

Parallel Contextual Bandits in Wireless Handover Optimization

Igor Colin

Noah's Ark Lab Paris
Huawei Technologies Ltd
igor.colin@huawei.com

Albert Thomas

Noah's Ark Lab Paris
Huawei Technologies Ltd
albert.thomas@huawei.com

Moez Draief

Noah's Ark Lab Paris
Huawei Technologies Ltd
moez.draief@huawei.com

Abstract—As cellular networks become denser, a scalable and dynamic tuning of wireless base station parameters can only be achieved through automated optimization. Although the contextual bandit framework arises as a natural candidate for such a task, its extension to a parallel setting is not straightforward: one needs to carefully adapt existing methods to fully leverage the multi-agent structure of this problem. We propose two approaches: one derived from a deterministic UCB-like method and the other relying on Thompson sampling. Thanks to its bayesian nature, the latter is intuited to better preserve the exploration-exploitation balance in the bandit batch. This is verified on toy experiments, where Thompson sampling shows robustness to the variability of the contexts. Finally, we apply both methods on a real base station network dataset and evidence that Thompson sampling outperforms both manual tuning and contextual UCB.

I. INTRODUCTION

The land area covered by a cellular wireless network, such as a mobile phone network, is divided into small areas called *cells*, each cell being covered by the antenna of a fixed *base station* (see Figure 1). Each base station is configured by a set of parameters that should be tuned so as to provide the best possible network coverage. Although default recommended values can be used, the best values of these parameters are likely to depend on the traffic (*e.g.*, the number of users) and the geographical location of the base stations. Furthermore, these parameters often need to be adjusted on a regular basis in order to adapt to the evolution of the traffic. The manual tuning of base station parameters may thus be highly time consuming, tedious and needs to preserve some level of quality of service. In addition, recent developments in cellular network standards lean towards a densification of base stations, encouraging operators to find automated solutions for optimal parameters configuration (see *e.g.*, [1]–[4]).

One possible way of modeling this problem is through contextual bandits [5]: in this framework, one aims at optimizing an objective that is depending on contextual features (*e.g.*, traffic and environment) while avoiding too much deterioration of the objective function (*e.g.*, quality of service). More specifically, at each time t , given side information about the current state of the wireless network (the *context*), the operator wants to choose the values of the parameters (an *action*) so as to obtain the best user experience (the *reward*). However, since the relation between an action and its associated reward is

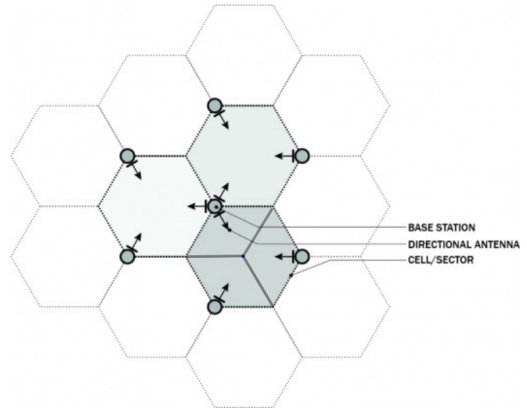


Fig. 1. Cellular network organization

initially unknown, one needs to *explore* the space of available actions in order to gain some knowledge about this relation before being able to *exploit* it. This exploration-exploitation tradeoff is common to every bandit problem, from multi-armed bandits [6]–[8] to lazy optimization through gaussian processes [9]–[12].

Contextual bandit problems have received a lot of attention in the past decade, either for theoretical guarantees [13], [14], delayed rewards framework [15] or pure exploration scenarios [16], [17]. Surprisingly, the multi-agent setting, such as the one induced by the base station parameter tuning, has hardly been investigated [18]. Efficiently extending existing methods to a parallel setting is not straightforward: naive implementation of an Upper Confidence Bound algorithm, for instance, may lead to suboptimal balance between exploration and exploitation when contexts are too similar. The goal of this paper is to formulate methods leveraging the multi-agent structure and to apply them to wireless handover optimization.

The paper is structured as follows. Sections II and III formally define the parallel contextual bandit problem. Section IV reviews related methods in the bandit literature. Section V develops two approaches for parallel contextual bandits. Finally, Section VI shows empirical performances of our methods, first on a toy example and then on a real wireless base station dataset.

II. DEFINITIONS AND NOTATIONS

For any integer $n > 0$, we denote by $[n]$ the set $\{1, \dots, n\}$ and by $|\mathcal{X}|$ the cardinality of any finite set \mathcal{X} . For any $t > 0$ and any sequence $(u_1, u_2, \dots) \in \mathcal{X}^{\mathbb{N}}$, \mathbf{u}_t will denote the sequence up to index t , that is (u_1, \dots, u_t) . For any element x of \mathbb{R}^d , $\|x\|$ will denote the euclidean norm and $|x|$ the ℓ_1 -norm. For a given positive definite matrix $A \in \mathbb{R}^{d \times d}$, $\|\cdot\|_A$ denotes the norm induced by the scalar product associated to A , that is for any $x \in \mathbb{R}^d$, $\|x\|_A = \sqrt{x^\top A x}$.

