Multi-armed Bandit Algorithms on System-on-Chip: Go Frequentist or Bayesian?

S. V. Sai Santosh, and Sumit J. Darak

Abstract—Multi-armed Bandit (MAB) algorithms identify the best arm among multiple arms via exploration-exploitation trade-off without prior knowledge of arm statistics. Their usefulness in wireless radio, IoT, and robotics demand deployment on edge devices, and hence, a mapping on system-on-chip (SoC) is desired. Theoretically, the Bayesian approach-based Thompson Sampling (TS) algorithm offers better performance than the frequentist approach-based Upper Confidence Bound (UCB) algorithm. However, TS is not synthesizable due to *Beta* function. We address this problem by approximating it via a pseudo-random number generator-based approach and efficiently realize the TS algorithm on Zynq SoC. In practice, the type of arms distribution (e.g., Bernoulli, Gaussian, etc.) is unknown and hence, a single algorithm may not be optimal. We propose a reconfigurable and intelligent MAB (RI-MAB) framework. Here, intelligence enables the identification of appropriate MAB algorithms for a given environment, and reconfigurability allows on-the-fly switching between algorithms on the SoC. This eliminates the need for parallel implementation of algorithms resulting in huge savings in resources and power consumption. We analyze the functional correctness, area, power, and execution time of the proposed and existing architectures for various arm distributions, word-length, and hardware-software co-design approaches. We demonstrate the superiority of the RI-MAB over TS and UCB only architectures.

Index Terms—Intelligent architecture, Zynq SoC, Thompson sampling, multi-Armed Bandit, partial reconfiguration

I. INTRODUCTION

Multi-armed bandit (MAB) algorithms are designed to identify the best arm among multiple arms without prior knowledge of type of arm distribution and their statistics [1]–[5]. They achieve this by optimizing an exploration-exploitation trade-off over a finite horizon (i.e. time slots) [4], [5]. Here, exploration refers to selection of all arms sufficient number of times to accurately learn their statistics and exploitation refers to selection of the best arm as often as possible. Their applications include online advertisement selection to increase the number of clicks [6], [7], clinical trials to identify best drugs [6], [7], news personalization [6], [7], decision making in financial markets [6], [7], and resource selection in wireless networks [8]–[10], internet of things (IoT) [11]–[14] and robotics [15], [16].

The performance metric for the MAB algorithm is the regret which is proportional to the number of selection of the sub-optimal arms and it should be as low as possible [1]–[5]. An optimal MAB algorithm guarantees a logarithmic regret which is the best one can achieve. The upper confidence bound (UCB) algorithm [4], Kullback Leibler UCB (KL-UCB) [2], and Thompson Sampling (TS) [3] are the popular optimal MAB algorithms. The KL-UCB algorithm is computationally

complex due to underlining optimization routine and hence, TS and UCB algorithms are preferred in [17].

In wireless radio, IoT, and robotics applications, MAB algorithms are used for decision-making tasks in the media-access control (MAC) layer. With tight integration of MAC and physical (PHY) layer, there is an option to accelerate the MAB and other MAC algorithms on hardware such as ASIC or FPGA instead of sequential processors based software implementation. The hardware implementation exploits the parallel architecture thereby offering lower execution time i.e. latency. Such applications also demand deployment on edge devices and hence, mapping of the MAB algorithms on system-on-chip (SoC) is desired. Recently, we discussed the implementation of UCB and KL-UCB algorithms on the heterogeneous Zynq SoC from Xilinx consisting of the dual-core ARM processor and 7-series FPGA [17]. In this work, we focus on the TS algorithm which has not been realized on SoC yet.

Considering two types of distributions, Bernoulli and Gaussian, frequentist approach based UCB algorithm remains the same while there are two variants of KL-UCB and Bayesian approach based TS algorithm, one each for Bernoulli and Gaussian distribution [2]–[7]. When arm distribution is known, we can select the appropriate algorithm and both guarantee lower regret than the UCB algorithm [3]. When arm distribution is unknown, the challenge is to decide the correct variant of the KLUCB/TS algorithm, and error in algorithm selection leads to significant degradation in performance. Specifically, the use of the Bernoulli variant of the KLUCB/TS algorithm for Gaussian distribution or vice-versa leads to high regret compared to the UCB algorithm [18], [19]. This demands intelligence to select the appropriate algorithm in an unknown environment.

In this paper, we design and implement a reconfigurable and intelligent architecture for MAB algorithms (RI-MAB) that can learn and select an appropriate algorithm in an unknown environment so as to minimize regret. The main contributions of this paper are summarized as below:

- 1) We propose a synthesizable TS algorithm for Bernoulli distribution (BTS) by approximating the *Beta* function via pseudo-random number generator-based approach and map it on Zynq SoC via hardware-software co-design. The architecture is optimized to reduce the computational complexity without compromising on the regret performance.
- 2) For an environment with unknown arm distribution, we propose RI-MAB architecture. Here, intelligence enables the identification of appropriate algorithm (UCB or BTS) for a given environment, and reconfigurability allows anytime on-the-fly switching between UCB and BTS algorithms via dynamic partial reconfiguration (DPR). The DPR on FPGA eliminates the need for

This work is supported by the funding received from core research grant (CRG) awarded to Dr. Sumit J. Darak from DST-SERB, GoI.

S. V. Sai Santosh, and Sumit J. Darak are with Electronics and Communications Department, IIT-Delhi, India-110020 (e-mail: [siripurapu17197,sumit]@iitd.ac.in)

parallel implementation of algorithms resulting in huge savings in resources and power consumption.

- 3) The functional correctness, resource requirement, power consumption, and execution time of the proposed BTS and RI-MAB architectures are analyzed for various arm distributions, word-length, and hardware-software co-design approaches.
- 4) We also demonstrate the superiority of the proposed RI-MAB architecture over BTS and UCB only architectures.

The rest of the paper is organized as follows. The MAB problem setup and a review of the relevant works are done in Section II followed by the synthesizable TS algorithm in Section III. The improved version of the TS algorithm and its architecture on SoC is presented in Section IV. In Section V, in-depth performance analysis and comparison with the UCB algorithm is done. The RI-MAB algorithm and its architecture are discussed in Section VI followed by its performance analysis in Section VII. Section VIII concludes the paper. Please refer to [20] for source codes and tutorial to reproduce results presented in this paper.

II. MAB ALGORITHMS AND STATE-OF-THE-ART REVIEW

In this section, we discuss the MAB setup, review state-of-the-art MAB algorithms, and feasibility on the SoC.

