.fr

Marius Lombard-Platet
Univ. PSL, DIENS, Be-Studys
marius.lombard-platet@ens.fr

Marta Soare
Univ. Orléans, LIFO
marta.soare@univ-orleans.fr

***Abstract*—We consider the problem of cumulative reward maximization in multi-armed bandits. We address the security concerns that occur when data and computations are outsourced to an honest-but-curious cloud i.e., that executes tasks dutifully, but tries to gain as much information as possible. We consider situations where data used in bandit algorithms is sensitive and has to be protected e.g., commercial or personal data. We rely on cryptographic schemes and propose UCB-DS, a distributed and secure protocol based on the UCB algorithm. We prove that UCB-DS computes the same cumulative reward as UCB while satisfying desirable security properties. In particular, cloud nodes cannot learn the cumulative reward or the sum of rewards for more than one arm. Moreover, by analyzing messages exchanged among cloud nodes, an external observer cannot learn the cumulative reward or the sum of rewards produced by some arm. We show that the overhead due to cryptographic primitives is linear in the size of the input. Our implementation confirms the linear-time behavior and the practical feasibility of our protocol, on both synthetic and real-world data.**

I. Introduction

The *stochastic multi-armed bandit* game is a sequential learning framework where a learning agent aims at maximizing its *cumulative reward* while successively interacting with an uncertain environment. At each time step, the agent chooses an action *(a bandit arm)* from a fixed set of actions with unknown associated values. The environment responds with a stochastic feedback *(reward)* drawn from the distribution associated with the chosen action. The agent uses the received feedback to update its estimate of the values for the chosen action and to decide which action to choose next. The agent has to continuously face the so-called exploration-exploitation dilemma and decide whether to *explore* by choosing actions with more uncertain associated values, or to *exploit* the information already acquired by choosing the action with the seemingly largest associated value. Cumulative reward maximization has been already extensively studied for several multi-armed bandit settings (see [1] for a survey) and for various applications, from clinical trials, to online advertising and recommendation systems. In this paper, we address the security concerns that occur when outsourcing the cumulative reward maximization data and computations to the cloud.

Our scenario is inspired by the *machine learning as a service* cloud computing model, for which security is known as a major concern [2]. As a motivating example, assume:

• A *data owner*: a company that wants to monetize some collected data, while keeping ownership over it. The collected data may be a large quantity of surveys on customer preferences for several products. By product, we mean any type of object or service. The K bandit arms are the surveyed products and only the data owner knows their associated rewards, based on the collected surveys.

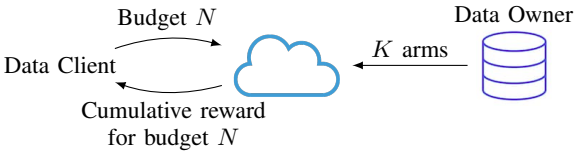

Fig. 1. Outsourcing data and computations.

• A *data client*: a company that wants to spend some budget to use some of the data owner's data. The data client may be a small company that cannot afford doing its own surveys, but wants to estimate the income that it could generate for the products surveyed by the data owner. The cumulative reward captures such information because it sums the rewards produced by each product. The budget N is the number of data owner's surveys used to compute the cumulative reward and the bandit algorithm has to decide *how* to choose these N surveys in order to maximize the cumulative reward. A larger budget gives a higher accuracy for the largest cumulative reward. The data client only sees the cumulative reward, without knowing the values associated to each arm.

We assume that the interaction between the data owner and the data client is done using the *cloud* (as shown in Fig. 1), where both data and computations are *outsourced*. The data owner does the data outsourcing, and the data client interacts directly with the cloud, by sending the budget and receiving the obtained cumulative reward. The outsourced data may be sensitive (e.g., personal, commercial, or medical data). We want the outsourced learning algorithm to be run while protecting data against unauthorized access. The problem that we address is how to allow the data client to obtain precisely the same cumulative reward as with a standard bandit algorithm, *Upper Confidence Bound (UCB)* [1], [3], within a reasonable computation time and while preserving the data security. Indeed, the outsourced data can be communicated over an untrustworthy network and processed on some untrustworthy machines, where malicious cloud users may learn sensitive data that belongs only to the data owner.

The privacy-preserving cumulative reward maximization is a hard problem. To solve it, the authors of [4], [5], [6] use *differential privacy* introduced in [7]. However, in these approaches, the returned reward is not the same reward obtained for the data client's budget using standard algorithms. This happens because differential privacy guarantees depend on noise be-

ing injected in the input/output. We take a complementary approach by relying on *cryptography* instead of differential privacy. To the best of our knowledge, our approach is original and its goal is to give security guarantees, while obtaining the same output as standard (non-secure) algorithms. The security for obtaining the same output has a price because the computation time may increase because of cryptographic primitives that are time-consuming in practice. More precisely, we require that the data owner (which can be seen as an oracle knowing the reward functions associated with each arm) encrypts her data before outsourcing it to the cloud. Then, the cumulative reward maximization algorithm is run directly in the encrypted domain, and the (encrypted) output should be exactly the same as for standard UCB, with the cost of an increased computation time.

From a theoretical point of view, the problem could be straightforwardly solved by using a fully homomorphic encryption scheme [8], which allows to compute any function directly in the encrypted domain. However, it remains an open question how to make such a scheme work fast and be accurate in practice. Indeed, the state-of-the-art fully homomorphic systems (SEAL[1] and HElib[2]) yield only approximate results when they work with real numbers, by using the CKKS scheme [9]. Hence, it is not currently possible to program an algorithm such as UCB in a fully homomorphic system and obtain exactly the same result as in the standard, non-encrypted UCB.

Consequently, our challenge is to rely on simpler cryptographic schemes and design a distributed protocol with several cloud node participants such that each of them can only learn the specific data needed for performing its task and nothing else e.g., if a participant does in clear computations on real numbers, these computations concern data of only one arm, and no other participant has access to this piece of data. Our distributed algorithm returns exactly the same cumulative reward as UCB, while satisfying desirable security properties such as: only the data client can see the cumulative reward, which cannot be learned by any cloud node participant nor by an external observer. We precisely characterize our security model and security guarantees later on in the paper. To achieve our goals, we rely on *indistinguishable under chosen-plaintext attack* (IND-CPA) cryptographic schemes: symmetric encryption AES-CBC [10], [11] and asymmetric partially homomorphic Paillier's scheme [12]. We formally prove the security of our protocol and we precisely characterize the number of needed cryptographic operations.

a) Related work: Each line in Table I corresponds to a standard problem in stochastic multi-armed bandits. The most popular problem is *cumulative reward maximization* and UCB is a standard algorithm for solving it [1], [3]. There is a recent line of research on enhancing algorithms such as UCB with differential privacy [4], [5], [6]. There are some fundamental differences between this line of work and our work based on cryptography. On the one side, the running time overhead of differentially-private algorithms is negligible, whereas our approach has an overhead in computation time coming from the use of cryptographic primitives. On the other side, the cumulative reward returned by differentially-private algorithms is different from the output of standard UCB. Indeed, to obtain differentially-private guarantees for a bandit algorithm, noise is added to the algorithm input or output. Thus, the cumulative reward obtained using a differentially-private algorithm is different from that obtained by the algorithm without privacy guarantees. This is reflected in the *regret analysis* of the algorithms (where the *regret* is given by the difference in the cumulative reward obtained by a learning agent and the best cumulative reward possible obtained by always playing the best arm): the regret of differentially-private bandit algorithms have as overhead an additive [6] or multiplicative factor [4], [5] with respect to the regret of their non-private version. In contrast, our cryptography-based algorithm is guaranteed to return exactly the same cumulative reward as standard UCB.

TABLE I
SUMMARY OF RELATED WORK AND POSITIONING OF OUR CONTRIBUTION.

	Differential privacy	Cryptography
Cumulative reward maximization aka cumulative regret minimization	[4], [5], [6]	**This paper**
Best arm identification aka simple regret minimization	Not yet studied to the best of our knowledge	[13]

The second line in Table I corresponds to a different bandit problem that is best arm identification [14], equivalent to minimizing the simple regret, that is the difference between the values associated with the arm that is actually the best and the best arm identified by the algorithm. From the cryptography point of view, there exists a distributed algorithm [13] that enhances the Successive rejects algorithm [14] for best arm identification with security guarantees that are similar to the ones from this paper. Naturally, the algorithms that are secured (Successive rejects [14] in [13] and UCB [3] in this paper) solve different problems, thus the corresponding secure protocols are different and cannot be reduced to one another.

All related works discussed thus far are for standard stochastic bandit models. Securing cumulative reward maximization algorithms using cryptography has been recently studied for a different bandit model i.e., linear bandits [15], where the arms are vectors and the rewards are unknown linear functions of the arms. The corresponding secure protocols are again different and cannot be reduced to one another.

b) Summary of contributions and paper organization: In Sect. II, we introduce some basic notions: standard UCB algorithm and some cryptographic tools. Then, Sect. III is the core of our contribution:

- We propose UCB-DS, a secure and distributed protocol for cumulative reward maximization that guarantees the same cumulative reward as standard UCB.
- We show that UCB-DS satisfies desirable security properties that we precisely characterize.
- We analyze the theoretical complexity of UCB-DS, by

[1]https://github.com/Microsoft/SEAL
[2]http://homenc.github.io/HElib/

quantifying the number of needed cryptographic primitives: $O(NK)$ AES-CBC encryptions/decryptions, K Paillier encryptions, and one Paillier decryption.
- We propose the UCB-DS2 refinement, with stronger security guarantees at the price of K more AES-CBC keys and $O(NK)$ more AES-CBC encryptions/decryptions, and the same number of Paillier encryptions/decryptions.

In Sect. IV, we include a proof-of-concept empirical evaluation that confirms the theoretical complexity, and shows the scalability and practical feasibility of our protocols, on synthetic and real-world data. Finally, we conclude our paper and outline directions for future work in Sect. V.

II. Preliminaries

We first recall the *UCB* algorithm [3]. Then, we briefly present two cryptographic schemes that we use to build our protocols: *Paillier asymmetric encryption* scheme and *AES-CBC symmetric encryption*, which are both *IND-CPA secure*.

a) Upper Confidence Bound (UCB): is a class of algorithms commonly used when facing the exploration-exploitation dilemma. Each bandit arm is associated with a distribution whose mean is unknown to the learning agent. When pulling an arm, the agent observes an independent reward drawn from the distribution associated to the chosen arm. Specifically, we consider rewards drawn from Bernoulli distributions with expected values $\mu_1, \ldots, \mu_K$ unknown to the agent. For a chosen arm i, a call to the function $pull(i)$ randomly returns 0 or 1 according to the associated Bernoulli distribution, i.e., the probability of returning 1 is μ_i and the probability of 0 is 1–μ_i. The agent sequentially selects the N arms to be pulled with the goal of maximizing the sum of rewards.

To guide the choice of the learner, arm scores have been proposed [16] to construct *upper confidence bounds* (UCB) based on the empirical mean of arm-specific rewards and the number of arm pulls. In the class of UCB algorithms, an important breakthrough was the introduction of algorithms with a finite-time analysis [3]. Specifically, in the UCB algorithm [3] presented in Fig. 2, for each arm i, the score B_i is an upper-confidence bound on μ_i, obtained as the sum between (i) the *exploitation* term given by the empirical mean of rewards observed from arm i, and (ii) the *exploration* term, which takes into account the uncertainty. Notice that after each observed reward, scores for all arms are updated, since the *exploration* term $\sqrt{\frac{2\ln(t)}{n_i}}$ depends on the total number of rewards observed up to current round t. Thus, an arm i being pulled few times (i.e., with small n_i) will have a relatively large *exploration* term. The score B_i is thus an *optimistic* estimate for the value associated to arm i, since it can be interpreted as the largest statistically plausible mean value associated to arm i, given the observed rewards. As shown in Fig. 2, UCB chooses to pull next the arm with the largest updated B_i score, thus following the principle of *optimism in the face of uncertainty*. This principle suggests to follow what seems to be the best arm, based on the *optimistically* constructed scores. The same principle is employed in various sequential decision making problems (see [17] for a survey).