III. PROBLEM STATEMENT

Let $d > 0$ and let $\mathcal{X} \subseteq \mathbb{R}^d$. For a given function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and a parameter $\theta_\star \in \mathbb{R}^d$, we define the reward of a contextual bandit as: for any context $x \in \mathcal{X}$,

$$r(x) = f(x; \theta_\star) + \varepsilon, \quad (1)$$

where ε is a R -sub-gaussian random variable independent of x , for some $R > 0$. The function f is called the expected reward. The contextual bandits problem consists in the following. At each iteration $0 < t < T$, a set of contexts \mathcal{X}_t is presented; one aims at selecting the context $x_t \in \mathcal{X}_t$ in order to minimize the expected regret. This boils down to finding the sequence $(x_1, \dots, x_T) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_T$ maximizing

$$\sum_{t=1}^T \mathbb{E}[r(x_t)] = \sum_{t=1}^T f(x_t; \theta_\star). \quad (2)$$

In most settings, the context space \mathcal{X} can be further expanded as a product of a state space \mathcal{S} and an action space \mathcal{A} . In other words, at a given iteration, one observes a state $s \in \mathcal{S}$ and wants to find the action maximizing the expected reward:

$$a^\star(s) \in \arg \max_{a \in \mathcal{A}} f(s, a; \theta_\star).$$

When the expected reward is parametrized, a natural strategy is to estimate the parameter θ_\star while limiting the regret as much as possible. This type of methods has been extensively studied in the multi-armed bandit—*i.e.*, context-free bandit—setting [6]–[8]. More recently, the contextual bandits setting has been investigated, whether the function f is linear [13], [14], logistic [19], [20] or unknown [9], [21], [22].

One way of addressing the problem efficiently is to use an Upper Confidence Bound (UCB) framework. It is a straightforward adaptation from UCB in the multi-armed bandit (MAB) case: given a state s and for each action a , one uses past observations to build associated confidence region on the expected reward and chooses the action associated to the greatest possible outcome. The method is formally stated in Figure 2.

In this paper, we focus on the problem where $n > 0$ contextual bandits run in parallel. We consider in addition that they share a similar function f , although their parameters $(\theta_\star^{(i)})_{1 \leq i \leq n}$ are not necessarily identical. The regret (2) is then replaced by the aggregated regret:

$$\sum_{t=1}^T \sum_{i=1}^n f(x_t^{(i)}; \theta_\star^{(i)}). \quad (3)$$

Require: Confidence parameter α

```

for  $t = 1, \dots, T$  do
  Receive state  $s_t$ 
  for  $a \in \mathcal{A}$  do
    Build a confidence region:  $C_\alpha(s_t, a)$ 
  end for
  Select  $a_t \in \arg \max_{a \in \mathcal{A}} \{\sup C_\alpha(s_t, a)\}$ 
  Observe reward  $r_t$ 
end for

```

Fig. 2. UCB for contextual bandit.

Without any additional assumption on the parameters $(\theta_\star^{(i)})_{1 \leq i \leq n}$, a straightforward strategy is to run one of the aforementioned policies independently on each bandit. However, if the parameters are selected from a restricted set, say $\theta_\star^{(i)} \in \Theta$ with $|\Theta| \ll n$, then one may wonder whether it is possible to use this structure to improve the regret minimization policy. This interrogation is investigated in [23] for the multi-armed bandit setting and in [24] for the contextual bandit setting, when the arms (resp. the contexts) are pulled (resp. chosen) sequentially, one at a time. In our setting however, we are interested in finding a strategy for choosing n contexts at each iteration, since the bandits run in parallel. In order to emphasize the interest of our approaches when this assumption is verified, we consider in the remainder of this paper a simpler setting, where every bandit shares the same parameter θ_\star and, consequently, the same expected reward function.

IV. RELATED WORK

Surprisingly, the parallel bandit setting has not been widely studied in the contextual bandit case. There is a lot of literature about closely related settings however; we detail each of these settings in the following sections.

A. Delayed reward

One way to model the parallel bandit problem is to consider bandits with delayed rewards [15]: we assume that the reward is not observed immediately after every action but rather delayed. If the rewards are received every n iterations, this setting is then equivalent to n contextual bandits with identical expected reward functions running in parallel. In the general online learning setting, the delayed feedback is modeled as follows. At a given iteration t , the environment chooses a state s_t and a set of admissible actions \mathcal{A}_t , just as in the standard setting. The reward however is only observed after a delay τ_t , possibly random, that is usually unknown in advance to the learner. One interesting result, enlightened in [15, Table 1], is the fact that the additional regret induced by the delay is additive in the stochastic feedback case and multiplicative in the adversarial setting. In our setting, the delay τ is in the set $\{0, \dots, n-1\}$, where n is the number of bandits. Therefore, the additional regret is proportional to the number of bandits.

The general online learning with delayed feedback problem was deeply investigated and can be extended to a wide range

Require: Confidence parameter α

```

for  $a \in \mathcal{A}$  do
  Initialize  $\hat{A}_a \leftarrow I, \hat{b}_a \leftarrow 0, \hat{w}_a \leftarrow 0$ 
end for
for  $t = 1, \dots, T$  do
  for  $i = 1, \dots, n$  do
    Select  $a_*^{(i)} \in \arg \max_{a \in \mathcal{A}} \hat{w}_a^\top s_t^{(i)} + \alpha \|s_t^{(i)}\|_{\hat{A}_a^{-1}}$ 
    Update  $A_{a_*^{(i)}} \leftarrow A_{a_*^{(i)}} + s_t^{(i)} s_t^{(i)\top}$ 
  end for
  Observe rewards  $r_t^{(1)}, \dots, r_t^{(n)}$ 
  for  $a \in \mathcal{A}$  do
    Update  $\hat{b}_a \leftarrow \hat{b}_a + \sum_{i=1}^n \mathbb{1}_{\{a=a_*^{(i)}\}} r_t^{(i)} s_t^{(i)}$ 
    Update  $\hat{w}_a \leftarrow \hat{A}_a^{-1} \hat{b}_a$ 
  end for
end for

```

Fig. 3. LinUCB-PR algorithm.

of applications (MDPs, distributed optimization, *etc.*), see *e.g.*, [15] and references therein for further details.