In MAB setup, each experiment consists of a horizon of N , $n \in \{1, 2, \dots, N\}$ sequential slots and the aim is to select the optimal arm from K , $k \in \{1, 2, \dots, K\}$ arms as often as possible. Let's denote the arm selected in slot n as I_n , and the reward received from the selected arm I_n in slot n as R_n . The reward of an arm k is generated from a distribution with mean, μ_k . The mean rewards are unknown and the performance metric, regret, is given as [2]–[7]

$$Regret = N * \mu^* - \sum_{n=1}^N R_n = N * \mu^* - \sum_{k=1}^K T_k * \mu_k \quad (1)$$

where μ^* is the mean reward of an optimal arm and T_k is the number of times the arm k selected in an experiment of horizon size N . Note that the distribution of arm rewards is unknown but fixed over a horizon. In this paper, we focus on Bernoulli (reward R_n is either 0 or 1) and Gaussian distribution (reward R_n is between 0 and 1) though the discussion can be extended to Exponential and Poisson distributions.

As discussed in Section I, UCB, KL-UCB, and TS algorithms are the popular regret-minimization MAB algorithms with logarithmic regret guarantees. In the case of the UCB and KL-UCB algorithms, the first phase is initialization where each arm is selected once in the first K time slots. Thereafter, in each slot, quality factor (QF) $Q(k, n)$ is calculated for each arm. For UCB, the QF, $Q_u(k, n)$, is given by, [4]

$$Q_u(k, n) = \frac{X(k, n)}{T(k, n)} + \sqrt{\frac{\alpha \log(n)}{T(k, n)}} \quad (2)$$

where

$$X(k, n) = X(k, n-1) + R_{n-1} * \mathbf{1}_{\{I_{n-1}=k\}} \quad \forall k \quad (3)$$

$$T(k, n) = T(k, n-1) + \mathbf{1}_{\{I_{n-1}=k\}} \quad \forall k \quad (4)$$

where $\mathbf{1}_{cond}$ is an indicator function and it is equal to 1 (or 0) if the condition, $cond$ is TRUE (or FALSE). The parameter

$X(k, n)$ is the total reward received using the arm k which has been selected for $T(k, n)$ time slots in total n time slots. The parameter, α , is the exploration factor that quantifies the aggression by which the UCB algorithm explores all arms and theoretically, it lies between 0.5 and 2. Then, the arm with the highest QF is selected and it is denoted by, I_n .

$$I_n = \arg \max_k Q_u(:, n) \quad (5)$$

After the arm I_n is played, its parameters, T and X , are updated using the received reward, R_n as shown in Eq. 3 and Eq. 4. The KL-UCB algorithm is similar to UCB except for the calculation of QF which is denoted as $Q_{kl}(k, n)$ [2]. As shown in Eq. 6, QF is computationally complex due to underlining optimization function [2], [17].

$$Q_{kl}(k, n) = \max \left\{ q \in [0, 1], d\left(\frac{X(k, n)}{T(k, n)}, q\right) \leq Y(k, n) \right\} \quad (6)$$

where

$$Y(k, n) = \frac{\log n + c \log(\log n)}{T(k, n)} \quad (7)$$

$$d(p, q) = p \log\left(\frac{p}{q}\right) + (1-p) \log\left(\frac{1-p}{1-q}\right) \quad (8)$$

The term, $d(p, q)$, in Eq. 8 denotes the KL divergence between p and q . The TS algorithm does not need an initialization phase and it uses *Beta* function for QF calculation. It is discussed later in Section II.

All these algorithms have been extended for various other settings. The multi-play setting is the same as above except that the aim is to identify the best L arms instead of one arm [21], [22]. In a time-limited pure exploration setting, the aim is to identify the best arm within a given number of time-slots such that the regret incurred during these slots is not considered i.e. pure exploration phase [23]–[25]. In a confidence-driven pure exploration setting, the aim is to identify the best arm with given confidence and in as few time slots as possible [23]–[25]. In a delayed and complex case, the reward of the selected arm in time slot n is delayed by few time slots and such delay is not deterministic [26]–[28]. Also, the received reward might be a function of arms selected in multiple time slots instead of a separate reward for each slot [26]–[28]. Existing works mainly focus on the design and performance analysis of these algorithms while the focus of this work is on efficient mapping of MAB algorithms on the SoC. Since all these extensions are based on UCB/KLUCB/TS algorithms, an efficient implementation of these three algorithms is the first and important step towards the realization of all other algorithms on the SoC.

In [29], we discussed the mapping of the UCB algorithm and its extensions on Zynq SoC via a hardware-software co-design approach. In [17] we proposed the modified KL-UCB algorithm by replacing optimization function in Eq. 6 with finite-iteration based synthesizable function. Though KL-UCB offers lower regret, the resource, latency, and power consumption of the KL-UCB is high compared to the UCB. To reduce the latency and power consumption without compromising on the regret performance, we proposed reconfigurable KL-UCB architecture that enables on-the-fly switch from KL-UCB to light-weight UCB after initial exploration [17]. In the proposed architecture, UCB QF calculation is accomplished using the KL-UCB QF blocks and hence, parallel implementation of two architectures is not needed. Since the TS is the most popular MAB algorithm, efficient mapping of the TS on SoC and

performance analysis for different word-length is an important research problem. Furthermore, an intelligence to identify the appropriate algorithm in an unknown environment is critical to get optimal regret performance. The work presented in this paper offers innovative solutions to these challenges.

III. SYNTHESIZABLE THOMPSON SAMPLING ALGORITHM FOR BERNOULLI DISTRIBUTION (SBTS)

The frequentist modeling-based UCB and KLUCB algorithms assume the mean reward of an arm is proportional to the average reward in repeated plays of a given experiment [2]–[7]. On the other hand, the Bayesian modeling-based TS algorithm assumes the mean reward of an arm is proportional to a degree of belief that the arm is optimal [3]. These beliefs are updated based on the observations from the environment via Baye's rule that takes a prior belief as an argument and returns a posterior belief for a given likelihood. Since the arm statistics are unknown, the uncertainty about arm optimality is modeled as probabilities and the arm with the highest probability of being optimal under the posterior distribution is selected [3].