Input: Budget N, number of arms K

Unknown environment: K distributions associated to the K arms, with expected values $\mu_1, \ldots, \mu_K$ unknown to the learning agent. The agent has access to a reward function $pull(.)$ that can be called N times. A call $pull(i)$ returns a random value from the distribution associated to arm i.

Output: Sum of observed rewards for all arms

/ Initialization: pull each arm once & initialize variables */*
for $1 \leq i \leq K$
  let $r = pull(i)$ */* Random reward for arm i */*
  let $s_i = r$ */* Sum of observed rewards for arm i */*
  let $n_i = 1$ */* Number of pulls of arm i */*
  let $B_i = \frac{s_i}{n_i} + \sqrt{\frac{2\ln(K)}{n_i}}$ */* B_i is an UCB on μ_i */*
/ Exploration-exploitation: pull one arm at each round t */*
for $K+1 \leq t \leq N$ */* Only a budget of $N-K$ is left */*
  let $i_m = arg\,max_{1 \leq i \leq K}(B_i)$ */* Ties broken randomly */*
  let $r = pull(i_m)$
  let $s_{i_m} = s_{i_m} + r$
  let $n_{i_m} = n_{i_m} + 1$
  for $1 \leq i \leq K$
    let $B_i = \frac{s_i}{n_i} + \sqrt{\frac{2\ln(t)}{n_i}}$
return $s_1 + \ldots + s_K$

Fig. 2. UCB Algorithm [3].

b) Paillier asymmetric encryption: Paillier's cryptosystem is *additive homomorphic* [12]. Let m_1 and m_2 be two plaintexts in $\mathbb{Z}_n$. The product of the two associated ciphertexts with the public key pk, denoted $c_1 = \mathcal{E}_{\mathsf{pk}}(m_1)$ and $c_2 = \mathcal{E}_{\mathsf{pk}}(m_2)$, is the encryption of the sum of m_1 and m_2. Indeed, we have: $\mathcal{E}_{\mathsf{pk}}(m_1) \cdot \mathcal{E}_{\mathsf{pk}}(m_2) = \mathcal{E}_{\mathsf{pk}}(m_1 + m_2)$. We also denote by $\mathcal{D}_{\mathsf{sk}}(c)$ the decryption of the cipher c by the secret key sk.

c) AES-CBC symmetric encryption: AES [10] is a NIST standard for encrypting messages of 128 bits. We use it with CBC mode (Cipher Block Chaining) and denote $c = \mathsf{Enc}(m)$ the encryption of m and $m = \mathsf{Dec}(c)$ the decryption of c with the same symmetric key shared between the participants.

Both Paillier and AES-CBC are **IND-CPA**: (i) Paillier is IND-CPA under the decisional composite residuosity assumption [12], and (ii) AES-CBC is IND-CPA under the assumption that AES is a pseudo-random permutation [11]. All theoretical security properties of our protocols also hold if we choose any other IND-CPA symmetric scheme instead of AES-CBC, and any other additive homomorphic IND-CPA asymmetric scheme instead of Paillier. Our choice to rely on the aforementioned schemes is due to practical reasons. AES-CBC is very efficient in practice and implemented in standard libraries for modern programming languages. Paillier is also supported by a number of libraries that can be used in practice.

III. UCB-DS: A Distributed and Secure Protocol Based on UCB Algorithm

We define the security model in Sect. III-A. We propose our secure protocol UCB-DS (Sect. III-B), and we analyze its correctness, security, and complexity (Sect. III-C). We introduce a refinement of UCB-DS in Sect. III-D.

A. Security Model

As outlined in Introduction and in Fig. 1, we assume that the data (i.e., the reward functions associated to K bandit arms) and the computations (i.e., the cumulative reward maximization algorithm) are outsourced to an *honest-but-curious* cloud. This means that the cloud executes tasks dutifully, but tries to extract as much information as possible from the data that it sees. Our model follows the classical formulation in [18] (Ch. 7.5, where *honest-but-curious* is denoted *semi-honest*), in particular (i) each cloud node is trusted: it correctly does the required computations, it does not sniff the network and it does not collude with other nodes, and (ii) an external observer has access to all messages exchanged over the network.

The data client indicates to the cloud her budget N and receives the cumulative reward R that the cloud computes using the K arms outsourced by the data owner and the data client's budget N. The data client does not have to do any computation, except for decrypting R when the data client receives this information encrypted from the cloud. We expect the following *security properties*:

1) No cloud node can learn the cumulative reward.
2) The data client cannot learn information about the rewards produced by each arm or which arm has been pulled at some round.
3) By analyzing the messages exchanged between different cloud nodes, an external observer cannot learn the cumulative reward, the sum of rewards produced by some arm, or which arm has been pulled at some round.

We give a brief intuition for each property. Property 1 implies that only the data client can see in clear the cumulative reward for which she spends a budget. Property 2 ensures that the data client can see only the information for which she pays, and nothing else. Otherwise, depending on the difficulty of the bandit problem, the data client could estimate the arm values based on the contribution of each arm to the cumulative reward, which would leak information that should be known only by the data owner. Property 3 states that if some curious cloud admin analyzes all messages exchanged over the network, then she should have no clue on any input, output, or intermediate data that is used by the cumulative reward maximization algorithm.

We design a distributed protocol that satisfies the aforementioned properties by exchanging only encrypted messages, and by distributing the computations among several cloud node participants, each of them having access only to the specific data that it needs for performing its task and nothing else. The challenge is to efficiently distribute tasks among as few cloud participants as possible, while minimizing the time needed for cryptographic primitives.

B. Overview of UCB-DS

In Fig. 3, we present an overview of UCB-DS. There are $K+1$ cloud participants: K arm nodes R_i and a node AS (Arm Selector) that is the controller of the protocol. We assume that the data owner and all cloud participants share the same symmetric AES-CBC key, used for encryption function Enc. The data client (DC) generates a Paillier's key pair (pk, sk) and for sake of clarity we denote $\mathcal{E}_{\mathsf{DC}}(m)$ for $\mathcal{E}_{\mathsf{pk}}(m)$. By $[\![x]\!]$, we denote the set $\{1,\ldots,x\}$, and by $y||z$ we denote the concatenation of y and z. UCB-DS works as follows:

- Fig. 3(a) (steps 0 and 1). For $i \in [\![K]\!]$, the data owner outsources to arm node R_i the reward function (encrypted with Enc) associated to arm i. The data client sends to the cloud her budget N.
- Fig. 3(b) (steps 2, 3, and 4). This is the core of the protocol, being done during $1+N-K$ iterations: once for the initialization phase of UCB and $N-K$ times for the exploration-exploitation phase of UCB cf. Fig. 2. For each iteration, all arm nodes interact to decide which arm should be pulled next and communicate this information to AS. The arm nodes communicate in a random order, which changes at each iteration. All messages exchanged between nodes are encrypted with Enc. Although each arm node stores information about its rewards, it never reveals this information to other nodes.
- Fig. 3(c) (steps 5 and 6). After spending the data client's budget, each arm node sends to AS the sum of rewards that it produced, encrypted with $\mathcal{E}_{\mathsf{DC}}$. Due to the additive homomorphic property of Paillier cryptosystem, AS is able to sum up the K partial rewards to compute the cumulative reward $\mathcal{E}_{\mathsf{DC}}(R)$ directly in the encrypted domain. Only the data client can decrypt this information.

We next detail each step and present pseudocode only when the step is not trivial.

a) Step 0: We recall (cf. Fig. 2) that the data owner knows $\mu_1,\ldots,\mu_K$ defining K Bernoulli distributions associated to the K arms. The data owner sends to each arm node R_i the encrypted value $\mathsf{Enc}(\mu_i)$, for $i \in [\![K]\!]$. Since the data owner and the cloud share the symmetric key, then each arm node R_i can decrypt and obtain μ_i. Moreover, each node R_i initializes to 0 the following two variables that it later on updates during the protocol: s_i (i.e., sum of rewards for arm i) and n_i (i.e., number of times the arm i has been pulled). Additionally, each arm node R_i initializes a variable $t = K-1$, which is later on updated and needed for the computation of B_i.

b) Step 1: The data client sends her budget N to AS.

Pseudocodes of Steps 2, 3, and 4 are presented in Fig. 4.

c) Step 2: It corresponds to everything except the last two lines in Fig. 4(a) and has $1+N-K$ iterations. At each iteration, AS sends to the R_i nodes a bit b_i indicating whether the arm i should be pulled or not. At the first iteration (that corresponds to the initialization phase of UCB cf. Fig. 2), AS sends $b_i = 1$ to each arm, and at the next $N-K$ iterations (that correspond to the exploration-exploitation phase of UCB cf. Fig. 2), AS sends $b_i = 1$ only to a chosen arm i_m and

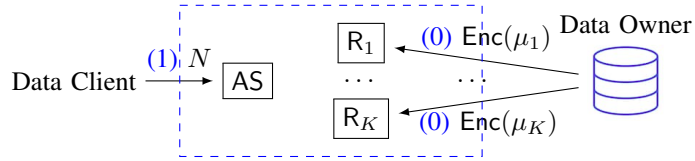

(a) Steps 0 and 1 are done only once at the beginning.

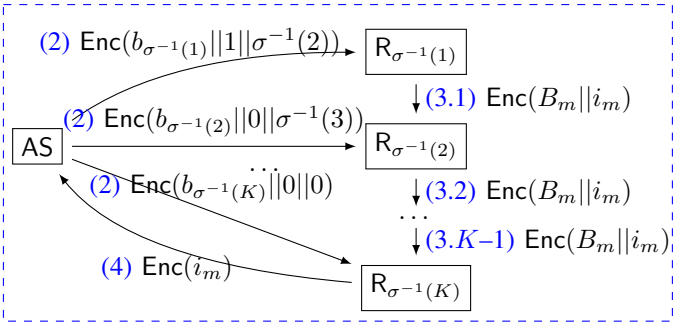

(b) Steps 2, 3, and 4 are done $1 + N - K$ times.

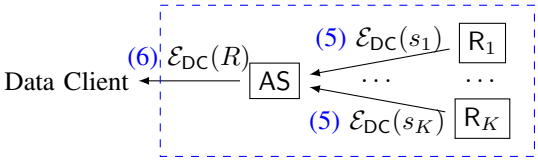

(c) Steps 5 and 6 are done only once at the end.