The more specific problem of multi-armed bandits with delayed feedback has been extensively studied in the past decade [25]–[27]. The particular structure of the problem allows for different approaches than general online learning that sometimes lead to improved convergence guarantees or decreased storage costs.

B. Piled rewards

Our setting offers more structure than a general contextual bandit with delayed rewards since the rewards are accumulated and then all disclosed simultaneously at a given iteration; to the best of our knowledge, this concept of “piled rewards” is only developed in [18] for the linear contextual bandit, that is a contextual bandit with a linear reward function:

$$f(x; \theta) = x^\top \theta.$$

The proposed algorithm is based on LinUCB, the essential difference being that the covariance matrix is updated as it cycles through the agents. The algorithm is detailed in Figure 3. This update trick is shown to shrink the confidence area as the different bandits are looped over, that is $\|s\|_{\hat{A}^{-1}}$ does not increase as \hat{A} is updated, for any state s . The LinUCB-PR is shown to have a $O(n^2 T |\mathcal{A}| \log(nT |\mathcal{A}| / \delta))$ regret with probability $1 - \delta$. This is not an improvement over LinUCB applied to the parallel setting but it behaves empirically better for large values of T , possibly due to the shrunk confidence region. This lack of improvement may be explained by the overconfidence induced by the intermediate updates. As explained in [11], one should build an overconfidence measure and moderate exploratory redundancy, before deriving any regret bound.

This approach is quite similar to the ones we will detail in the next sections. We will not limit our attention to the linear setting though—the logistic setting will be of particular




Fig. 4. Illustration of GP-UCB algorithm. The posterior mean of the gaussian process is the solid blue line and the associated variance is the filled area. The next evaluation candidate is the one maximizing the *optimistic* prediction of f .

interest—and we will develop several methods to tackle this problem in a more general fashion.

C. Gaussian processes

Lazy optimization with gaussian processes is a particular application of Bayesian optimization where one aims at finding the maximum of a possibly non-convex objective function f . The idea is to use a gaussian process as a prior on f and then to sequentially refine the posterior as objective values are observed. A popular method for optimizing in the gaussian process framework is GP-UCB, which is an extension of UCB to gaussian processes. Indeed, at each iteration, one queries the point presenting the highest upper confidence bound based on the posterior of the objective. This is illustrated on Figure 4. A typical application lies in hyperparameter tuning for machine learning algorithms [12]. The objective is the negative empirical risk (or any fitting score), which is typically expensive to evaluate. The need to parallelize gaussian process algorithms therefore arose naturally: while being expensive to evaluate, one may have access to additional computational power in order to perform several evaluations simultaneously. However, since GP-UCB selects the optimal candidate in a deterministic fashion, the extension of GP-UCB to a parallel setting is not straightforward. This issue received a lot of attention recently, and many types of approaches have been developed to tackle it. We detail three methods, as they seem to grasp the main ideas of parallelizing, but the reader may find many derivations in, *e.g.*, [28]–[33].

First, a method based on pure exploration techniques, namely GP-UCB-PE, has been proposed in [10]. The idea is to select the GP-UCB candidate and define a confidence region around that candidate. Then, the subsequent queries will be selected in order to “maximize the exploration”, that is each query will select the point in the confidence region with highest posterior variance. Since the posterior variance does not need the objective value to be updated, this ensures distinct candidates among the batch. With a batch size n , resulting

expected regret is improved by a factor \sqrt{n} in terms of time and is similar for large T in terms of number of queries. This method is somewhat similar to the pure exploration for batched bandits proposed in [27] in the sense that a part of the batch (here $n - 1$ agents) is dedicated to exploring as much as possible in order to guarantee an improved behavior of the remainder of the batch (here the first agent).

Another approach for parallelizing gaussian processes was introduced in [11] as GP-BUCB. As is the case for GP-UCB-PE, this method relies on the fact that posterior variance only depends on the points selected, not their associated values. The next elements are then chosen following a twofold criterion: the standard UCB criterion on the updated process and a overconfidence criterion. The purpose of the latter is to avoid the exploratory redundancy already mentioned in Section IV-B—see [11, Section 4.1] for additional details. Regret analysis of this algorithm yields bounds roughly similar to GP-UCB-PE in terms of number of queries.

Finally, the last method we focus on relies on Thompson sampling—see *e.g.*, [34], [35]. At each iteration, the posterior gaussian process is updated. Then, n functions are sampled from this distribution and the n candidates are the maximizers of these sampled functions. This class of methods will be of particular interest in the next section because it seems well-suited for global regret minimization: as opposed to GP-UCB-PE and GP-BUCB, it does not necessarily involve improving the performance of only one or few agents in the batch.

V. PARALLEL BANDITS WITH IDENTICAL PARAMETERS

In this section, we develop the two main approaches for contextual bandits and see how they can be adapted in the parallel bandit setting.