In the MAB setup, posterior belief becomes a prior in subsequent time slots, and the distributions which exhibit such behavior are known as conjugate prior. For example, Beta distribution is a conjugate prior for Bernoulli likelihood function [3]. Similarly, Gamma and Pareto distributions are a conjugate prior for Poisson and Gamma distributions, respectively [3]. Thus, the Bayesian approach needs to explicitly specify prior beliefs upfront in the form of the distribution and hence, each likelihood distributions have a specific variant of the TS algorithm. None of these TS variants are realized on the SoC yet and the proposed work on the implementation of the BTS algorithm on SoC is the first contribution in this direction.

The mapping of the BTS algorithm on the SoC consists of three steps: 1) Parameter update (Eq. 3 and Eq. 4), 2) QF Calculation, 3) Arm selection (Eq. 5). Since steps 1 and 3 are identical to UCB and KL-UCB algorithms, we request readers to refer to [17], [29] for in-depth understanding and implementation. Due to limited page constraints, the discussion is focused only on Step 3: QF calculation though our implementation and tutorials include all three steps. In the BTS algorithm, the QF for each arm, denoted by $Q_{ts}(k, n)$, is calculated by drawing the random sample from the Beta distribution with parameters, $\alpha = X(k, n)$ and $\beta = T(k, n) - X(k, n)$. For the Bernoulli distribution, α refers to a number of successes and β refers to a number of failures. The probability distribution function (PDF), y_{beta} , of Beta distribution is given by [3],

$$y_{beta} = f(x|\alpha, \beta) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1} I_{[0,1]}(x) \quad (9)$$

where $B(\cdot)$ is the Beta function. The indicator function $I_{[0,1]}(x)$ ensures that only values of $x \in (0, 1)$ have nonzero probability. The QF calculation of the BTS algorithm involves two steps: 1) Integration of the PDF, y_{beta} given in Eq. 9, over x to generate the cumulative distribution function (CDF), $F(x|\alpha, \beta)$ and computation of its inverse $F^{-1}(x|\alpha, \beta)$. 2) Generation of a uniformly distributed random number x and its substitution into inverse CDF. The value obtained is the random number sampled from the Beta distribution and it is considered as the QF of the arm.

The implementation of the above steps is computationally intensive and not well suited for hardware implementation due to the need for *gamma* random generators followed by

the division of random numbers. Please refer to the in-built Matlab function, *betarnd* for more details. In the proposed approach, we approximate the *betarnd* function using an alternative synthesizable function. In each time slot n , we generate $T(k, n)$ uniform random numbers for each arm k . These random numbers are sorted in ascending order and $X(k, n)^{th}$ random number in the sorted array is taken as the QF of the arm k . This approach needs the generation of random numbers in hardware and we implement a popular Mersenne Twister pseudo-random number generator (PRNG) due to its high throughput [30]–[32]. We referred to it as a synthesizable BTS (SBTS) algorithm.

In Fig. 1, the functionality of the proposed SBTS algorithm is verified by comparing its regret performance with the BTS algorithm realized using the *betarnd* function. We consider $K = 6$, $N = 10000$ and 150 experiments. In each experiment, arm statistics are chosen randomly and the plots in Fig. 1 include the cumulative regret averaged over 150 experiments, and standard deviation (shown with a shaded region). We observed that the proposed approach selected the best arm on an average 9436 number of times ($\approx 94\%$) compared to 9445 times ($\approx 94\%$) in the *betarnd* based BTS algorithm. The regret of both algorithms is nearly identical validating the correctness of the proposed QF calculation approach.

Fig. 1: Cumulative regret averaged over 150 independent experiments for the BTS algorithm with QF using *betarnd* and proposed SBTS algorithm.

IV. IMPROVED SBTS ALGORITHM FOR EFFICIENT MAPPING ON SoC

The proposed SBTS QF calculation approach is synthesizable on the SoC but it suffers from two drawbacks:

- 1) In each time slot, a large number of random numbers need to be generated. For K arms, we need to generate $\sum_{k=1}^K T(k, n) = n$ random numbers in each time slot n . This is followed by K sorting operations, one for each arm. In the worst case (when $n=N$ i.e. end of the horizon), we need to generate and sort N random numbers in a slot. Thus, the time required to generate random numbers is not fixed and increases with n . This is not desirable for most applications.
- 2) Even if each random number is represented using fewer bits, say 8 bits, we need total storage of N bytes (For example, 10 Kbytes when $N=10000$) which is not feasible due to cost and area constraints of the majority of the embedded applications.

Ideally, MAB execution time should be fixed and as small as possible. The time taken by the MAB algorithm to select

the arm affects the time available for subsequent tasks. For instance, in wireless radio, communication is time-slotted which means arm (channel) selection is followed by transmission in each time slot. The higher the time taken for channel selection, the lower is the time available for actual data communication resulting in lower throughput. In the BTS algorithm, the time required to calculate QF for all arms increases with time due to an increase in the number of random numbers and subsequent sorting tasks. Please refer to Section for the detailed resource requirement and execution time comparison. To overcome the above drawbacks of the SBTS algorithm, we present a further improvements to simplify the sorting operation and minimize the number of random number generation in each slot.

A. SBTS-ES: SBTS Algorithm with Efficient Sorting

In the MAB setup with K arms, the SBTS algorithm involves sorting of K arrays consisting of random numbers between 0 and 1. For arm k , array size is $T(k, n)$ with its maximum value as N . After sorting, random value at the $X(k, n)^{th}$ index of the sorted array is considered as QF of the arm. For accurate QF calculation, floating-point representation of random numbers is preferred which results in computationally complex sorting operation.

In the proposed SBTS-ES algorithm, the sorting operation is simplified by grouping the random numbers in pre-defined ranges and keeping the track of number of random numbers generated in each range. For illustration, consider an array, β_k , of size 10 such that $\beta_k(1)$ represents the number of random numbers out of $T(k, n)$ lies between 0 and 0.1, $\beta_k(2)$ represents the number of random numbers out of $T(k, n)$ lies between 0.1 and 0.2. In the same fashion, $\beta_k(10)$ represents the number of random numbers out of $T(k, n)$ lies between 0.9 and 1. For k^{th} arm with $T(k, n) = 4$ and $X(k, n) = 2$, four random numbers are generated in time slot n . Lets assume these random numbers as $\{0.342, 0.012, 0.753, 0.553\}$. Then, $\beta_k = \{1, 0, 0, 1, 0, 1, 0, 1, 0, 0\}$. The $X(k, n)^{th}$ non-zero value lies in the $\beta_k(3)$ and hence, QF of the arm is equal to the mean of $\beta_k(3)$ range i.e. $\frac{0.3+0.4}{2} = 0.35$. For random numbers as $\{0.342, 0.012, 0.083, 0.553\}$, we have $\beta_k = \{2, 0, 0, 1, 0, 1, 0, 0, 0, 0\}$ and hence, the QF of the arm is equal to mean of $\beta_k(1)$ range i.e. $\frac{0+0.1}{2} = 0.05$.