Fig. 3. Overview of UCB-DS. The dashed rectangle is the cloud.

sends $b_i = 0$ to all other arms. Moreover, at each iteration, AS generates a permutation $\sigma : [\![K]\!] \to [\![K]\!]$ (i.e., a function for which every element occurs exactly once as an image value), based on which AS computes two more components that it sends to R_i: $first_i$ that indicates whether the arm node is the first of the ring hence it should initialize B_m and i_m, and $next_i$ that indicates to which node the updated B_m and i_m should be sent during Step 3. The arm node that receives 0 on the $next$ component is the last one of the ring and sends i_m to AS, which thus knows which arm should be pulled next. All information that AS sends to R_i are thus useful for the ring computation of i_m in Step 3. The permutation changes at each AS iteration because it is important to have a random order during the ring communication. Without a random order, it may happen that the last arm is much better than all others and it is almost always pulled, hence it has a very good estimate of the cumulative reward.

d) Step 3: This step corresponds to everything except the last two lines in Fig. 4(b). Note that the variable t stores how many arm pulls have been done in total since the beginning of the protocol. As discussed for Step 0, each arm initialized $t = K - 1$, hence $t = K$ after the first iteration of AS, which allows to compute the first B_i values at the end of the initialization phase. Then, during the next $N - K$ iterations of AS, the variable t is incremented, which allows to compute B_i values during the exploration-exploitation phase. To decide which arm has the highest B_i and should be pulled at the next iteration, the arm nodes R_i do a distributed ring computation,

let $i_m = 0$
for $j \in [\![N - K + 1]\!]$
  let σ=random permutation of $[\![K]\!]$
  for $i \in [\![K]\!]$
    / b_i is a bit indicating if arm i should be pulled */*
    if $i_m = 0$ **or** $i_m = i$ **then let** $b_i = 1$ **else let** $b_i = 0$
    / $first_i$ is a bit indicating if i is the first arm node in the ring cf. σ */*
    if $\sigma(i) = 1$ **then let** $first_i = 1$ **else let** $first_i = 0$
    / $next_i$ indicates the next arm node in the ring, or 0 if i is the last cf. σ */*
    if $\sigma(i) \neq K$ **then let** $next_i = \sigma^{-1}(\sigma(i) + 1)$ **else let** $next_i = 0$
    send $\mathsf{Enc}(b_i||first_i||next_i)$ to arm node R_i
  receive ciphertext from arm node $\mathsf{R}_{\sigma^{-1}(K)}$
  / ciphertext is* $\mathsf{Enc}(i_m)$ **/*
  let $i_m = \mathsf{Dec}(\text{ciphertext})$

(a) Pseudocode of AS.

receive ciphertext1 from AS
/ ciphertext1 is* $\mathsf{Enc}(b_i||first_i||next_i)$ **/*
let $b_i||first_i||next_i = \mathsf{Dec}(\text{ciphertext1})$
let $t = t + 1$
if $b_i = 1$ */* Pull arm i and update its variables */*
  let $r = pull(i)$
  let $s_i = s_i + r$
  let $n_i = n_i + 1$
let $B_i = \frac{s_i}{n_i} + \sqrt{\frac{2\ln(t)}{n_i}}$
if $first_i = 0$
  receive ciphertext2 from preceding arm node in ring
  / ciphertext2 is* $\mathsf{Enc}(B_m||i_m)$ **/*
  $B_m||i_m = \mathsf{Dec}(\text{ciphertext2})$
if $first_i = 1$ **or** $B_m < B_i$
  let $i_m = i$
  let $B_m = B_i$
if $next_i \neq 0$
  send $\mathsf{Enc}(B_m||i_m)$ to R_{next_i}
else
  send $\mathsf{Enc}(i_m)$ to AS

(b) Pseudocode of R_i, for $i \in [\![K]\!]$.

Fig. 4. Pseudocode of AS and R_i during steps 2, 3, and 4 cf. Fig. 3(b).

where the first arm node according to permutation σ (i.e., the only arm node that received $first_i$=1) initializes max value B_m and argmax i_m. At each ring iteration (Steps 3.1, . . ., 3.K-1, cf. Fig. 3(b)), the current arm node sends updated B_m and i_m to the next arm node cf. σ. Even though B_m and i_m do not change, it is important to re-encrypt $\mathsf{Enc}(B_m||i_m)$ before sending it to the next node to prevent an external observer from knowing when there is a change in the max and argmax (and hence learn information about which arms are pulled more often). Finally, once the ring computation reaches the last arm node relative to σ (i.e., the only one that received $next_i = 0$),

we go to Step 4.

e) Step 4: This step corresponds to the last two lines in Fig. 4(b) (the last arm node in the ring sends $\mathsf{Enc}(i_m)$ to AS), followed by the last two lines in Fig. 4(a) (AS receives and decrypts the index of the arm to be pulled at the next iteration).

f) Step 5: Once the budget is spent and no more arm has to be pulled, each arm node R_i (for $i \in [\![K]\!]$) encrypts with $\mathcal{E}_{\mathsf{DC}}$ its sum of rewards s_i and sends the result $\mathcal{E}_{\mathsf{DC}}(s_i)$ to AS.

g) Step 6: The node AS takes the K ciphertexts $\mathcal{E}_{\mathsf{DC}}(s_i)$ received at Step 5, and computes $\mathcal{E}_{\mathsf{DC}}(R) = \mathcal{E}_{\mathsf{DC}}(\sum_{i=1}^{K} s_i) = \prod_{i=1}^{K}(\mathcal{E}_{\mathsf{DC}}(s_i))$, thanks to the additive homomorphic property of Paillier cryptosystem. Then, AS sends $\mathcal{E}_{\mathsf{DC}}(R)$ to the data client, who is able to decrypt using sk and hence obtains R.

C. Analysis of UCB-DS

Next, we analyze the correctness (Sect. III-C1), security (Sect. III-C2), and complexity (Sect. III-C3) of UCB-DS.

1) Correctness: We point out that UCB-DS outputs exactly the same cumulative reward as UCB. The computations done in Fig. 4 to maximize the reward are the same as the one done in Fig. 2. Indeed, if we take UCB-DS and remove all encryptions/decryptions (both symmetric and asymmetric), and all messages are communicated in clear between participants, then we obtain a protocol that we call UCB-D, which outputs exactly the same result as UCB-DS. This happens because of the consistency property of the chosen cryptographic schemes i.e., if we encrypt a message M using Enc (or $\mathcal{E}_{\mathsf{DC}}$, respectively) to obtain a ciphertext C, then if we decrypt C using Dec (or $\mathcal{D}_{\mathsf{DC}}$, respectively), then we obtain exactly M. Next, to reduce UCB-D to UCB, we simply remove all distributions of tasks among participants and rewrite UCB-D as a sequential algorithm to obtain exactly UCB. In particular, the random permutation σ (that is generated at each round to decide in which order to iterate over arms) reduces to the randomness in the argmax function used in standard UCB cf. Fig. 2 when, if several arms have maximal B_i-value, then the argmax should be randomly picked among those arms.

2) Security: In Table II, we summarize what each participant in UCB-DS knows/does not know. The main properties of our protocol are:

- No cloud node can learn the cumulative reward and additionally:
  - Only AS and the pulled arm know which arm is pulled at each round. Arms that are not pulled can guess the pulled arm with average probability $\frac{1}{2} + \frac{1}{2K}$.
  - Only arm node R_i knows the sum of rewards for arm i.
- Only DC knows the cumulative reward, and she knows nothing else.
- An external observer cannot learn the cumulative reward, the sum of rewards for some arm, or which arm has been pulled at some round.

These properties subsume the list of desirable security properties listed in Sect. III-A. We omit here formal statements and proofs for all these security properties, which are available in [19].

TABLE II
WHAT EACH PARTICIPANT OF UCB-DS KNOWS AND DOES NOT KNOW.

Participant	*Knows*	*Does not know*
AS	• Arm pulled at each round	• Sum of rewards for some arm and cumulative reward
R_i	• Sum of rewards for arm i • Arm pulled at each round, with average probability $\frac{1}{2} + \frac{1}{2K}$	• Sum of rewards of other arm $j \neq i$ and cumulative reward
DC	• Cumulative reward	• Arm pulled at each round • Sum of rewards for some arm
External observer	• Nothing	• Arm pulled at each round • Sum of rewards for some arm and cumulative reward

TABLE III
NUMBER OF CRYPTOGRAPHIC OPERATIONS USED IN UCB-DS.

	Encryptions		*Decryptions*	
AES-CBC	K	(step 0)	K	(step 0)
	$(N-K+1)K$	(step 2)	$(N-K+1)K$	(step 2)
	$(N-K+1)(K-1)$	(step 3)	$(N-K+1)(K-1)$	(step 3)
	$(N-K+1)$	(step 4)	$(N-K+1)$	(step 4)
Paillier	K	(step 5)	1	(step 6)

3) Complexity: We detail in Table III the number of cryptographic operations used in each step of UCB-DS. By summing up, we obtain $O(NK)$ AES-CBC encryptions/decryptions, K Paillier encryptions, and one Paillier decryption. Hence, we have a number of AES-CBC operations linear in N, whereas the number of Paillier operations does not depend on N. These are desirable complexity properties. In particular, the number of Paillier operations (which are quite slow to evaluate in practice) depends only on K that is typically much smaller than N in bandit scenarios. Our implementation (cf. Sect. IV) follows the aforementioned theoretical analysis and confirms the linear time behavior and the scalability of UCB-DS.

D. Refinement

We propose the UCB-DS2 refinement, which adds slightly stronger security guarantees to UCB-DS, for few more cryptographic operations (but the similar asymptotic behavior as UCB-DS). A property of UCB-DS (cf. Table II) is that an arm node R_i knows with average probability of $\frac{1}{2} + \frac{1}{2K}$ what arm is pulled at the next round. This happens because during the ring computation, every arm sees in clear the partial argmax i_m. Our UCB-DS2 refinement removes this leakage.

The idea of UCB-DS2 is that, in addition to UCB-DS, we also encrypt the partial argmax i_m during the ring computation. This modification requires to introduce new keys. We recall that UCB-DS assumes an AES-CBC key that is shared between the data owner and all cloud participants and that is used for the functions $\mathsf{Enc}/\mathsf{Dec}$. For UCB-DS2, if we want that an arm node R_i cannot decrypt the partial argmax i_m received from the previous arm node in the ring, we need to encrypt i_m with some other key. This is why in UCB-DS2 we introduce K new AES-CBC keys, each of them shared between AS and a single R_i arm node. Each such key defines

functions Enc_i/Dec_i. We omit here the pseudocode and the analysis of UCB-DS2 that we include in [19].

IV. Experiments

We show that the overhead due to cryptographic primitives is reasonable, hence our protocols are feasible in practice. More precisely, we show the scalability of our protocols with respect to both parameters N and K through an experimental study using synthetic and real data. We compare:

- UCB = Standard UCB [3], outlined in Fig. 2.
- UCB-D = UCB with distribution of tasks among participants in the spirit of UCB-DS cf. Sect. III-B, but with all messages exchanged in clear among participants (i.e., UCB-D does not use any cryptographic primitive). The only overhead w.r.t. UCB is due to distribution of tasks.
- UCB-DS = Distributed Secure UCB cf. Sect. III-B.
- UCB-DS2 = Refinement of UCB-DS cf. Sect. III-D.

We implemented the algorithms in Python 3. For AES-CBC we used the *PyCryptodome* library[3] and keys of 256 bits. For Paillier, we used the *phe* library[4] in the default configuration with keys of 2048 bits. We did our experiments on a laptop with CPU Intel Core i7 of 2.80GHz and 16GB of RAM, running Ubuntu. Each reported result is averaged over 100 runs. In each run, we executed all algorithms using the same random seeds, needed for drawing arm rewards and for generating the permutation used to iterate in a random order over the arms when choosing the argmax arm to be pulled at the next round.

We make available on a public GitHub repository[5] our source code, together with the data that we used, the generated results from which we obtained our plots, and scripts that allow to install the needed libraries and reproduce our plots.