A. UCB contextual bandits

Upper Confidence Bound (UCB) approaches for contextual bandits rely on the following framework, first introduced for the linear case in [36]. At a given iteration $t > 0$, one is able to build a confidence region B_{t-1} for the parameter θ_* , based on the previously selected contexts and associated rewards $(x_1, r_1), \dots, (x_{t-1}, r_{t-1})$. A set of contexts \mathcal{X}_t is proposed and one then chooses the context leading to the best possible reward on B_{t-1} :

$$x_t \in \arg \max_{x \in \mathcal{X}_t} \left\{ \max_{\theta \in B_{t-1}} f(x; \theta) \right\}. \quad (4)$$

Methods based on UCB present the advantage of being easy to implement and fast to compute. They obviously require the knowledge of a “good” region B_{t-1} in the sense that B_{t-1} should be as tight as possible with respect to the selected confidence level.

In the specific case of parallel bandits with identical parameters, the confidence region is based on every bandit historical contexts and rewards: $(x_1^{(i)}, r_1^{(i)}), \dots, (x_{t-1}^{(i)}, r_{t-1}^{(i)})$, for $i \in [n]$. Then, n contexts are selected, one from every $\mathcal{X}_t^{(i)}$, using the UCB rule. Although this method will be preferable to independent policies on each bandits, it may lack

Require: B_0 , update rule for B

for $t = 1, \dots, T$ **do**

for $i = 1, \dots, n$ **do**

Select $x_t^{(i)}$ according to (4)

end for

Observe the rewards $r_t^{(1)}, \dots, r_t^{(n)}$

Update B_t with observed rewards and selected contexts

end for

Fig. 5. UCB algorithm for parallel bandits.

the expected exploration/exploitation balance if the contexts sets $\mathcal{X}_t^{(1)}, \dots, \mathcal{X}_t^{(n)}$ are too similar to enforce different choices amongst the bandits. Indeed, in the—extreme—setting where the contexts sets are identical, the selected contexts will be identical at every bandit: the policy will enforce either a full exploration or a full exploitation scheme, being no different from a setting with only one bandit. The potential regret improvement with respect to independent policies relies solely on the variety of the contexts. This method is formally stated in Figure 5

B. Bayesian contextual bandits

In this section, we focus on bayesian approaches for contextual bandits, and more specifically Thompson-based approaches. In such setting, one defines a prior probability $p(\theta)$ on the parameter to estimate. At iteration $t > 0$, the posterior probability is then of the form

$$p(\theta | \mathbf{x}_{t-1}, \mathbf{r}_{t-1}) \propto p(\theta) p(\mathbf{r}_{t-1} | \mathbf{x}_{t-1}, \theta),$$

where $\mathbf{x}_{t-1} = (x_1, \dots, x_{t-1})$, $\mathbf{r}_{t-1} = (r_1, \dots, r_{t-1})$ and $p(\mathbf{r}_{t-1} | \mathbf{x}_{t-1}, \theta)$ is the likelihood function. For the sake of simplicity, we use the notation $p(\theta | t - 1)$ for the posterior $p(\theta | \mathbf{x}_{t-1}, \mathbf{r}_{t-1})$. Using this relation, one may sample a parameter θ_t from the posterior probability, either using a closed form or an approximation—*e.g.*, MCMC or Laplace approximation. Finally, the context selected from \mathcal{X}_t is the context maximizing the expected reward parametrized by θ_t :

$$x_t \in \arg \max_{x \in \mathcal{X}_t} \{ f(x; \theta_t) \}. \quad (5)$$

Bayesian approaches may offer more flexibility when confidence bounds are not tight but are often much slower to compute, even with rough approximations.

In the specific case of parallel bandits with identical parameters, the posterior is built on every bandit historical contexts and rewards. Then, there are two ways of adapting the regular Thompson approach. First, one may sample one θ_t and select the contexts $(x_t^{(i)})_{1 \leq i \leq n}$ according to the rule (5). Another method is to sample n parameters $\theta_t^{(1)}, \dots, \theta_t^{(n)}$ independently from the same posterior and then to select every context according to its respective parameter. The former is similar to the adaption of UCB methods to a parallel setting. The latter however benefits from the parallel setting, since it will enforce a better balance between exploitation and exploration at a limited cost (sampling being usually cheap *e.g.*, when

Require: Prior $p(\theta)$
for $t = 1, \dots, T$ **do**
 Sample a parameter $\theta_t \sim p(\theta|t-1)$
 for $i = 1, \dots, n$ **do**
 Select $x_t^{(i)}$ according to (5)
 end for
 Observe the rewards $r_t^{(1)}, \dots, r_t^{(n)}$
 Update the posterior with observed rewards and selected contexts
end for

Fig. 6. Naive Thompson-based algorithm for parallel bandits.

Require: Prior $p(\theta)$
for $t = 1, \dots, T$ **do**
 for $i = 1, \dots, n$ **do**
 Sample a parameter $\theta_t^{(i)} \sim p(\theta|t-1)$
 Select $x_t^{(i)}$ according to (5)
 end for
 Observe the rewards $r_t^{(1)}, \dots, r_t^{(n)}$
 Update the posterior with observed rewards and selected contexts
end for

Fig. 7. Multisampling Thompson-based algorithm for parallel bandits.

using Laplace approximation). In the setting where all contexts sets are identical, the randomness of the sampled parameters will preserve the exploitation/exploration balance, even if no reward is observed until every context is chosen. Both approaches are detailed in Figure 6 and 7.