In Fig. 2, we consider $K = 3$ arms. In the first K time slots, each arm is selected once. In time slot 4, $T(k, n)$ random numbers are generated for k^{th} arm. In Fig. 2, one random

number is generated separately for each arm. After updating respective β_k , QF is calculated for each arm and the third arm is selected due to the highest QF. As shown in Fig. 2, received reward is 0 in time slot 4 and hence, only $T(3, 5)$ is updated i.e. $T(3, 5) = T(3, 4) + 1$. The rest of the parameters do not change. Note that β_k of all arms is initialized to zero at the beginning of each time slot. In time slot 5, one random number is generated for the first two arms and two random numbers are generated for the third arm. The same process of QF calculation, arm selection, and parameter update are repeated in each time slot till the end of the horizon. The advantages of the proposed SBTS-ES are:

- 1) Instead of $\sum_{k=1}^{K=N} T(k, n)$ i.e. at most N floating-point random numbers, the storage of only $|\beta|K$ integer numbers with word length of $\lceil \log_2 N \rceil$ is needed. Here, $|\beta| = |\beta_1| = |\beta_2| = \dots = |\beta_K|$. SBTS needs $32KN$ bits while SBTS-ES needs $|\beta|K \lceil \log_2 (N+1) \rceil$. For $N \geq 100$ and $2 \leq |\beta| \leq 100$, SBTS memory requirement is at least 3 kilo Bytes (KB) higher than SBTS-ES. For $N > 1000$ and $N > 2000$, respectively, the difference is at least 50 KB and 80 KB, respectively.
- 2) In SBTS, sorting operations need multiple read and write memory operations for each random number which limits the execution time due to the limited number of memory ports. SBTS-ES needs only a single read and write operation per random number from a given index of β .
- 3) In SBTS, each random number is compared with at most $(T(k, n) - 1)$ random numbers and hence, at most $(N - 1) \sum_{k=1}^K T(k, n) \approx N^2$ floating point comparisons are done in each time slot. In SBTS-ES, only $2KN|\beta|$ comparisons are needed per slot. It is obvious to note that $N \gg 2K|\beta|$ for sufficiently large horizon. For instance, for $K \leq 20$ and $|\beta| \leq 100$, SBTS-ES is superior for $N > 4000$.

B. SBTS-ESSR: SBTS-ES Algorithm with Single Random Number Sample

In the SBTS-ES algorithm, total $\sum_{k=1}^K T(k, n)$ floating-point random numbers need to be generated in each time slot. This means total N random numbers will be generated in the last time slot of the horizon. Generating such a huge number of random numbers is a time-consuming, memory-intensive, and inefficient approach. Since only one arm is selected in each

Fig. 2: Illustrative example demonstrating the functionality of the SBTS-ES algorithm.

time slot, the parameter $T(k, n)$ of all arms except the selected arm will remain unchanged. Though $T(k, n)$ random numbers are needed to calculate QF of an arm k , SBTS, and SBTS-ES algorithms discard previous generated random numbers. This is inefficient since instead of generating all random numbers in each slot, we can use the random numbers from previous slots as well. Furthermore, separate random number generators for each arm can be avoided.

In the proposed SBTS-ESSR approach, β is not initialized at the beginning of each slot. In each time slot, a single random number is generated followed by updation of parameter β for all arms. To incorporate a new random number, we discard any one of the entries in β and update it with a newly generated random number. For illustrations, consider two arms with $\beta_1 = \{0, 1, 0, 0, 1, 1, 0, 0, 1, 1\}$ and $\beta_2 = \{0, 1, 0, 1, 1, 0, 0, 1, 0, 0\}$ in time slot 9. Then, $T(1, 9) = 5$ and $T(2, 9) = 4$. Assume $X(1, 9) = 2$ and $X(2, 9) = 3$. In the SBTS algorithm, 5 random numbers are generated for arm 1 followed by sorting and selection of $X(1, 9)^{th}$ random number as the QF of the arm. The same process is repeated for arm 2 with 4 random numbers. In the SBTS-ES algorithm, existing β is discarded and 9 random numbers are generated. Parameters, β_1 , and β_2 are updated using these random numbers followed by QF calculation as discussed in Section. In the SBTS-ESSR algorithm, the first step is to randomly remove one sample from β_1 and β_2 instead of discarding them completely. Then, a single random number is generated which is used to update β_1 and β_2 . After that, QF is selected using the same approach as in the SBTS-ES algorithm. It is important to note that SBTS-ESSR is a functional equivalent to SBTS-ES since the former generates $T(k, n)$ random numbers in a one-time slot while the latter uses $(T(k, n) - 1)$ random numbers generated in the previous time slots and only one random number is generated in the current time slot. Compared to SBTS-ES, SBTS-ESSR reduces the number of random number generations in each time slot as well as the number of comparisons by a factor N i.e. from $2KN/|\beta|$ to $2K/|\beta|$. Furthermore, the execution time of the SBTS-ESSR algorithm is same in each time slot compared to SBTS and SBTS-ES algorithms where execution time in each slot increases with the increase in the index of the time slot.

The SBTS-ESSR algorithm is given in Algorithm 1. In the beginning, all elements of X and T are initialized to 1 assuming an initial uniform prior i.e. all arms have equal probability of being optimal. For clarity of notations, the subscript n is removed in X and T . In each time slot of the horizon, the QF of each arm is calculated (Line 2) and the arm with the highest QF is selected (Line 3). The selected arm is played and the algorithm receives the reward from the environment (Line 4). Based on the reward, parameters X and T are updated (Line 5).

The QF generation is explained using Subroutine 1. For a given $L = |\beta|$ and K , β is a matrix where each column belongs to one arm. In the first time slot, β is initialized in the same way as T (Lines 1-3). Otherwise, β is updated for the arm selected in the previous time slot (Lines 4-6) by generating a single random number. Then, one sample is removed from each column of β , and the corresponding row index is selected randomly (Lines 9 - 10). Next, a new random number is generated and β of all arms is updated (Lines 11-13). Using updated β and parameter X , the QF is calculated for each arm (Lines 14-15).