As expected, in each experiment, all four algorithms output exactly the same cumulative reward. The property that our secure algorithms return exactly the same cumulative reward as standard UCB is in contrast with differentially-private multi-armed bandit algorithms [4], [5], [6], where the returned cumulative rewards are different from that of standard UCB. Consequently, a shallow empirical comparison between these works and ours boils down to comparing apples and oranges: (i) on the one hand, the running time of differentially-private bandit algorithms is roughly the same as for standard UCB and is never reported in their experiments, whereas (ii) on the other hand, for our algorithms the cumulative reward is always the same as for standard UCB and consequently there is no point for us in doing any plot on the cumulative reward. Nevertheless, we carefully analyzed all experimental settings (N, K, μ) used in the related work, that we adapt for our scalability experiments, as we detail next.

a) Scalability with respect to N: In this experiment, we rely on six scenarios from the related work [4], [6] to fix K and μ, and to vary N. In Fig. 5, we show the results only for Scenario 1 [4]. We omit here the other scenarios, which yield similar results, included in [19]. We vary N from 10^2 to 10^5 that is also the maximum budget from [4], [6]. UCB and UCB-D have very close running times, and up to two orders of magnitude smaller than UCB-DS and UCB-DS2, which are also very close. All algorithms have a similar linear time behavior. The overhead between secure and non-secure algorithms comes naturally from the cryptographic primitives. Moreover, the two lines corresponding to the secure algorithms are not parallel with the other two lines because, cf. Sect. III-C3, the overhead due to Paillier encryptions depends only on K (that is fixed in the figure) and not on N (that varies in the figure), hence the Paillier overhead is more visible for small N. The running times of UCB-DS/UCB-DS2 for the largest considered budget N=10^5 is of $\sim$100 seconds, which remains practical. In Fig. 5, we also zoom on the time taken by each participant of UCB-DS for N=10^5. We observe that AS takes the lion's share, which is expected because at each round AS sends encrypted messages to all R_i participants, whereas each R_i sends an encrypted message only to one other participant. As expected, all R_i take roughly the same time. The shares taken by the data owner and the data client are the smallest among all participants, which is a desirable property because we require them to do as few computations as possible, whereas the bulk of the computation is outsourced to the cloud.

b) Scalability with respect to K: In this experiment, we fix N=10^5, and we vary K∈$\{5, 10, 15, 20\}$ and implicitly μ with μ_1=0.9 and $\mu_{2\leq i\leq K}$=0.8. We present results in Fig. 6. We observe, as in the previous experiment, a linear time behavior and a similar zoom on the time taken by each participant.

c) Real-world data: We also stress-tested our algorithms on real-world data, using the same data and setup as [20]. After a pre-processing step (detailed in [19]), we transformed real-world data in three bandit scenarios: Jester-small (K=10) and Jester-large (K=100) based on Jester dataset [21], and MovieLens (K=100) based on MovieLens dataset [22]. We ran each of these scenarios with N=10^5 that is the largest budget considered in [20]. Our results (cf. Fig. 7) essentially confirm the behavior observed in the synthetic experiments i.e., there are roughly two orders of magnitude between non-secure and secure algorithms. In the largest considered scenarios (Jester-large and MovieLens, both with K=100), where standard UCB takes around twenty seconds, both UCB-DS and UCB-DS2 take around one thousand seconds, that we believe acceptable as waiting time for the data client before getting the cumulative reward result for which she pays.

V. Conclusions and Future Work

We tackled the problem of cumulative reward maximization in multi-armed bandits, in a setting where data and computations are outsourced to some honest-but-curious cloud. We proposed UCB-DS, a distributed and secure protocol based on UCB, which yields exactly the same cumulative reward as UCB while enjoying desirable security properties that we precisely characterize. In particular, no cloud node or external observer can learn the cumulative reward, which can be seen only by the data client who pays a budget. We rely on

[3]https://pycryptodome.readthedocs.io/en/latest/src/cipher/classic.html

[4]https://python-paillier.readthedocs.io/en/develop/

[5]https://github.com/radu1/secure-ucb

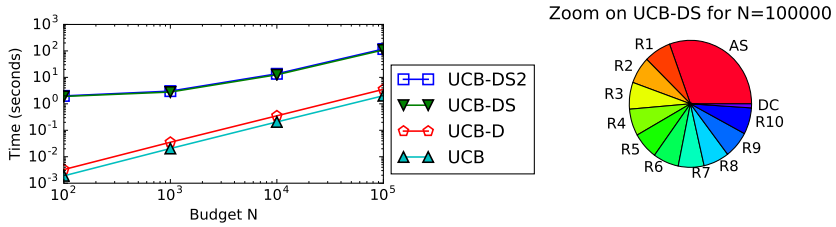

Fig. 5. Scalability with respect to N, for fixed $K = 10$, $\mu_1 = 0.9$, $\mu_2 = \ldots = \mu_{10} = 0.8$. In the zoom, we do not show DO (its share is close to 0).

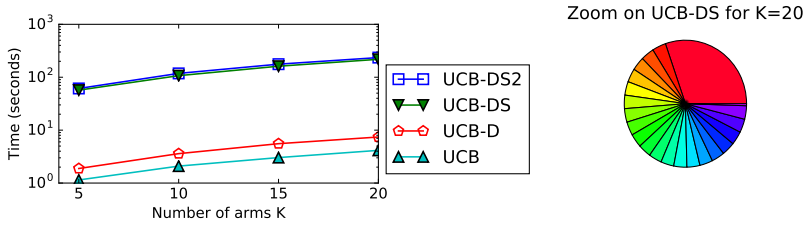

Fig. 6. Scalability with respect to K, for fixed $N = 10^5$. In the zoom (labels not shown because they would be colliding): the AS takes the lion's share, $\mathsf{R}_{1 \leq i \leq 20}$ take the 20 equal shares, DC is barely visible, and DO is not shown since its share is close to 0.

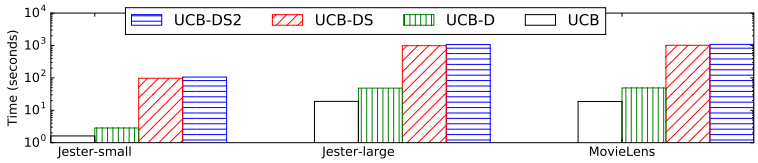

Fig. 7. Running times on three real-world data scenarios from [20].

distribution of tasks and on cryptographic schemes to achieve the security properties of UCB-DS, and we characterize the overhead of cryptography from both theoretical and empirical points of view. Our experiments show the scalability and practical feasibility of UCB-DS, and of its refinement UCB-DS2.

As future work, we plan to extend our scenario such that multiple data clients concurrently submit budgets to the cloud and receive corresponding cumulative rewards. In such a scenario, parallelism between nodes could be leveraged to improve the system's throughput.

REFERENCES

[1] S. Bubeck and N. Cesa-Bianchi, "Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems," *Foundations and Trends in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.

[2] F. Bourse, M. Minelli, M. Minihold, and P. Paillier, "Fast Homomorphic Evaluation of Deep Discretized Neural Networks," in *CRYPTO*, 2018, pp. 483–512.

[3] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time Analysis of the Multiarmed Bandit Problem," *Machine Learning*, vol. 47, no. 2-3, pp. 235–256, 2002.

[4] P. Gajane, T. Urvoy, and E. Kaufmann, "Corrupt Bandits for Preserving Local Privacy," in *ALT*, 2018, pp. 387–412.

[5] N. Mishra and A. Thakurta, "(Nearly) Optimal Differentially Private Stochastic Multi-Arm Bandits," in *UAI*, 2015, pp. 592–601.

[6] A. C. Y. Tossou and C. Dimitrakakis, "Algorithms for Differentially Private Multi-Armed Bandits," in *AAAI*, 2016, pp. 2087–2093.

[7] C. Dwork, "Differential Privacy," in *ICALP*, 2006, pp. 1–12.

[8] C. Gentry, "Fully Homomorphic Encryption Using Ideal Lattices," in *STOC*, 2009, pp. 169–178.

[9] J. H. Cheon, A. Kim, M. Kim, and Y. S. Song, "Homomorphic Encryption for Arithmetic of Approximate Numbers," in *ASIACRYPT*, 2017, pp. 409–437.

[10] "Advanced Encryption Standard (AES)," https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.197.pdf, 2001, fIPS Publication 197.

[11] M. Bellare, A. Desai, E. Jokipii, and P. Rogaway, "A Concrete Security Treatment of Symmetric Encryption," in *FOCS*, 1997, pp. 394–403.

[12] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *EUROCRYPT*, 1999, pp. 223–238.

[13] R. Ciucanu, P. Lafourcade, M. Lombard-Platet, and M. Soare, "Secure Best Arm Identification in Multi-Armed Bandits," in *ISPEC*, 2019, pp. 152–171.

[14] J. Audibert, S. Bubeck, and R. Munos, "Best Arm Identification in Multi-Armed Bandits," in *COLT*, 2010, pp. 41–53.

[15] R. Ciucanu, A. Delabrouille, P. Lafourcade, and M. Soare, "Secure Cumulative Reward Maximization in Linear Stochastic Bandits," in *ProvSec*, 2020.

[16] R. Agrawal, "Sample Mean Based Index Policies with $O(\log(n))$ Regret for the Multi-Armed Bandit Problem," *Advances in Applied Probability*, vol. 27, no. 4, pp. 1054–1078, 1995.

[17] R. Munos, "From Bandits to Monte-Carlo Tree Search: The Optimistic Principle Applied to Optimization and Planning," *Foundations and Trends in Machine Learning*, vol. 7, no. 1, pp. 1–129, 2014.

[18] O. Goldreich, *The Foundations of Cryptography - Volume 2: Basic Applications*. Cambridge University Press, 2004.

[19] R. Ciucanu, P. Lafourcade, M. Lombard-Platet, and M. Soare, "Secure Outsourcing of Multi-Armed Bandits," in *TR at https://hal.inria.fr/ /hal-02953292*, 2020.

[20] P. Kohli, M. Salek, and G. Stoddard, "A Fast Bandit Algorithm for Recommendation to Users With Heterogenous Tastes," in *AAAI*, 2013.

[21] K. Y. Goldberg, T. Roeder, D. Gupta, and C. Perkins, "Eigentaste: A Constant Time Collaborative Filtering Algorithm," *Information Retrieval*, vol. 4, no. 2, pp. 133–151, 2001.

[22] F. M. Harper and J. A. Konstan, "The MovieLens Datasets: History and Context," *ACM TiiS*, vol. 5, no. 4, pp. 19:1–19:19, 2016.

APPENDIX A
SECURITY PROOFS

In Sect A-A, we detail the definition of the security tools from Sect. II. Then, we prove the security of UCB-DS and UCB-DS2 in Sect A-B and A-C, respectively.

A. Additional Information on Security Tools

In this section, we detail the definition of the security tools briefly introduced in Sect. II in order to provide enough background to formally prove the security of our protocols. Before introducing the two cryptographic schemes, we point out that each of them has a *security parameter* λ that is input to key generation. By 1^λ we denote the unary representation of λ, which is a standard notation in cryptography. Our security theorems are always asymptotic i.e., they describe the behavior when λ becomes infinitely large. In practice, the security parameter is the length of the keys, for both Paillier and AES-CBC.

a) Paillier asymmetric encryption: Paillier's cryptosystem [12] is an asymmetric partial homomorphic encryption scheme defined by a triple of polynomial-time algorithms $(\mathcal{G}, \mathcal{E}, \mathcal{D})$ and a security parameter λ such that:

- $\mathcal{G}(1^\lambda)$ generates two prime numbers p and q according to λ, sets $n = p \cdot q$ and $\Lambda = \mathrm{lcm}(p-1, q-1)$ (i.e., the least common multiple), generates the group $(\mathbb{Z}^*_{n^2}, \cdot)$, randomly picks $g \in \mathbb{Z}^*_{n^2}$ such that $M = (L(g^\Lambda \mod n^2))^{-1} \mod n$ exists, with $L(x) = (x-1)/n$. It sets $\mathsf{sk} = (\Lambda, M)$, $\mathsf{pk} = (n, g)$, it returns $(\mathsf{sk}, \mathsf{pk})$.
- $\mathcal{E}_{\mathsf{pk}}(m)$ randomly picks $r \in \mathbb{Z}^*_n$, computes $c = g^m \cdot r^n \mod n^2$, and outputs c.
- $\mathcal{D}_{\mathsf{sk}}(c)$ computes $m = L(c^\Lambda \mod n^2) \cdot M \mod n$, and outputs m.