Previous integrations of Thompson sampling in a parallel setting [34], [35] as well as its well-studied empirical behavior [19] suggests that it will behave better—regret-wise—than UCB in the contextual bandit setting. However, the theoretical analysis of such methods is far from trivial and is out of the scope of this paper: even MAB theoretical bounds were provided only a few years ago [37]–[39] and bounds for the contextual case are limited to linear payoffs [14]. Consequently, the next section is devoted to exhibiting the aforementioned differences between the two methods and then applies both algorithms to the handover parameter tuning on a real base stations dataset.

VI. EXPERIMENTS

We first illustrate on a toy example the advantages, mentioned at the end of the previous section, of the multisampling Thompson-based algorithm over UCB in the case of parallel bandits. We then present the results obtained when applying UCB and Thompson sampling to the online optimization of handover parameters in a wireless cellular network. As explained in the introduction, this problem can be naturally modeled as parallel contextual bandits. In all the experiments, the implemented linear UCB algorithm is the Optimism in the Face of Uncertainty Linear bandit (OFUL) algorithm, described in [36].

A. Toy example

To illustrate the benefits of the multisampling Thompson-based algorithm described in Figure 7 over the UCB algorithm described in Figure 5 we consider the toy example of a linear contextual bandit model. In this case, the expected reward is a linear function of the context x and we assume that the stochastic reward r is given by

$$r(x, \theta_*) = x^\top \theta_* + \varepsilon \quad (6)$$

where $\varepsilon \sim \mathcal{N}(0, R^2)$. We also consider here that the context $x = (s, a)$ where $s \in \mathbb{R}^{d-1}$ corresponds to the state received at each iteration t and $a \in \mathbb{R}$ to the action that has to be chosen. This is the case for the wireless cellular network application described in the introduction and section VI-B. We take a parameter θ_* of the form $\theta_* = (\theta_*^s, 1)$, where $\theta_*^s \in \mathbb{R}^{d-1}$ and the true parameter θ_*^s is sampled from a multivariate Gaussian distribution $\mathcal{N}(0, \mathbf{I}_{d-1})$ and then normalized to a unit norm vector.

At each iteration, a state $s \in \mathbb{R}^{10}$ is sampled from a multivariate Gaussian distribution $\mathcal{N}(0, \sigma_s^2 \cdot \mathbf{I}_{10})$ and the algorithm must choose between one of the 5 actions $a_i = i$, $0 \leq i \leq 4$. The associated reward is generated according to (6) where the variance R^2 of the noise term ε is set to 2.5. For both strategies, multisampling Thompson and UCB, a penalization term $0.01 \cdot \|\theta\|$ is added to the linear regression.

1) *Influence of the variance of the states:* We run $n = 20$ bandits in parallel and compare the regrets obtained with the multisampling Thompson-based algorithm and the linear UCB algorithm for different variances. The regrets are computed at a time horizon $T = 500$ for 100 random repetitions of the algorithm, the randomness coming from the strategies themselves and the generation of the states at each iteration. The results are shown in Figure 8.




Fig. 8. Cumulated regrets obtained at a time horizon $T = 500$ when running $n = 20$ bandits in parallel for different values of the variance of the states σ_s^2 .

One may observe that the regret of UCB decreases as the variance increases. This is explained by the fact that when

the variance of the contexts is small, UCB will choose the same action for each of the n bandits that are ran in parallel. This either leads to a full exploration or a full exploitation within each batch of size n . On the contrary, the multisampling Thompson-based algorithm is more robust to the change in variance and performs better than UCB. As explained in the previous section, the multisampling Thompson-based algorithm allows for a better balance between exploration and exploitation as n parameters $\theta^{(1)}, \dots, \theta^{(n)}$ are sampled from the posterior $p(\theta|\mathbf{x}_{t-1}, \mathbf{r}_{t-1})$.

2) *Influence of the number of bandits n* : We now study the impact of the number n of bandits ran in parallel. Here, the variance of the states is set to a fixed value $\sigma_s^2 = 0.01$. As above, the regrets are computed at a time horizon $T = 2000$ for 100 random repetitions of the algorithm, the randomness coming from the strategies themselves and the generation of the contexts at each step. The results are shown in Figure 9.




Fig. 9. Cumulated regrets obtained at time horizon $T = 2000$ with a variance of the states $\sigma_s^2 = 0.01$ for different values of the number of bandits ran in parallel n .

As the number of bandits n increases, the overall regret of UCB degrades whereas the performance of the multisampling Thompson-based algorithm remains stable. This is due to the fact that for large values of the number of bandits n , the multisampling Thompson-based algorithm preserves an exploration/exploitation balance compared to UCB.

As explained in section IV-A, the parallel bandits problem can be seen as bandits with delayed rewards. For the non-contextual case, it was also found in [19] that Thompson sampling was more robust to the delay than UCB.

B. Online optimization of base station parameters

The main motivation of the multisampling Thompson strategy developed in this paper is to tune base station parameters of a cellular wireless network so as to provide a good connectivity for all the users. Here we focus on parameters related to *handovers*. A handover occurs when the connection between a user and a cell is transferred to a neighboring cell




Fig. 10. Impact of event A2 threshold.

in order to ensure the continuity of the radio network coverage and prevent interruptions of communication. The reader can refer to Chapter 2 of [40] for an account on handovers. We here only present the different steps of a handover procedure.