The environment generates the reward in each slot based

Algorithm 1 SBTS-ESSR Algorithm

Input: K, N

Initialize: $X = [1]_{1 \times K}, T = [1]_{1 \times K}$

Output: *Regret*

- 1: **for** $n = 1, 2, \dots, N$ **do**
 - 2: Calculate $Q_{ts}(:, n) = \text{QF_SBTS_ESSR}(X, T, K, n, I_{n-1})$
 - 3: Select and play arm, $I_n = \arg \max_k Q_{ts}(:, n)$.
 - 4: Receive Reward, $R_n = \text{Reward}(I_n)$.
 - 5: Update X and T : $X(I_n) = X(I_n) + R_n$, $T(I_n) = T(I_n) + 1$
 - 6: **end for**
 - 7: Calculate regret using Eq. 1.
-

Subroutine 1: QF for SBTS-ESSR (QF_SBTS_ESSR)

Input: X, T, K, n, I_{n-1}

Initialize: $L = \text{No. of divisions in the } [0, 1] \text{ space i.e. } |\beta|$

Output: $Q_{1 \times K}$

- 1: **if** $n == 1$ **then**
 - 2: $\beta = [1]_{L \times K}$
 - 3: **else**
 - 4: Generate a random number, p between 0 and 1.
 - 5: $\beta_{index} = l$ if $(l - 1)/L \leq p < l/L$ where $l = 1, 2, \dots, L$.
 - 6: Update $\beta(\beta_{index}, I_{n-1}) = \beta(\beta_{index}, I_{n-1}) + 1$.
 - 7: **end if**
 - 8: **for** $k = 1, 2, \dots, K$ **do**
 - 9: Generate an integer random numbers, s , between 0 and L .
 - 10: Update $\beta(s, :) = \beta(s, :) - 1$.
 - 11: Generate a random number, p between 0 and 1.
 - 12: $\beta_{index} = l$ if $(l - 1)/L \leq p < l/L$ where $l = 1, 2, \dots, L$.
 - 13: Update $\beta(\beta_{index}, :) = \beta(\beta_{index}, :) + 1$.
 - 14: Compute $Q_{index} : \arg \min_l \sum_{l=1}^L \beta(:, k) \geq X(k)$
 - 15: $Q(k) = (2Q_{index} - 1)/(2L)$
 - 16: **end for**
-

on the selected arm and the corresponding process is given in Subroutine 2. Rewards are either 1 or 0 i.e. Bernoulli distribution.

Subroutine 2: Reward (Environment)

Input: I_n

Initialize: μ (Known to environment only)

Output: R_n

- 1: Generate random number, p between 0 and 1.
 - 2: **if** $p \leq \mu(I_n)$ **then**
 - 3: $R_n = 1$
 - 4: **else**
 - 5: $R_n = 0$
 - 6: **end if**
-

In Fig. 3, we repeat the experiments similar to Fig. 1 and compare the regret of the three proposed algorithms. As expected, the regret of the SBTS-ESSR is highest followed by SBTS-ES and SBTS. However, the difference in the regret is less than 12 for a horizon size of $N = 10000$. On average, the best arm was selected 9421 (94%), 9346 (93.5%), 9271 (92%) number of times. These results validate the functional correctness of the proposed algorithms in the MAB setup i.e. ability to identify and select the best arm as many times as possible.

The proposed algorithms are mapped on the ZSoC platform and the corresponding architecture is shown in Fig. 4. The architecture is designed and implemented using Vivado 2019.1,

Fig. 3: Cumulative regret averaged over 150 independent experiments.

Vivado High-Level Synthesis (HLS), Vivado Software Development Kit (SDK), and DPR toolbox. The parameter update, QF calculation, and arm selection are realized in FPGA (PL) while reward generation is part of the ARM processor (PS). We have explored various other configurations via hardware-software co-design. Also, the WL of various blocks realized in FPGA is carefully chosen so as to optimize the resource utilization and power consumption without compromising on the regret performance. Corresponding results are presented in Section V. The section of the proposed architecture realized on FPGA is made reconfigurable via DPR. Specifically, the number of active arms, K , and $|\beta|$ can be dynamically configured via processor configuration access port (PCAP) using the partial bit-streams pre-loaded in the SD card [33], [34]. The required bitstreams are sent to the FPGA, through the bare-metal application deployed on the ARM processor, for reconfiguration using the device configuration (DevC) direct memory access (DMA). Please refer to [20] for source codes and tutorial explaining the building blocks of the proposed architectures.

Fig. 4: Proposed DPR-enabled architecture for SBTS and its extensions with reconfigurable $|\beta|$ and no. of arms.

V. PERFORMANCE ANALYSIS: SBTS ALGORITHMS

In this section, we verify the functional correctness of the proposed SBTS algorithms on Zynq SoC and compare its regret performance with state-of-the-art UCB algorithms for different WLs. The rewards are assumed to have Bernoulli distribution. All results are obtained after averaging over 100 different experiments to consider the non-deterministic nature of the online machine learning algorithms. Later, the resource utilization, power consumption, and execution time of these

algorithms are analyzed. MAB algorithms such as KL-UCB, UCB_v, and UCB_t are not considered since UCB offers regret which is close to these algorithms with significant savings in resources, power consumption and execution time [17], [29].

A. Regret Comparison

Similar to Fig. 1 and Fig. 3, we repeat the experiments for $K = 4$ and $K = 8$ for algorithms realized on ZSoC with single-precision floating-point (SP-FL) WL. In Fig. 5, we consider four algorithms: 1) UCB, 2) SBTS, 3) SBTS-ES ($|\beta| = \{10, 20\}$), and 4) SBTS-ESSR ($|\beta| = \{10, 20\}$). It can be observed that the proposed SBTS algorithms offer significantly lower regret than UCB. This is expected since the TS algorithm has shown to outperform the UCB algorithm in analytical and simulation results. It can be observed that the regret of SBTS-ES and SBTS-ESSR algorithms decreases with an increase in $|\beta|$. This is because higher $|\beta|$ allows accurate calculation of QF leading to a reduction in the selection of sub-optimal arms.

Fig. 5: Average cumulative regret comparison for randomly generated arm distributions, (a) $K = 4$, and (b) $K = 8$.

Next, we analyze the regret performance for four cases of carefully selected arm distributions.