Paillier's cryptosystem is *additive homomorphic*. Let m_1 and m_2 be two plaintexts in $\mathbb{Z}_n$. The product of the two associated ciphertexts with the public key $\mathsf{pk} = (n, g)$, denoted $c_1 = \mathcal{E}_{\mathsf{pk}}(m_1) = g^{m_1} \cdot r_1^n \mod n^2$ and $c_2 = \mathcal{E}_{\mathsf{pk}}(m_2) = g^{m_2} \cdot r_2^n \mod n^2$, is the encryption of the sum of m_1 and m_2. Indeed, we have:

$$\begin{aligned}
\mathcal{E}_{\mathsf{pk}}(m_1) \cdot \mathcal{E}_{\mathsf{pk}}(m_2) &= c_1 \cdot c_2 \mod n^2 \\
&= (g^{m_1} \cdot r_1^n) \cdot (g^{m_2} \cdot r_2^n) \mod n^2 \\
&= (g^{m_1+m_2} \cdot (r_1 \cdot r_2)^n) \mod n^2 \\
&= \mathcal{E}_{\mathsf{pk}}(m_1 + m_2)\ .
\end{aligned}$$

b) AES-CBC symmetric encryption: AES [10] is a NIST standard for symmetric encryption that encrypts messages of 128 bits. AES is used as a block cipher, for instance using CBC mode (Cipher Block Chaining). The AES-CBC cryptosystem is a symmetric encryption scheme defined by a triple of polynomial-time algorithms $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ and a security parameter λ such that:

- $\mathsf{KeyGen}(1^\lambda)$ generates Key, a uniformly random symmetric key of 128, 192 or 256 bits, according to λ.
- $\mathsf{Enc}(m)$ splits m in blocks of 128 bits $m_1, \ldots, m_n$ (padding bits may be added if m_n is smaller than 128 bits). By $x \oplus y$ we denote the standard bit-wise xor operation between two numbers x and y. Then, Enc computes $c_i = \mathsf{E}(\mathsf{Key}, m_i \oplus c_{i-1})$ for $1 \leq i \leq n$ where $c_0 = IV$ is a random 128-bits number and E is the AES encryption [10]. Then Enc returns the tuple $(c_0, \ldots, c_n)$.
- $\mathsf{Dec}(c)$ splits c in blocks of 128 bits $c_0, \ldots, c_n$ and computes $m_i = \mathsf{D}(\mathsf{Key}, c_i) \oplus c_{i-1}$ for $1 \leq i \leq n$ where $c_0 = IV$ and D is the AES decryption [10]. Then, $\mathsf{Dec}(c)$ returns $m = (m_1, \ldots, m_n)$.

c) IND-CPA (INDistinguishability under Chosen-Plaintext Attack) [11]: Let Π= (KeyGen, Encrypt, Decrypt) be a cryptographic scheme. The *probabilistic polynomial-time (PPT) adversary* $\mathcal{A}$ tries to break the security of Π. The IND-CPA game, denoted by $\mathrm{EXP}(\mathcal{A})$, works as follows: the adversary $\mathcal{A}$ chooses two messages (m_0, m_1) and receives a challenge $c = \mathsf{Encrypt}(LR_b(m_0, m_1))$ from the *challenger* who selects a bit $b \in \{0, 1\}$ uniformly at random, and where $LR_b(m_0, m_1)$ is equal to m_0 if b=0, and m_1 otherwise. The adversary, knowing m_0, m_1 and c, is allowed to perform any number of polynomial computations or encryptions of any messages, using the encryption oracle, in order to output a guess b' of the encrypted message in c chosen by the challenger. Intuitively, Π is IND-CPA if there is no PPT adversary that can guess b with a probability significantly better than $\frac{1}{2}$. By $\alpha = \Pr[b' \leftarrow \mathrm{EXP}(\mathcal{A}); b = b']$, we denote the probability that $\mathcal{A}$ correctly outputs her guessed bit b' when the bit chosen by the challenger in the experiment is b. A scheme is IND-CPA secure if $\alpha - \frac{1}{2}$ is negligible function in λ, where a function γ is negligible in λ, denoted $\mathrm{negl}(\lambda)$, if for every positive polynomial $p(\cdot)$ and sufficiently large λ, $\gamma(\lambda) < 1/p(\lambda)$.

Both cryptographic schemes mentioned earlier in this section are IND-CPA: (i) Paillier is IND-CPA under the decisional composite residuosity assumption [12], and (ii) AES-CBC is IND-CPA under the assumption that AES is a pseudo-random permutation [11].

In our theorems, the notion of "better than random" is consistent with the aforementioned IND-CPA property. We also point out an additional notation used in the proofs. Similarly to Landau Big O notation, where by convention $O(f)$ can describe any function bounded above by f, we abuse notation and denote by $\mathrm{negl}(\lambda)$ any function negligible in λ. Notably, we have $\mathrm{negl}(\lambda) + \mathrm{negl}(\lambda) = \mathrm{negl}(\lambda)$ and we may write $x + \mathrm{negl}(\lambda)$ instead of $x - \mathrm{negl}(\lambda)$.

TABLE IV
WHAT EACH PARTICIPANT OF UCB-DS KNOWS AND DOES NOT KNOW, WITH POINTERS TO THE RELEVANT THEOREMS.

Participant	*Knows*	*Does not know*
AS	• Arm pulled at each round	• Sum of rewards for some arm and cumulative reward (Th. 1)
R_i	• Sum of rewards for arm i • Arm pulled at each round, with average probability $\frac{1}{2} + \frac{1}{2K}$ (Th. 2)	• Sum of rewards of other arm $j \neq i$ and cumulative reward (Th. 3)
DC	• Cumulative reward	• Arm pulled at each round (Th. 4) • Sum of rewards for some arm (Th. 5)
External observer	• Nothing	• Arm pulled at each round (Th. 6) • Sum of rewards for some arm and cumulative reward (Th. 7)

B. Security Proofs for Sect. III-C2

In this section, we provide formal statements and proofs for the security properties of UCB-DS that we have already outlined in Sect. III-C2. In Table IV, we summarize what each participant in UCB-DS knows/does not know, with pointers to the relevant theorems.

Before formally stating the theorems, we point out some assumptions. First, we recall (cf. Sect. III-A) that the participants are honest-but-curious and do not collude. We discuss the impact of collusions at the end of this section. Second, during the ring computation (cf. Step 3 in Sect. III-B), each arm learns an intermediate max value B_m, together with intermediate arm argmax i_m; we assume that the knowledge on intermediate B_m and i_m by each arm does not leak significant information on the sum of rewards. Our refinement UCB-DS2 (cf. Sect. III-D) hides i_m during the ring computation to relax the second hypothesis.

Before discussing the security properties for each participant, we introduce some notations needed for the theorem statements:

- $n_{i,t}$ = the number of times arm i has been pulled until round t.
- $s_{i,t}$ = the sum of rewards obtained by arm i until round t.
- $data^t_{\mathsf{A}}$ = the data to which participant A has access until round t, where A can be a participant from Fig. 3 or the external observer (ext). If t is omitted, this denotes the data to which A has access at the end of the protocol.
- $\mathcal{A}^{pb(.)}(d)$ = the answer of a Probabilistic Polynomial-Time (PPT) adversary $\mathcal{A}$ that knows d and tries to solve the problem pb. Depending on the problem, pb can also take some *input*.
- By *negligible in* λ, we denote that our security theorems are always asymptotic i.e., they describe the behavior when the security parameter λ of the cryptographic schemes becomes infinitely large.

We next provide theorems that state each non-trivial property from Table IV. We first state an useful lemma, which intuitively says that guessing the cumulative reward with probability better than random is equivalent to guessing the sum of rewards of some arm with probability better than random.

Lemma 1. *Let $\mathcal{A}$ be a PPT adversary trying to find the cumulative reward R, and let $\mathcal{B}$ be a PPT adversary trying to find the sum of rewards of some arm. Let d be some data, $cr(.)$ be the problem of guessing the cumulative reward, and $sum(.)$ be the problem of guessing the sum of rewards of some arm. We have the following statement: $\mathcal{A}^{cr(.)}(d)$ has a non-negligible advantage $\Leftrightarrow$ $\mathcal{B}^{sum(.)}(d)$ has a non-negligible advantage.*

Proof. $\Leftarrow$ Assume that $\mathcal{B}$ can guess the sum of rewards of some arm with probability better than random. Then, $\mathcal{A}$ can call $\mathcal{B}$, and hence get the sum of rewards of one arm with probability better than random. From this sum, $\mathcal{A}$ can guess a lower bound on the cumulative reward, hence eliminating some possibilities, and thus guessing the cumulative reward with probability better than random.

$\Rightarrow$ If $\mathcal{A}$ can guess the cumulative reward with probability better than random, then $\mathcal{B}$ can use this cumulative reward as an upper bound on the sum of rewards of some arm, thus having a probability better than random of guessing the sum of rewards of some arm. □

a) Security of AS*:* By construction of UCB-DS, AS knows the arm pulled at each round. We state that AS cannot learn the sum of rewards produced by some arm.

Theorem 1. *For an arm $i \in [\![K]\!]$ and a round $t \in [\![N-K+1]\!]$, an honest-but-curious* AS *cannot learn $s_{i,t}$, given $data^t_{\mathsf{AS}}$, with a probability better than random. More precisely, for all PPT adversaries $\mathcal{A}$,*

$$\left| \Pr\left[(i, \hat{s}_{i,t}) \leftarrow \mathcal{A}^{sum(t)}(data^t_{\mathsf{AS}}); \hat{s}_{i,t} = s_{i,t} \right] - p_S(n_{i,t}, s_{i,t}) \right|$$

is negligible in λ, where $\mathcal{A}^{sum(t)}(data^t_{\mathsf{AS}})$ returns $(i, \hat{s}_{i,t})$ in which $\hat{s}_{i,t}$ is $\mathcal{A}$'s guess on $s_{i,t}$ for the arm i (chosen by $\mathcal{A}$), and $p_S(n_{i,t}, s_{i,t})$ is the probability of obtaining a sum of rewards $s_{i,t}$ from $n_{i,t}$ pulls of arm i until round t.

Proof. Before Step 5 of UCB-DS, AS has access at each round to the indices of the pulled arms. Thus, AS knows $n_{i,t}$ i.e., the number of times the arm i has been pulled until round t. The set of all possible sums of rewards for arm i until round t is $\{0, 1, \ldots, n_{i,t}\}$. We denote by $p_S(n_{i,t}, s_{i,t})$ the probability of obtaining the sum of rewards $s_{i,t}$ from $n_{i,t}$ pulls of arm i until round t. Next, we show that the advantage of AS based of $data^t_{\mathsf{AS}}$ is $p_S(n_{i,t}, s_{i,t})$ plus an amount negligible in λ.

Since AS has no knowledge on μ_i, the property stated in the theorem is respected at each round before Step 5, i.e., for all $t < N - K + 1$.