1) *Handovers*: A handover can occur between two cells using the same frequency (*intra-frequency*) or between two cells using different frequencies (*inter-frequency*). For both types of handovers, the user triggers the handover when it receives a better signal from a neighboring cell than from the serving cell. To trigger such an event, the user needs to measure the signal received from neighboring cells. This is automatic for an intra-frequency handover. However inter-frequency measurements are only triggered when the signal received by the serving cell is lower than a pre-specified threshold. Such an event is called *event A2* and is the event of interest in this paper. If the threshold is too low, this results in a late handover and a bad data rate or *throughput* between the cell and the user. On the contrary, if the threshold is too high, unnecessary inter-frequency measurements are triggered and this also results in a bad throughput (Figure 10). Indeed, when the user is performing inter-frequency measurements on a neighboring cell, it can no longer exchange information with the serving cell. Tuning the parameter of the base station associated to this threshold can therefore improve the *throughput* between a cell and a user and this should be done for each cell of a wireless network.

Optimization of handover parameters has already been studied in the wireless literature (see e.g. [3], [4], [41]–[43]). However most of the literature mainly focuses on the optimization of handover performance metrics such as early handovers, late handovers or ping-pong handovers whereas the focus here is on the throughput.

2) *Data*: Data coming from $n = 105$ cells have been recorded every hour during 5 days. For each hour and for each cell the value of the threshold of event A2 and five traffic data features are available. The traffic features are: downlink average number of active users, average number of users, channel quality index of cell edge users, and two features related to the traffic of small data packets. The goal is to recommend the values of the threshold of event A2 so as to achieve the best possible throughput the next hour. As we are only concerned about handovers we aim to maximize the

throughput of cell edge users, i.e. users that are located at the edge of a cell and therefore have a low throughput. We thus define the quantity to optimize as the proportion of users which have a throughput lower than 5 MB/s (Megabytes per second). Obviously, the lower this proportion is the better the parameters are. With the terminology used in this paper and in the contextual bandits literature, the threshold corresponds to the *action*, the quantity that we want to optimize to the *reward* and traffic data features to the *state*.

We emphasize here on the fact that at each hour t we should recommend a threshold value for each one of the 105 cells. Assuming that the reward of each cell is parametrized by a same parameter θ , this is therefore equivalent to running $n = 105$ bandits in parallel as described in section V.

3) *Parallel logistic contextual bandits*: The reward corresponding to a proportion of users, it appears natural to use a logistic regression model. We therefore recall here the logistic contextual bandit setting. For any $x, \theta \in \mathbb{R}^d$, the expected reward is given by $f(x; \theta) = \sigma(x^\top \theta)$ where $\sigma : z \mapsto (1 + \exp(-z))^{-1}$ is the sigmoid function and the reward follows a Bernoulli distribution:

$$p(r|x, \theta) = f(x; \theta)^r (1 - f(x; \theta))^{1-r},$$

with $r \in \{0, 1\}$.

Given past observations $(\mathbf{r}_t, \mathbf{x}_t)$, the penalized negative likelihood estimator is defined by

$$\hat{\theta} = \arg \min_{\theta} \left\{ \sum_{s=1}^t \log(1 + \exp(-r_s \theta^\top x_s)) + \Omega(\theta) \right\}.$$

From a Bayesian standpoint, if $\Omega = (\lambda_2/2) \|\cdot\|^2$ (ridge penalty), the corresponding prior is Gaussian. If $\Omega = \lambda_1 |\cdot| + (\lambda_2/2) \|\cdot\|^2$ (elastic net penalty), the corresponding prior is a mixture of a Gaussian distribution and a Laplacian distribution.

Compared to the linear contextual bandits framework, the posterior $p(\theta|t)$ is here intractable. A common way to draw samples from this posterior is to use the Laplace approximation (see e.g., section 4.4 in [44]): the posterior is approximated by a Gaussian distribution $\mathcal{N}(\mu_t, \Sigma_t)$, with parameters:

$$\begin{cases} \mu_t &= \hat{\theta} \\ \Sigma_t &= \sum_{s=1}^t (1 - \sigma(x_s^\top \hat{\theta})) \sigma(x_s^\top \hat{\theta}) x_s x_s^\top. \end{cases} \quad (7)$$

In practice, Σ can be expensive to compute, so it is common to only use the diagonal coefficients and this is what we do in the experiment.

The assumption that the cells are sharing the same parameter θ_* may be a bit strong. We show in the next section how to deal with different parameters $\theta_*^{(i)}$, $1 \leq i \leq n$.