- 1) $K = 4$, $\mu_1 = \{0.1, 0.3, 0.5, \mathbf{0.7}\}$
- 2) $K = 4$, $\mu_2 = \{\mathbf{0.54}, 0.53, 0.52, 0.51\}$
- 3) $K = 8$, $\mu_3 = \{0.1, 0.5, \mathbf{0.8}, 0.7, 0.4, 0.2, 0.6, 0.3\}$
- 4) $K = 8$, $\mu_4 = \{0.21, 0.22, 0.26, \mathbf{0.28}, 0.24, 0.25, 0.27, 0.23\}$

The optimal arm in each case is highlighted in bold font. Compared to μ_1 and μ_3 , the difference between the arm statistics is small in μ_2 and μ_4 . This makes the learning and identification of the optimal arm challenging. As shown in Fig. 6, the regret of the SBTS, SBTS-ES, and SBTS-ESSR algorithms is lower than that of the UCB algorithm for all arm distributions. In each case, it is verified that the optimal arm is chosen the highest number of times by all algorithms. The appropriate selection of $|\beta|$ is important as it affects the precision of QF selection. This results in multiple arms with identical QF values and hence, frequent selection of sub-optimal arms. For instance, $|\beta| = 10$ is not sufficient for μ_2 as it offers high regret as shown in Fig. 6 (d). Based on extensive performance analysis, $15 \leq |\beta| \leq 20$ leads to a higher number of selection of optimal arm i.e. lower regret and the gain in performance is not significant for $|\beta| > 20$. The proposed architecture in Fig. 4 allows on-the-fly selection of $|\beta|$ via DPR. Next, we compare the effect of WL on the regret performance of the SBTS-ESSR ($|\beta| = 20$) and UCB algorithms. In Fig 7, we compare the regret of these algorithms at the end of the horizon of size, $N = 10000$ for μ_1 and μ_3 . We compare the regret for SP-FP and fixed-point implementations with a total WL of 27, 11, and 6 bits. In each case, the number of bits to represent integer and fractional parts are chosen carefully

Fig. 6: Average cumulative regret of different algorithms for (a) μ_1 , (b) μ_2 , (c) μ_3 , and (d) μ_4 .

so as to minimize regret. It can be observed that the regret degrades with the decrease in WL. However, degradation is not significant till WL=11. For WL=6, regret is high and this happens due to insufficient bits to represent QF which in turn leads to the selection of sub-optimal arms. Thus, the selection of appropriate WL is important since lower WL leads to significant savings in resources but it should not come at the cost of regret performance.

B. Complexity Comparison

In Table I, we compare the resource utilization (LUT, FFs, DSP and BRAM), and power consumption of four different architectures realized on ZSoC. These architectures correspond to SBTS, SBTS-ES, SBTS-ESSR, and UCB algorithms. Each architecture is made reconfigurable via DPR at the arm level i.e. the number of arms can be dynamically configured (i.e., the number of arms can be tuned to any values less than or equal to K_{max}). Each architecture is realized with three different WLs: SP-FP, fixed point WL of 27, and 11 bits. The reconfigurable architecture is compared with non-reconfigurable Velcro approach-based architecture with fixed K_{max} arms and SP-FP WL.

Fig. 7: Average regret of SBTS-ESSR and UCB algorithms with different WLs for (a) μ_1 , and (b) μ_3 .

As shown in Table I, resource utilization and power consumption of DPR based architecture depend on the number of active arms, K compared to the Velcro approach, which corresponds to architecture with K_{max} arms, i.e., all blocks are active all the time compared to dynamic activation and deactivation of arms in the DPR based architecture. It can be observed that the SP-FP version of the DPR-based architecture offers around savings of 18-39% in LUTs, 17-38% in FFs, 25-44% in DSP48Es, and 20-40% in BRAM over the Velcro approach for $K < K_{max}$. Further, they offer a 19-40% reduction in the dynamic power consumption over the Velcro approach for $K < K_{max}$. Using the fixed-point implementation with WL=27, one can achieve up to 83%, 79%, 100%, 100% savings in the consumption of LUTs, FFs, DSP48Es, and BRAM_18Ks respectively for $K < K_{max}$. This can be achieved with almost identical performance as that of the SP-FP architecture. With WL=11, further improvement in savings of up can be achieved with a slight degradation in the regret. In terms of the dynamic power consumed, the architectures with fixed-point WLs offer up to 94% savings. Among SBTS algorithms, the SBTS-ESSR algorithm offers significant savings in resource utilization as well as power consumption. Compared to the UCB algorithm, SBTS-ESSR is computationally efficient, offers lower regret and power consumption. This makes the proposed SBTS-ESSR a superior alternative to the state-of-the-art UCB al-

TABLE I: Comparison of Resource Utilization and Power Consumption

No.	Algorithm	$ \beta $	DPR/Velcro Precision	No. of LUTs			No. of FFs			No. of DSPs			No. of BRAMs			Dynamic Power (W)		
				K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5
1	SBTS, tunable $K \leq K_{max}$	NA	DPR	14952	19488	24024	11804	15350	18896	18	24	32	72	96	120	0.258	0.339	0.42
			SP-FP	-38%	-19%		-38%	-19%		-44%	-25%		-40%	-20%		-39%	-20%	
			Velcro		24024			18896			32			120			0.42	
2	SBTS-ES, tunable $K \leq K_{max}$	10	DPR	14634	19064	23494	11441	14866	18291	18	24	32	24	32	40	0.178	0.234	0.29
			SP-FP	-38%	-19%		-38%	-19%		-44%	-25%		-40%	-20%		-39%	-20%	
			Velcro		23494			18291			32			40			0.29	
		20	DPR	20196	26480	32764	13488	17542	21636	18	24	32	24	32	40	0.198	0.261	0.324
			SP-FP	-39%	-20%		-38%	-19%		-44%	-25%		-40%	-20%		-39%	-20%	
			Velcro		32764			21636			32			40			0.324	
3	SBTS-ESSR, tunable $K \leq K_{max}$	10	DPR	7386	9400	11414	4631	5786	6941	18	24	32	27	36	45	0.046	0.061	0.076
			SP-FP	-36%	-18%		-34%	-17%		-44%	-25%		-40%	-20%		-40%	-20%	
			Velcro		11414			6941			32			45			0.076	
		20	DPR	7041	8940	10839	4811	6026	7241	18	24	32	27	36	45	0.096	0.127	0.158
			SP-FP	-36%	-18%		-34%	-17%		-44%	-25%		-40%	-20%		-40%	-20%	
			DPR	1857	2169	2481	1792	2110	2428	0	0	0	3	4	5	0.010	0.012	0.014
			WL=27	-83%	-80%	-78%	-76%	-71%	-67%	-100%	-100%	-100%	-94%	-92%	-89%	-94%	-93%	-92%
			DPR	1864	2167	2470	1792	2110	2428	0	0	0	3	4	5	0.011	0.013	0.015
			WL=11	-83%	-80%	-78%	-76%	-71%	-67%	-100%	-100%	-100%	-94%	-92%	-89%	-93%	-92%	-91%
			Velcro		10839			7241			32			45			0.158	
4	UCB, tunable $K \leq K_{max}$	NA	DPR	7664	9866	12068	7060	9093	11126	51	68	85	4.5	6	7.5	0.127	0.165	0.203
			SP-FP	-37%	-19%		-37%	-19%		-40%	-20%		-40%	-20%		-38%	-19%	
			DPR	5761	7375	8989	4390	5574	6758	0	0	0	0	0	0	0.019	0.023	0.027
			WL=27	-53%	-39%	-26%	-61%	-50%	-40%	-100%	-100%	-100%	-100%	-100%	-100%	-91%	-89%	-87%
			DPR	2425	2927	3429	2443	2978	3513	0	0	0	0	0	0	0.013	0.015	0.017
			WL=11	-80%	-76%	-72%	-79%	-74%	-69%	-100%	-100%	-100%	-100%	-100%	-100%	-94%	-93%	-92%
			Velcro		12068			11126			85			7.5			0.203	