We next prove the property for the last round i.e., $t = N - K + 1$. At the end of UCB-DS, at Step 5, AS receives the values $\mathcal{E}_{\mathsf{DC}}(s_{1,t}), \ldots, \mathcal{E}_{\mathsf{DC}}(s_{K,t})$. We prove that retrieving any information about any $s_{i,t}$ from these ciphertexts breaks the IND-CPA property of Paillier's cryptosystem [12]. At this point of UCB-DS, $data^t_{\mathsf{AS}}$ consists of $\mathcal{E}_{\mathsf{DC}}(s_{1,t}), \ldots, \mathcal{E}_{\mathsf{DC}}(s_{K,t})$ and the list of arms that have been pulled at each round. Assume there exists a PPT adversary $\mathcal{A}$ able, from $data^t_{\mathsf{AS}}$ to find $s_{i,t}$ for some i with non negligible advantage x:

$$\left| \Pr\left[(i, \hat{s}_{i,t}) \leftarrow \mathcal{A}^{sum(t)}(data^t_{\mathsf{AS}}); \hat{s}_{i,t} = s_{i,t} \right] - p_S(n_{i,t}, s_{i,t}) \right| = x + \text{negl}(\lambda).$$

In the worst case, each $i \in [\![K]\!]$ has an equal probability of being chosen by $\mathcal{A}$. We also assume that if $data^t_{\mathsf{AS}}$ does not correspond to the data collected by AS during a run of UCB-DS (for instance, if one piece of $data^t_{\mathsf{AS}}$ has been replaced by another unrelated message), then $\mathcal{A}$ does not give any advantage. If such an adversary $\mathcal{A}$ exists, then we show how to construct an adversary $\mathcal{B}$ able to break the IND-CPA property of Paillier.

Let us build an IND-CPA game, in which $\mathcal{B}$ chooses two values m_0, m_1, and sends them to the challenger. The challenger randomly selects $b \in \{0, 1\}$ and answers with $\mathcal{E}_{\mathsf{DC}}(m_b)$. $\mathcal{B}$ wins the IND-CPA game if $\mathcal{B}$ guesses b with a non-negligible advantage.

To do so, $\mathcal{B}$ first creates a simulation of an UCB-DS execution i.e., $\mathcal{B}$ creates nodes DC', AS', R'_i, and DO', with Bernoulli distributions defined by μ'_i of its choice. Then, $\mathcal{B}$ runs an execution of UCB-DS on these nodes. Because $\mathcal{B}$ controls all the nodes, it knows the sums of rewards $s'_{1,t}, \ldots, s'_{K,t}$, as well as a list L of arms pulled at each round.

As input for the IND-CPA game, $\mathcal{B}$ chooses $m_1 = s'_{1,t}$ and another value m_0, different from all $s'_{i,t}$, sends both values to the challenger, and receives $\mathcal{E}_{\mathsf{DC}}(m_b)$. Then, $\mathcal{B}$ computes $\mathcal{E}_{\mathsf{DC}}(s'_{i,t})$ for each i, and calls $\mathcal{A}^{sum(t)}([\mathcal{E}_{\mathsf{DC}}(m_b), \mathcal{E}_{\mathsf{DC}}(s'_{2,t}), \ldots, \mathcal{E}_{\mathsf{DC}}(s'_{K,t}), L])$.

The strategy of $\mathcal{B}$ is as follows: if $\mathcal{A}$ returns $(1, m_1)$, then $\mathcal{B}$ answers 1. Otherwise, $\mathcal{B}$ answers randomly. We next derive the probability of a correct answer by $\mathcal{B}$.

- If $i \neq 1$ (probability $1 - \frac{1}{K}$), then $\mathcal{B}$ answers randomly and is correct with probability $\frac{1}{2}$. Hence this branch offers a probability of success of $(1 - \frac{1}{K})\frac{1}{2}$.
- If $i = 1$ (probability $\frac{1}{K}$), let us consider the value of b.
  - If $b = 0$ (probability $\frac{1}{2}$), then we have two cases:
    - If the output of $\mathcal{A}$ is $(1, m_1)$ (probability $p_S(n_{1,t}, s_{1,t})$), then $\mathcal{B}$ answers 1 and it is wrong, hence the probability of success is 0.
    - Otherwise (probability $1 - p_S(n_{1,t}, s_{1,t})$), $\mathcal{B}$ answers randomly and is correct with probability $\frac{1}{2}$. The probability of success of this branch is $\frac{1}{K}\frac{1}{2}(1 - p_S(n_{1,t}, s_{1,t}))\frac{1}{2}$.
  - If $b = 1$ (probability $\frac{1}{2}$), then we have two cases:
    - If the output of $\mathcal{A}$ is $(1, m_1)$ (probability $p_S(n_{1,t}, s_{1,t}) + x + \text{negl}(\lambda)$), then $\mathcal{B}$ correctly answers 1. The probability of success of this branch is $\frac{1}{K}\frac{1}{2}(p_S(n_{1,t}, s_{1,t}) + x + \text{negl}(\lambda))$.
    - Otherwise (probability $1 - p_S(n_{1,t}, s_{1,t}) - x - \text{negl}(\lambda)$), $\mathcal{B}$ answers randomly and is correct with probability $\frac{1}{2}$. The probability of success of this branch is $\frac{1}{K}\frac{1}{2}(1 - p_S(n_{1,t}, s_{1,t}) - x - \text{negl}(\lambda))\frac{1}{2}$.

By aggregating the aforementioned cases, the probability α of success of $\mathcal{B}$ is:

$$\begin{aligned}
\alpha =& \left(1 - \frac{1}{K}\right)\frac{1}{2} + \frac{1}{K}\frac{1}{2}\left(1 - p_S(n_{1,t}, s_{1,t})\right)\frac{1}{2} + \frac{1}{K}\frac{1}{2}\left(p_S(n_{1,t}, s_{1,t}) + x + \text{negl}(\lambda)\right) \\
&+ \frac{1}{K}\frac{1}{2}\left(1 - p_S(n_{1,t}, s_{1,t}) - x - \text{negl}(\lambda)\right)\frac{1}{2} \\
=& \frac{1}{2} - \frac{1}{2K} + \frac{1}{4K} - \frac{p_S(n_{1,t}, s_{1,t})}{4K} + \frac{p_S(n_{1,t}, s_{1,t})}{2K} + \frac{x}{2K} + \frac{1}{4K} - \frac{p_S(n_{1,t}, s_{1,t})}{4K} - \frac{x}{4K} + \text{negl}(\lambda) \\
=& \frac{1}{2} + \frac{x}{4K} + \text{negl}(\lambda)
\end{aligned}$$

Hence, $\mathcal{B}$ has an advantage of $\frac{x}{4K}$ in the IND-CPA game, which is non negligible. This is a contradiction with the fact that Paillier is IND-CPA secure. Consequently, there does not exist any PPT adversary $\mathcal{A}$ that violates the property stated in the theorem. □

As a corollary, by Lemma 1 and Theorem 1, we infer that AS cannot learn the cumulative reward with probability better than random.

b) Security of R_i*:* By construction of UCB-DS, each arm node R_i knows its sum of rewards. Moreover, due to the properties of the ring computation, R_i knows with average probability $\frac{1}{2} + \frac{1}{2K}$ the arm to be pulled at the next round (Theorem 2), but it cannot learn the sum of rewards of any other arm (Theorem 3).

Theorem 2. *At the end of round $t \in [\![N-K]\!]$ and before the start of round $t+1$, given $data^t_{\mathsf{R}_i}$, the average probability that an honest-but-curious R_i guesses the arm to be pulled at round $t+1$ is $\frac{1}{2} + \frac{1}{2K}$.*

Proof. After round t, an arm i can either guess randomly (with a success probability of $\frac{1}{K}$), or use the data to which it has access: the partial max B_m, the partial argmax index i_m, and the next arm in the ring communication. The knowledge of the next arm is useless, as it does not bring any information about any B value. Similarly, the knowledge of B_m does not leak more information than i_m. Hence, the only useful piece is i_m. Based on this only useful piece of data and on the earlier assumption that any information derived from partial argmax data from the previous rounds is negligible, we infer that the best policy for the arm is to bet that arm i_m is the arm to be pulled at the next round. Let us consider an arm at position $\sigma^{-1}(i)$, where σ is the ring permutation used at the round t. Its guess is correct if and only if the next arm to be selected, say j, has position $\sigma^{-1}(j) \leq \sigma^{-1}(i)$. Hence, the arm at position $\sigma^{-1}(i)$ has a success probability of $\frac{\sigma^{-1}(i)}{K}$. On average, an arm has a success probability of

$$\frac{1}{K}\sum_{i=1}^{K}\frac{\sigma^{-1}(i)}{K} = \frac{1}{K^2}\frac{K(K+1)}{2} = \frac{K+1}{2K} = \frac{1}{2} + \frac{1}{2K}$$

which concludes the proof. □

Theorem 3. *For an arm $i \in [\![K]\!]$ and a round $t \in [\![N-K+1]\!]$, an honest-but-curious R_i cannot learn $s_{j,t}$ for some other arm $j \neq i$, given $data^t_{\mathsf{R}_i}$, with a probability better than random. More precisely, for all PPT adversaries $\mathcal{A}$,*

$$\left|\Pr\left[(j, \hat{s}_{j,t}) \leftarrow \mathcal{A}^{sum(t)}(data^t_{\mathsf{R}_i}); \hat{s}_{j,t} = s_{j,t}\right] - p_R(n_{i,t}, t, s_{j,t})\right|$$

is negligible in λ, where $\mathcal{A}^{sum(t)}(data^t_{\mathsf{R}_i})$ returns a tuple $(j, \hat{s}_{j,t})$ in which $j \neq i$ is chosen by $\mathcal{A}$ and $\hat{s}_{j,t}$ is $\mathcal{A}$'s guess of the sum of rewards for arm j, and $p_R(n_{i,t}, t, s_{j,t})$ is the probability of arm j to have sum of rewards $s_{j,t}$ at round t seen that arm i has been pulled $n_{i,t}$ times.

Proof. If an arm i has been pulled $n_{i,t}$ times until round t, then another arm j has been pulled at most $t - n_{i,t}$ times. Hence, a baseline probability of R_i to guess the sum of rewards of any other arm j is the $p_R(n_{i,t}, t, s_{j,t})$ defined in the theorem statement. The arm node R_i cannot possibly guess the sum of rewards for arm j with a better probability because it does not see any useful information that it can leverage. In particular, the only information that R_i receives about the rewards of any other arm is the partial max value B_m (derived from the sum of arm i_m using the number of pulls of i_m, to which R_i does not have access) received during Step 3. As mentioned earlier, we assume that the information that one arm can derive from one such random B value does not provide any advantage. □

As a corollary, by Lemma 1 and Theorem 3, we infer that R_i cannot learn the cumulative reward with probability better than random.

c) Security of DC*:* The data client knows the cumulative reward that she can decrypt after Step 6. Moreover, the data client cannot learn the arms selected at some round (Theorem 4) or the sum of rewards for some arm (Theorem 5).