4) *Different parameters*: The parameters $(\theta_*^{(1)}, \dots, \theta_*^{(n)})$ of each cell are not necessarily identical in practice. However they should benefit from each other's observations. We thus consider that for each $i \in [n]$, $\theta_*^{(i)}$ can be decomposed into the sum of a global parameter θ_* , shared by all the cells, and a local parameter $\tilde{\theta}_*^{(i)}$: $\theta_*^{(i)} = \theta_* + \tilde{\theta}_*^{(i)}$. Let us denote $\theta =$

$(\theta, \tilde{\theta}^{(1)}, \dots, \tilde{\theta}^{(n)}) \in \mathbb{R}^{d \times (n+1)}$. The new penalized negative likelihood minimizer is:

$$\hat{\theta} = \arg \min_{\theta} \left\{ \sum_{s=1}^t \sum_{i=1}^n \log(1 + \exp(-r_s^{(i)} x_s^{(i)\top} (\theta + \tilde{\theta}^{(i)}))) + \Omega(\theta) \right\},$$

where again Ω may be a ridge penalty $\Omega(\theta) = \lambda_2/2 \|\theta\|^2 + \lambda_2'/2 \sum_{i=1}^n \|\tilde{\theta}^{(i)}\|^2$ or an elastic net penalty:

$$\Omega(\theta) = \lambda_1 |\theta| + \frac{\lambda_2}{2} \|\theta\|^2 + \lambda_1' \sum_{i=1}^n |\tilde{\theta}_i| + \frac{\lambda_2'}{2} \sum_{i=1}^n \|\tilde{\theta}^{(i)}\|^2.$$

The diagonal covariance matrix of the Laplace approximation is then:

$$\Sigma_t^{diag} = \begin{pmatrix} \Sigma_t & 0 & \dots & 0 \\ 0 & \tilde{\Sigma}_t^{(1)} & & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & \tilde{\Sigma}_t^{(n)} \end{pmatrix},$$

where Σ_t and the $\tilde{\Sigma}_t^{(i)}$ are defined by substituting the elements of $\hat{\theta}_t$ in (7).

5) *Results for wireless handover optimization*: The Thompson-based algorithm is applied with the bayesian logistic regression presented above and the OFUL algorithm is applied on the logit transform σ^{-1} of the rewards. As the true reward function $r(\cdot, \theta_*)$ is unknown, we first fit a logistic regression model on a training data set to learn parameters $(\hat{\theta}^{(1)}, \dots, \hat{\theta}^{(n)})$ which we then use as surrogates for the true parameters $(\theta_*^{(1)}, \dots, \theta_*^{(n)})$ in order to evaluate the different strategies: multisampling Thompson-based algorithm, OFUL and the strategy used to collect the data. The cumulated expected regret for each strategy is shown in Figure 11 where one can see that Thompson sampling performs better than OFUL and the strategy used to collect the data.

VII. CONCLUSION

In this paper, we exhibited two approaches for handling multi-agent scenarios in the contextual bandits framework. The first one, based on UCB, is a naive extension of the single-agent case; the second one relies on Thompson sampling in order to preserve the exploration-exploitation balance in the bandits batch. Our synthetical experiments enlightened the advantages of Thompson sampling in the parallel setting, as was suggested by theoretical and empirical studies [19]. Furthermore, application on wireless handover parameters tuning exhibited a clear superiority of Thompson sampling, in comparison of both manual tuning and UCB-like approach.

Extending this framework to a setting where contextual bandits are not identical but rather regrouped in clusters (as in [23], [24]) may be a promising way of generalizing this approach to larger networks. Also, deriving theoretical bounds for Thompson sampling in the parallel setting could lead to additional insights on how to improve existing methods.




Fig. 11. Multisampling Thompson-based algorithm versus OFUL for the online optimization of handover parameters. The current strategy denotes the strategy used to collect the data.