gorithm for the environment with Bernoulli rewards.

Execution time of the algorithm is an importance-performance metric that depends on efficient implementation and underlining architecture. In Table II, we consider three different configurations obtained by realizing the algorithm using: 1) PS and PL with optimal partitioning, 2) Only PS (ARM + NEON), 3) Only ARM. It can be observed that the first approach offers the lowest execution time due to the parallel execution of the QF function in PL compared to sequential PS execution. Furthermore, the gain improves as K increases. Among various TS algorithms, SBTS-ESSR offers the lowest execution time as expected. In applications like wireless networks, MAB algorithms are realized in upper layers (MAC/Network) i.e. in ARM or other processors while the PHY is present in the SoC. The proposed architecture enables shifting of the MAB algorithms from MAC to PHY layers thereby resulting in an accelerator factor ranging from 465.6-776 for UCB and 2.5-33.6 for TS. For larger $K > 20$, the acceleration factor will be significantly higher. Between UCB and TS algorithms, UCB is faster on ZSoC due to hardware-friendly arithmetic operations compared to PRNG in SBTS but UCB incurs high regret. On the PS-only architectures (i.e. ARM and ARM+NEON), SBTS-ESSR offers the lowest execution time.

TABLE II: Comparison of Execution Time in Milliseconds

No.	Algorithm	Nb	Precision	ZSoC (in ms)	PS (ARM) (in ms)			PS (ARM+NEON) (in ms)		
				K={3,4,5}	K=3	K=4	K=5	K=3	K=4	K=5
1	SBTS	NA	SP-FP	12504	26758	35697	44580	13060	17389	21746
2	SBTS-ES	20	SP-FP	18008	35454	47267	59079	6791	8981	11212
		10		10010	25134	33509	41882	6549	8730	10913
3	SBTS-ESSR	20	SP-FP	3.9	38	46	54	15	16	18
			WL=27	1.5						
			WL=11	1.5						
		10	SP-FP	2.7	33	40	46	11	13	14
4	UCB	NA	SP-FP	0.1	29	37	47	20	26	33
			WL=27	0.1						
			WL=11	0.1						

VI. RECONFIGURABLE AND INTELLIGENT MAB (RI-MAB)

The SBTS-ESSR algorithm is well-suited only for arm rewards with Bernoulli distribution and hence, it may not outperform the UCB algorithm when arm distribution is unknown and random [3]. For example, in wireless Radio, arm i.e. wireless channel distribution is usually Bernoulli at high signal-to-noise ratios (SNR) and Gaussian at medium and low SNRs. In IoT and robotics applications, the type of arm distribution is unknown. Such applications demand architecture which can learn, identify and deploy appropriate MAB algorithm.

Very few works have addressed this problem and [18] is one of the recent works in which the aggregator algorithm selects the arm chosen by one of the candidate MAB algorithms in each time slot. The aim is to identify an optimal algorithm for a given unknown environment and this is achieved by learning from the past performance of various algorithms i.e. algorithm exploration-exploitation trade-off in addition to arm exploitation-exploration trade-off. The major problem with [18] is the need for hardware implementation of all candidate MAB algorithms in parallel (referred here as Velcro MAB). Such architecture is area and power inefficient. The proposed RI-MAB algorithm augmented with DPR-based reconfigurable architecture overcomes this problem.

The proposed RI-MAB algorithm is given in Algorithm 2. The number of candidate algorithms, number of arms, and horizon size is A , K , and N , respectively. The RI-MAB algorithm maintains the probability distribution on all candidate algorithms to indicate their optimality. Since all algorithms are equally likely to be optimal at the beginning of each experiment, the prior belief, π_0 , is initialized as uniform distribution (Line 2). To update the belief and identify the optimal algorithm, the RI-MAB algorithm performs epoch-based exploration in the initial N_{learn} time slots (Lines 5-19). In this phase, each algorithm is selected for 2^e time slots before incrementing the parameter e by 1 (Line 12-17). This allows a sufficient number of learning samples of each algorithm before finalizing the algorithm for the rest of the horizon i.e. after N_{learn} (Line 20). Such approach allows only one MAB algorithm to be active in hardware in each