Theorem 4. *For each round $t \in \{2, \ldots, N-K+1\}$, the data client* DC *cannot guess which arm is pulled at round t with probability better than random.*

Proof. The data client does not receive any message until the end of UCB-DS (Step 6). By construction of UCB-DS, all arms are pulled at the first round, then from round 2 and until the end of UCB-DS i.e., round $N-K+1$, there is a single arm pulled at each round. In particular, the data client does not receive any information on which arm is pulled at some round, hence her best strategy is to answer randomly, with a probability of success of $\frac{1}{K}$. □

Theorem 5. *For an arm $i \in [\![K]\!]$, the data client* DC *cannot guess the sum s_i of rewards for the arm i with probability better that random.*

Proof. Similarly to the previous proof, we observe that the data client DC does not receive any message until the end of UCB-DS (Step 6). In particular, DC does not get any information about which arm is selected at some round. Because all arm probability distributions are equiprobable to DC, it is also true that all partitions of the cumulative reward R are equiprobable to DC, thus DC has no advantage in guessing the partition of rewards. Hence, the the probability of DC guessing a correct partition of the rewards is equal to $\frac{1}{p(R)}$, where $p(R)$ is the number of partitions of R. This observation also proves that DC cannot guess the individual sum of rewards of some arm i. If it was the case, then DC would know that some of the partitions are more likely e.g., if DC can guess the sum of rewards s_i of the arm i, then all partitions not having s_i as the value for arm i would be discarded, which is a contradiction. □

d) External observer: An external observer sees all messages exchanged between nodes, from which we show that she cannot learn which arm is pulled at some round (Theorem 6) or the sum of rewards for some arm (Theorem 7).

Theorem 6. *For each round $t \in \{2, \ldots, N-K+1\}$, an honest-but-curious external observer cannot learn which arm is pulled at round t, given $data^t_{ext}$, with probability better than random. More precisely, for all PPT adversaries $\mathcal{A}$,*

$$\left|\Pr[\mathcal{A}^{pa(t)}(data^t_{ext}) = i^t_m] - \frac{1}{K}\right| \textit{ is negligible in } \lambda,$$

where $\mathcal{A}^{pa(t)}(data^t_{ext})$ returns the guess of $\mathcal{A}$ on which arm is pulled at round t, and i^t_m is the true arm pulled at round t.

Proof. By construction of UCB-DS, all arms are pulled at the first round, then from round 2 and until the end of UCB-DS i.e., round $N-K+1$, there is a single arm pulled at each round. We next show that if there exists a PPT adversary with a non negligible advantage in guessing the arm pulled at some round $2 \leq t \leq N-K+1$, then this would break the IND-CPA property of AES-CBC.

An external observer (denoted ext in the sequel) sees all encrypted messages that are exchanged among UCB-DS participants. We denote by $data^t_{ext}$ this collection of data after round t. We assume, toward a contradiction, that there exists a PPT adversary $\mathcal{A}$ able from $data_{ext}$ to find the arm i^t_m pulled at some round t with a non negligible advantage x:

$$\left|\Pr[\mathcal{A}^{pa(t)}(data^t_{ext}) = i^t_m] - \frac{1}{K}\right| = x + \text{negl}(\lambda).$$

We also assume that if $data^t_{ext}$ does not correspond to an actual collection of encrypted messages that ext sees, then the advantage for such an input is negligible.

We next show that by using the adversary $\mathcal{A}$, we can construct an adversary $\mathcal{B}$ able to break the IND-CPA property of AES-CBC. To do so, $\mathcal{B}$ creates a simulation of an UCB-DS execution, similarly to the proof of Theorem 1. Even though the messages of such a simulation are encrypted, $\mathcal{B}$ knows the keys hence the state of each arm. In particular, $\mathcal{B}$ knows in plain text the message sent by AS to the arm pulled at round t. This message is of the form $m_1 = (1||first_{i,t}||next_{i,t})$, with 1 being the Boolean value saying the arm has to be pulled.

As input for the IND-CPA game, $\mathcal{B}$ sends the aforementioned m_1 and another message $m_0 = (0||first_{i,t}||next_{i,t})$ that it generates based on m_1. Then, $\mathcal{B}$ receives back $\mathsf{Enc}(m_b)$, where b is a random bit selected uniformly by the challenger. Next, $\mathcal{B}$ calls $\mathcal{A}^{pa(t)}(data'_{ext})$, where $data'_{ext}$ is the collection of encrypted messages from the $\mathcal{B}$'s simulation, except that it replaces $\mathsf{Enc}(m_1)$ by $\mathsf{Enc}(m_b)$. The strategy of $\mathcal{B}$ is: if $\mathcal{A}$ returns the correct i^t_m, then $\mathcal{B}$ returns 1, otherwise answer randomly.

- If $b=0$ (probability $\frac{1}{2}$), then $\mathcal{A}$ does not receive a correct simulation because no arm is pulled at round t. According to our assumption, $\mathcal{A}$ does not give any advantage.
  - If $\mathcal{A}$ returns the correct i^t_m (probability $\frac{1}{K}$), then $\mathcal{B}$ answers 1 and is wrong.
  - Otherwise (probability $1-\frac{1}{K}$), then $\mathcal{B}$ answers randomly and is correct with probability $\frac{1}{2}$. This branch yields a probability of success of $\frac{1}{2}(1-\frac{1}{K})\frac{1}{2}$.
- If $b=1$ (probability $\frac{1}{2}$), then the advantage given by $\mathcal{A}$ can be leveraged by $\mathcal{B}$.
  - If $\mathcal{A}$ returns the correct i^t_m (probability $\frac{1}{K}+x+\text{negl}(\lambda)$), then $\mathcal{B}$ correctly answers 1. The probability of success of this branch is $\frac{1}{2}(\frac{1}{K}+x+\text{negl}(\lambda))$.
  - Otherwise (probability $1-\frac{1}{K}-x-\text{negl}(\lambda)$), $\mathcal{B}$ answers randomly and is correct with probability $\frac{1}{2}$. This branch yields a probability of success of $\frac{1}{2}(1-\frac{1}{K}-x-\text{negl}(\lambda))\frac{1}{2}$.

By aggregating the aforementioned cases, the probability α of success of $\mathcal{B}$ is:

$$\begin{aligned}
\alpha &= \frac{1}{2}(1-\frac{1}{K})\frac{1}{2} + \frac{1}{2}(\frac{1}{K}+x+\text{negl}(\lambda)) + \frac{1}{2}(1-\frac{1}{K}-x-\text{negl}(\lambda))\frac{1}{2} \\
&= \frac{1}{4} - \frac{1}{4K} + \frac{1}{2K} + \frac{x}{2} + \frac{1}{4} - \frac{1}{4K} - \frac{x}{4} + \text{negl}(\lambda) \\
&= \frac{1}{2} + \frac{x}{4} + \text{negl}(\lambda)
\end{aligned}$$

Hence, $\mathcal{B}$ has an advantage of $\frac{x}{4}$ in the IND-CPA game, which is non negligible. This contradicts the fact that AES-CBC is IND-CPA secure. Hence, we conclude that there does not exist any PPT adversary $\mathcal{A}$ that violates the property stated in the theorem. □

Theorem 7. *For an arm $i \in [\![K]\!]$ and a round $t \in [\![N-K+1]\!]$, an honest-but-curious external observer cannot learn $s_{i,t}$, given $data_{ext}^t$, with a probability better than random. More precisely, for all PPT adversaries $\mathcal{A}$,*

$$\left|\Pr\left[(i, \hat{s}_{i,t}) \leftarrow \mathcal{A}^{sum(t)}(data_{ext}^t); \hat{s}_{i,t} = s_{i,t}\right] - p_Q(t, s_{i,t})\right|$$

is negligible in λ, where $\mathcal{A}^{sum(t)}(data_{ext}^t)$ returns $(i, \hat{s}_{i,t})$ in which $\hat{s}_{i,t}$ is $\mathcal{A}$'s guess on $s_{i,t}$ for the arm i (chosen by $\mathcal{A}$), and $p_Q(t, s_{i,t})$ is the probability of obtaining a sum of rewards $s_{i,t}$ from at most t pulls of arm i until round t.

Proof. The external observer collects $data_{ext}^t$, which consists of several encrypted messages, some of them being encrypted with Enc (AES-CBC) and some other being encrypted with $\mathcal{E}_{\mathsf{DC}}$ (Paillier). We prove that these messages do not provide an advantage bigger than the advantage of an adversary in a classical IND-CPA game on Enc or $\mathcal{E}_{\mathsf{DC}}$. For simplicity, we assume that the $data_{ext}^t$ only contains two encrypted messages, $\mathsf{Enc}(m)$ and $\mathcal{E}_{\mathsf{DC}}(n)$. The proof can obviously be adapted if $data_{ext}^t$ consists of more than two messages.

The goal of the adversary is to extract at least a bit of information from either m or n. The entropy of this system is minimal when $m = n$. Hence, when $m = n$, the adversary has the highest probability of guessing at least a bit from either m or n (which are the same in this case). As a consequence, in the general case, the advantage of an adversary having to guess a bit about m or n, knowing $\mathsf{Enc}(m)$ or $\mathcal{E}_{\mathsf{DC}}(n)$ is bounded above by the advantage of an adversary having to guess a bit about m, knowing $\mathsf{Enc}(m)$ and $\mathcal{E}_{\mathsf{DC}}(m)$.

Let us prove that the advantage of a PPT adversary in this latter case (having to guess a bit about m from $\mathsf{Enc}(m)$ and $\mathcal{E}_{\mathsf{DC}}(m)$) is negligible.

We assume, toward a contradiction, that there exists a PPT adversary $\mathcal{A}$ able to win the game where, given $\mathsf{Enc}(m)$ and $\mathcal{E}_{\mathsf{DC}}(m)$, $\mathcal{A}$ recovers a bit of information about m with a non-negligible advantage x: given $\mathsf{Enc}(m)$ and $\mathcal{E}_{\mathsf{DC}}(m)$, the probability that $\mathcal{A}$ outputs a correct guess about a bit of m is equal to $\frac{1}{2} + x + \text{negl}(\lambda)$.

We use this adversary to create another adversary $\mathcal{B}$ able to break the IND-CPA property of the encryption schemes Enc (or $\mathcal{E}_{\mathsf{DC}}$, respectively). As usually in the IND-CPA game, $\mathcal{B}$ chooses two messages m_0 and m_1, and sends them to the challenger. Then, $\mathcal{B}$ receives the challenge $\mathsf{Enc}(m_b)$ (or $\mathcal{E}_{\mathsf{DC}}(m_b)$, respectively), and calls $\mathcal{A}(\mathsf{Enc}(m_b), \mathcal{E}_{\mathsf{DC}}(m_0))$ (or $\mathcal{A}(\mathsf{Enc}(m_0), \mathcal{E}_{\mathsf{DC}}(m_b))$, respectively). If $\mathcal{A}$ returns a correct guess about m_0, then $\mathcal{B}$ returns 0. Otherwise, it returns 1.

- If $b = 0$ (happens with probability $\frac{1}{2}$), then $\mathcal{A}$ has a non negligible advantage in guessing a bit about m.
  - $\mathcal{A}$ outputs a correct guess about one bit of m_0 with probability $\frac{1}{2} + x + \text{negl}(\lambda)$. In this case, $\mathcal{B}$ is correct. This branch happens with probability $\frac{1}{2}(\frac{1}{2} + x + \text{negl}(\lambda))$.
  - If $\mathcal{A}$ does not answer correctly (happens with probability $\frac{1}{2} - x - \text{negl}(\lambda)$), then $\mathcal{A}$ is correct with probability $\frac{1}{2}$. This branch happens with probability $\frac{1}{2}(\frac{1}{2} - x - \text{negl}(\lambda))\frac{1}{2}$.
- If $b = 1$ (happens with probability $\frac{1}{2}$), then $\mathcal{A}$ has no advantage.
  - If $\mathcal{A}$ returns a correct guess about one bit of m_0 (happens with probability $\frac{1}{2}$), then $\mathcal{B}$ is wrong.
  - If not (happens with probability $\frac{1}{2}$), then $\mathcal{A}$ returns a random guess and is correct with probability $\frac{1}{2}$. This branch of events happen with probability $\frac{1}{2^3}$.