REFERENCES

- [1] I. Siomina, P. Varbrand, and D. Yuan, "Automated optimization of service coverage and base station antenna configuration in umts networks," *IEEE Wireless Communications*, vol. 13, no. 6, pp. 16–25, Dec 2006.
- [2] A. Awada, B. Wegmann, I. Viering, and A. Klein, "Optimizing the radio network parameters of the long term evolution system using taguchi's method," *IEEE Transactions on Vehicular Technology*, vol. 60, no. 8, pp. 3825–3839, 2011.
- [3] V. Capdevielle, A. Feki, and A. Fakhreddine, "Self-optimization of handover parameters in lte networks," in *2013 11th International Symposium and Workshops on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2013, pp. 133–139.
- [4] I. N. M. Isa, M. Dani Baba, A. L. Yusof, and R. A. Rahman, "Handover parameter optimization for self-organizing lte networks," in *2015 IEEE Symposium on Computer Applications Industrial Electronics (ISCAIE)*, 2015, pp. 1–6.
- [5] T. Lu, D. Pál, and M. Pál, "Contextual multi-armed bandits," in *AISTATS*, 2010, pp. 485–492.
- [6] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Advances in applied mathematics*, vol. 6, no. 1, pp. 4–22, 1985.
- [7] 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.
- [8] S. Bubeck, G. Stoltz, C. Szepesvári, and R. Munos, "Online optimization in x-armed bandits," in *NIPS*, 2009, pp. 201–208.
- [9] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger, "Gaussian process optimization in the bandit setting: No regret and experimental design," in *ICML*, 2010, pp. 1015–1022.
- [10] E. Contal, D. Buffoni, A. Robicquet, and N. Vayatis, "Parallel gaussian process optimization with upper confidence bound and pure exploration," in *ECML-PKDD*, 2013, pp. 225–240.
- [11] T. Desautels, A. Krause, and J. W. Burdick, "Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization," *JMLR*, vol. 15, no. 1, pp. 3873–3923, 2014.
- [12] F. Pedregosa, "Hyperparameter optimization with approximate gradient," in *ICML*, 2016, pp. 737–746.
- [13] W. Chu, L. Li, L. Reyzin, and R. Schapire, "Contextual bandits with linear payoff functions," in *AISTATS*, 2011, pp. 208–214.
- [14] S. Agrawal and N. Goyal, "Thompson sampling for contextual bandits with linear payoffs," in *ICML*, 2013, pp. 127–135.
- [15] P. Joulani, A. Gyorgy, and C. Szepesvári, "Online learning under delayed feedback," in *ICML*, 2013, pp. 1453–1461.
- [16] J.-Y. Audibert and S. Bubeck, "Best arm identification in multi-armed bandits," in *COLT*, 2010.
- [17] L. Xu, J. Honda, and M. Sugiyama, "Fully adaptive algorithm for pure exploration in linear bandits," in *AISTATS*, 2018.
- [18] K.-H. Huang and H.-T. Lin, "Linear upper confidence bound algorithm for contextual bandit problem with piled rewards," in *PAKDD*, 2016, pp. 143–155.
- [19] O. Chapelle and L. Li, "An empirical evaluation of thompson sampling," in *NIPS*, 2011, pp. 2249–2257.
- [20] O. Chapelle, E. Manavoglu, and R. Rosales, "Simple and scalable response prediction for display advertising," *ACM TIST*, vol. 5, no. 4, p. 61, 2015.
- [21] A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. Schapire, "Taming the monster: A fast and simple algorithm for contextual bandits," in *ICML*, 2014, pp. 1638–1646.
- [22] M. Valko, N. Korda, R. Munos, I. Flounas, and N. Cristianini, "Finite-time analysis of kernelised contextual bandits," in *UAI*, 2013, pp. 654–663.
- [23] O.-A. Maillard and S. Mannor, "Latent bandits," in *ICML*, 2014, pp. 1–136–1–144.
- [24] A. Gopalan, O.-A. Maillard, and M. Zaki, "Low-rank bandits with latent mixtures," *CoRR*, vol. abs/1609.01508, 2016.
- [25] S. Guha, K. Munagala, and M. Pal, "Multiarmed bandit problems with delayed feedback," *arXiv preprint arXiv:1011.1161*, 2010.
- [26] C. Vernade, O. Cappé, and V. Perchet, "Stochastic bandit models for delayed conversions," *arXiv preprint arXiv:1706.09186*, 2017.
- [27] V. Perchet, P. Rigollet, S. Chassang, E. Snowberg *et al.*, "Batched bandit problems," *The Annals of Statistics*, vol. 44, no. 2, pp. 660–681, 2016.
- [28] E. A. Daxberger and B. K. H. Low, "Distributed batch Gaussian process optimization," in *ICML*, vol. 70, 2017, pp. 951–960.
- [29] Z. Wang, C. Li, S. Jegelka, and P. Kohli, "Batched high-dimensional bayesian optimization via structural kernel learning," in *ICML*, 2017.
- [30] J. González, Z. Dai, P. Hennig, and N. Lawrence, "Batch bayesian optimization via local penalization," in *AISTATS*, 2016, pp. 648–657.
- [31] A. Shah and Z. Ghahramani, "Parallel predictive entropy search for batch global optimization of expensive objective functions," in *NIPS*, 2015, pp. 3330–3338.
- [32] R. T. Haftka, D. Villanueva, and A. Chaudhuri, "Parallel surrogate-assisted global optimization with expensive functions – a survey," *Structural and Multidisciplinary Optimization*, vol. 54, no. 1, pp. 3–13, Jul 2016.
- [33] J. Wu and P. Frazier, "The parallel knowledge gradient method for batch bayesian optimization," in *NIPS*, 2016, pp. 3126–3134.
- [34] J. M. Hernández-Lobato, J. Requeima, E. O. Pyzer-Knapp, and A. Aspuru-Guzik, "Parallel and distributed thompson sampling for large-scale accelerated exploration of chemical space," 2017.
- [35] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos, "Parallelised bayesian optimisation via thompson sampling," 2018.
- [36] Y. Abbasi-Yadkori, D. Pál, and C. Szepesvári, "Improved algorithms for linear stochastic bandits," in *NIPS*, 2011, pp. 2312–2320.
- [37] S. Agrawal and N. Goyal, "Analysis of thompson sampling for the multi-armed bandit problem," in *COLT*, 2012, pp. 39–1.
- [38] —, "Further optimal regret bounds for thompson sampling," in *AISTATS*, 2013, pp. 99–107.
- [39] E. Kaufmann, N. Korda, and R. Munos, "Thompson sampling: An asymptotically optimal finite-time analysis," in *ALT*, 2012, pp. 199–213.
- [40] A. Karandikar, N. Akhtar, and M. Mehta, *Mobility Management in LTE Heterogeneous Networks*, 1st ed. Springer Publishing Company, Incorporated, 2017.
- [41] P. Muoz, R. Barco, and I. de la Bandera, "On the potential of handover parameter optimization for self-organizing networks," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 5, pp. 1895–1905, 2013.
- [42] S. S. Mwanje and A. Mitschele-Thiel, "Distributed cooperative q-learning for mobility-sensitive handover optimization in lte son," in *2014 IEEE Symposium on Computers and Communications (ISCC)*, vol. Workshops, 2014, pp. 1–6.
- [43] N. Sinclair, D. Harle, I. A. Glover, J. Irvine, and R. C. Atkinson, "An advanced son algorithm applied to handover management within lte," *IEEE Transactions on Vehicular Technology*, vol. 62, no. 5, pp. 1883–1894, 2013.
- [44] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.