By aggregating these cases, the probability α of success of $\mathcal{B}$ is:

$$\begin{aligned}\alpha &= \frac{1}{2}(\frac{1}{2} + x + \text{negl}(\lambda)) + \frac{1}{2}(\frac{1}{2} - x - \text{negl}(\lambda))\frac{1}{2} + \frac{1}{8} \\ &= \frac{1}{4} + \frac{1}{2}x + \frac{1}{8} - \frac{1}{4}x + \frac{1}{8} + \text{negl}(\lambda) \\ &= \frac{1}{2} + \frac{1}{4}x + \text{negl}(\lambda)\end{aligned}$$

Hence, $\mathcal{B}$ has a non-negligible advantage of $\frac{1}{4}x$ in the IND-CPA game against Enc (or $\mathcal{E}_{\mathsf{DC}}$, respectively), which is a contradiction with its IND-CPA property. Guessing a bit about the encrypted message is equivalent to guessing the reward with a probability better than random (i.e., better than $p_Q(t, s_{i,t})$ cf. our theorem statement), which concludes our proof. □

As a corollary, by Lemma 1 and Theorem 7, we infer that the external observer cannot learn the cumulative reward with probability better than random.

receive ciphertext from arm node $\mathsf{R}_{\sigma^{-1}(K)}$
/ ciphertext is now* $\mathsf{Enc}(\mathsf{Enc}_{i_m}(i_m))$ **/*
let ciphertext2 = Dec(ciphertext)
/ ciphertext2 is* $\mathsf{Enc}_{i_m}(i_m)$*, but* AS *does not know* i_m*, hence it has to try each* i **/*
for $i \in [\![K]\!]$
if Dec_i(ciphertext2) $= i$
/ ciphertext2 decrypts as a correct arm index */*
let $i_m = i$
break

(a) For AS, take all except the last 2 lines in Fig. 4(a), then take the above.

if $first_i = 0$
receive ciphertext2 from preceding arm node in ring
/ ciphertext2 is now* $\mathsf{Enc}(B_m||\mathsf{Enc}_{i_m}(i_m))$ **/*
$B_m||\mathsf{Enc}_{i_m}(i_m) = \mathsf{Dec}$(ciphertext2)
if $first_i = 1$ **or** $B_m < B_i$
let $i_m = i$
let $B_m = B_i$
if $next_i \neq 0$
send $\mathsf{Enc}(B_m||\mathsf{Enc}_{i_m}(i_m))$ to R_{next_i}
else
send $\mathsf{Enc}(\mathsf{Enc}_{i_m}(i_m))$ to AS

(b) For R_i (with $i \in [\![K]\!]$), take first 8 lines in Fig. 4(b), then take the above.

Fig. 8. Modifications to UCB-DS pseudocode cf. Fig. 4 to obtain UCB-DS2.

e) Impact of collusions: As pointed out earlier, an hypothesis behind our security theorems is that cloud nodes do not collude. By collusion we mean that cloud nodes put together all their data. If at least 2 of the R_i nodes collude, they could learn their respective algorithm inputs (i.e., bandit arm values that only the data owner is supposed to know at the same time) and outputs (i.e., cloud nodes could sum up the partial sums of rewards known by each node), hence UCB-DS would not satisfy the desirable security properties. In addition to following a standard security model (cf. discussion in Sect. III-A), we believe that the no-collusion hypothesis is necessary if we want a secure cumulative reward maximization algorithm that produces exactly the same output as standard UCB, which manipulates real numbers i.e., B_i needs average, ln, and sqrt. Indeed, as already mentioned in the introduction, it is not currently possible in practice to use fully-homomorphic encryption on real numbers without result approximation. Hence, to minimize data leakage, our choice is to do computations on real numbers in clear, and to distribute reward functions and B_i-value computations among K cloud nodes (one per arm), each of them having access in clear only to data pertaining to its arm.

C. Security Proof for Sect. III-D

A property of UCB-DS, stated in Theorem 2, is that an arm node R_i knows with average probability of $\frac{1}{2} + \frac{1}{2K}$ what arm is pulled at the next round. This happens because during the ring computation, every arm sees in clear the partial argmax i_m. The UCB-DS2 refinement of UCB-DS removes the aforementioned leakage and hence allows relaxing the second hypothesis from the beginning of Sect. A-B.

We show in Fig. 8 the modifications to Step 3 and 4 of UCB-DS cf. Fig. 4 that allow to obtain UCB-DS2. In the worst case, these modifications cost $(N-K+1)(K-1)$ encryptions at Step 3 and $(N-K+1)K$ decryptions at Step 4, which does not change the overall asymptotic behavior outlined in Sect. III-C3.

All theorems from Sect. A-B also hold for UCB-DS2, except Theorem 2 that is replaced by the next theorem, which formally states the stronger security guarantees of UCB-DS2.

Theorem 8. *In UCB-DS2, at the end of round* $t \in [\![N-K]\!]$ *and before the start of round* $t+1$*, given* $data^t_{\mathsf{R}_i}$*, an honest-but-curious arm node* R_i *cannot learn the arm to be pulled at round* $t+1$ *with probability better than random.*

Proof. At each round t, the arm node R_i receives $\mathsf{Enc}(B_m||\mathsf{Enc}_{i_m}(i_m))$ and decrypts into $B_m||\mathsf{Enc}_{i_m}(i_m)$. By hypothesis, B_m does not leak any information about the next arm to be pulled. The only way for R_i to guess the next arm with probability better than random is to use some information contained in $\mathsf{Enc}_{i_m}(i_m)$. However, since Enc_{i_m} is IND-CPA, it is impossible to learn any information on i_m with non negligible advantage. Hence, the strategy of R_i to guess the arm pulled at round $t+1$ is not better than random. □

APPENDIX B
ADDITIONAL INFORMATION ON EXPERIMENTS

A. Additional Experiments for Scalability with Respect to N

We present additional experimental results, using the same experimental setup as in Sect. IV. More precisely, in Fig. 9, we show the results for five additional scenarios from the related work, which confirm all the observations that we have already discussed for the scenario in the paragraph Scalability with respect to N of Sect. IV.

B. Pre-processing of Real-World Data

For this experiment, we use the same data and experimental setup as [20]. More precisely, we use data from Jester[6][21], a collection of ratings ranging from -10 (very not funny) to 10 (very funny), given by 25K users on 100 jokes. Exactly as [20], we pre-process this dataset by assigning the lowest score to the unrated jokes, and then we extract two bandit scenarios:

- Jester-small: $K = 10$, corresponding to the 10 most rated jokes, where μ_i = (# of ratings $\geq$ threshold 3.5 for joke i) / (# of users).
- Jester-large: $K = 100$, corresponding to all 100 jokes, where μ_i is computed similarly as for Jester-small, except that the threshold here is set to 7.

Moreover, we use data from MovieLens[7][22], more precisely the "MovieLens 100K Dataset" that contains ratings ranging from 1 (bad) to 5 (very good) given by 1K users on a set movies, from which, exactly as [20], we look only at the first 100 movies and derive the following bandit scenario:

- MovieLens: $K = 100$, corresponding to the first 100 movies, where μ_i = (# of ratings $\geq$ threshold 4 for movie i) / (# of users).

[6]http://eigentaste.berkeley.edu/dataset/

[7]https://grouplens.org/datasets/movielens/

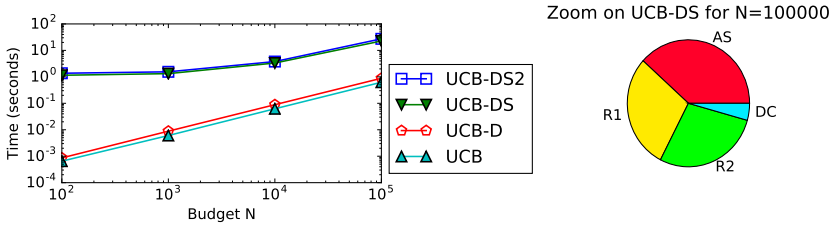

(a) Scenario 2 [4], [6]: $K = 2$, $\mu_1 = 0.9$, $\mu_2 = 0.6$.

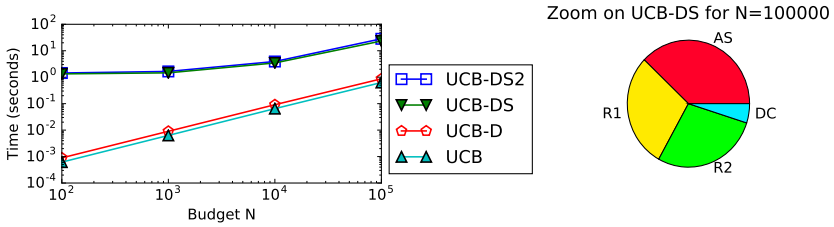

(b) Scenario 3 [4] $K = 2$ $\mu_1 = 0.9$, $\mu_2 = 0.8$

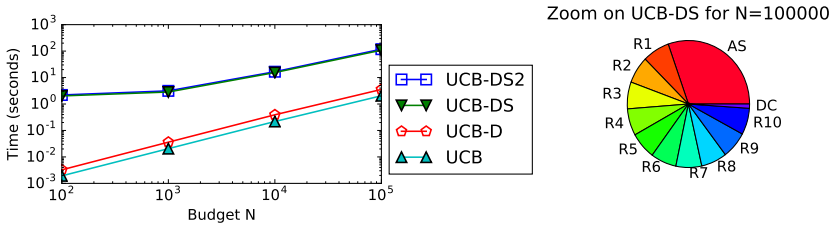

(c) Scenario 4 [4] $K = 10$: $\mu_1 = 0.9$, $\mu_2 = \mu_3 = \mu_4 = 0.8$, $\mu_5 = \mu_6 = \mu_7 = 0.7$, $\mu_8 = \mu_9 = \mu_{10} = 0.6$.

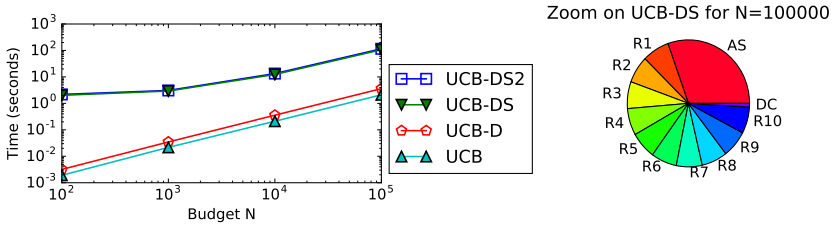

(d) Scenario 5 [4] $K = 10$: $\mu_1 = 0.9$, $\mu_2 = \ldots = \mu_{10} = 0.6$.

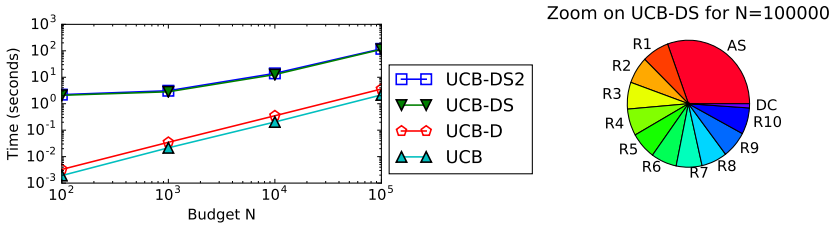

(e) Scenario 6 [6] $K = 10$: $\mu_1 = 0.55$, $\mu_2 = 0.2$, $\mu_3 = \ldots = \mu_{10} = 0.1$.

Fig. 9. Scalability with respect to N for five more bandit scenarios from the related work. In the zoom, we do not show DO (its share is close to 0).