Algorithm 2 RI-MAB Algorithm

```

1: Input:  $K, N, N_{learn}, A$ 
2: Initialize:  $\pi_0 = U(\{1, 2, \dots, A\}), e = 1, alg = 1, X = [1]_{1 \times K}, T = [1]_{1 \times K}, I_0 = K, r = 0$ 
3: Output: Regret
4: for  $n = 1 : 1 : N$  do
5:   if  $n \leq N_{learn}$  then
6:      $I_n = \text{ARM\_SEL}(alg, n, I_{n-1}, X, T, K)$ 
7:     Receive Reward,  $R_n = \text{Reward}(I_n)$  #Subroutine 2
8:     Update the learning rate,  $\eta_n = \sqrt{\log A / (n \times K)}$ 
9:     Compute the unbiased reward using the belief of the selected algorithm,  $\hat{R}_n = \frac{R_n}{\pi_{n-1}(alg)}$ 
10:    Update the belief of the selected algorithm:  $\pi_n(alg) = \exp(\eta_n \hat{R}_n) \times \pi_{n-1}(alg)$ 
11:    Normalize the belief of all algorithms:  $\pi_n(i) = \pi_n(i) / \sum_{a=1}^A \pi_n(alg), i \in \{1, 2\}$ 
12:     $r = r + 1$ 
13:    if  $r == 2^e$  then
14:       $r = 0$ 
15:       $alg = (alg + 1)$  #Switch Algorithm
16:      if  $alg > A$  then
17:         $e = e + 1$  #Increase epoch size
18:         $alg = 1$  #Reset to first algorithm
19:    end if
20:  end if
21:  else #Select algorithm with higher belief
22:     $alg = \arg \max \pi_{N_{learn}} - 1$ 
23:     $I_n = \text{ARM\_SEL}(alg, n, I_{n-1}, R_{n-1})$ 
24:    Receive Reward,  $R_n = \text{Reward}(I_n)$  #Subroutine 2
25:  end if
26:  Update  $X$  and  $T$ :  $X(I_n) = X(I_n) + R_n, T(I_n) = T(I_n) + 1$ 
27: end for
28: Calculate regret using Eq. 1.

```

Subroutine 3: ARM_SEL (Arm Selection Using Chosen Algorithm)

Input: alg, n, I_{n-1}, X, T, K

Output: I_n

```

1: if  $alg == 0$  then #UCB Algorithm
2:    $Q(k, n) = \frac{X(k)}{T(k)} + \sqrt{\frac{\alpha \log(n+K)}{T(k)}}, \forall k \in \{1, 2, \dots, K\}$ 
3: else #SBTS-ESSR Algorithm (Subroutine 1)
4:   Calculate  $Q(:, n) = \text{QF\_SBTS\_ESSR}(X, T, K, n, I_{n-1})$ 
5: end if
6: Select arm,  $I_n = \arg \max_k Q(:, n)$ .

```


time slot and increasing length epoch reduces the number of algorithm switching. Even though only one algorithm is active, the parameters, X , and T , are common across all candidate algorithms which means there is no compromise on arm learning aspects of MAB setup.

During algorithm exploration (Lines 6-18), RI-MAB updates the parameter, π_n in each time slot based on the selected algorithm and received reward. Similar to [18], we obtain the unbiased estimate of received reward using the probability of the arm selection i.e. belief of the selected algorithm (Line 9). Then, the belief of the selected algorithm is updated using the exponential multiplicative factor (Line 10) and learning rate (Line 8) [18]. In the end, the beliefs of all algorithms are normalized.

The functionality of the RI-MAB algorithm is explained using Fig 8. In the beginning, the prior belief, $\pi_0 = \{0.5, 0.5\}$. The algorithm starts by selecting UCB for the first $2^1 = 2$ slots. It should be noted that we simulate the algorithm with Bernoulli arm rewards for ease of understanding. In slot 1, UCB selects the arm, $I_1 = 4$ and receives a reward, $R_1 = 1$. This results in an increase in the belief value of UCB by a factor of $\exp(\eta_1 * 1)$. The same happens in slot 2 when the belief of UCB further increases by receiving $R_2 = 1$ for the selected arm, $I_2 = 2$. In slots 3 and 4, the algorithm selects SBTS-ESSR, which receives rewards, $R_3 = 0$ and $R_4 = 0$ on arms 4 and 2 respectively. Hence, the belief of TS is not updated in both slots (as $\exp(\eta_n * 0) = 1$). This is when the parameter e is incremented by 1. Skipping over to slots 19 and 24, the algorithm selects UCB, which receives a reward, $R_{19} = 1$ and $R_{24} = 1$ respectively, improving the belief of UCB. In slots 128 and $N_{learn} = 500$, the selected candidate algorithm, SBTS-ESSR, receives a reward of 1, which improves its belief. In summary, the belief of the selected candidate algorithm increases when it receives a reward, $R_n = 1$, and is not updated when $R_n = 0$. After slot $N_{learn} = 500$, the algorithm finalizes SBTS-ESSR for the rest of the horizon, pertaining to its higher belief in slot N_{learn} .

The proposed RI-MAB algorithm is mapped on SoC and the corresponding architecture is shown in Fig. 9. Compared to Fig. 4, an additional algorithm selection unit is included in the ARM processor and the QF calculation block is reconfigured via DPR depending on the selected candidate algorithm.

VII. PERFORMANCE AND COMPLEXITY ANALYSIS: RI-MAB

We first begin with the regret comparison of the proposed RI-MAB architecture with the UCB, SBTS-ESSR, and Velcro MAB architectures. To begin with, we consider $N = 10000$,

Fig. 9: Proposed DPR-enabled architecture for RI-MAB with six candidate algorithms along with reconfigurable $|\beta|$ and K . $K = 8$ and choose μ_3 as the probability distributions. Note that rewards are modeled as Gaussian instead of Bernoulli distribution.

In Fig. 10 (a), we compare the regret of the four algorithms at different instances of the horizon. The regret is averaged over 20 independent experiments. The final regret (i.e. regret at the end of the horizon) for each experiment is shown in Fig. 10 (c). In all 20 experiments, SBTS-ESSR offers the lowest regret while UCB incurs the highest regret though both are able to identify the optimal arm. Velcro MAB offers performance

Fig. 10: Average cumulative regret for μ_3 at the end of (a) 20 experiments, and (b) 200 experiments. Final regret at the end of horizon for (c) 20 experiments, and (d) 200 experiments.

	Slot 1	Slot 2	Slot 3	Slot 4	Slot 19	Slot 24	Slot 128	Slot N_{learn}	Slot $N_{learn} +$
Algorithm, alg	UCB ✓ SBTS-ESSR	UCB ✓ SBTS-ESSR	UCB SBTS-ESSR ✓	UCB SBTS-ESSR ✓	UCB ✓ SBTS-ESSR	UCB ✓ SBTS-ESSR	UCB SBTS-ESSR ✓	UCB SBTS-ESSR ✓	UCB SBTS-ESSR ✓
Arm Selection, I_n	4	2	4	1	4	3	4	4	-
Reward, R_n	1	1	0	0	1	1	1	1	R
Belief, π_t	0.73 ↑ 0.27	0.79 ↑ 0.21	0.79 0.21	0.79 0.21	0.38 ↑ 0.62	0.47 ↑ 0.53	0.36 0.64 ↑	0.18 0.82 ↑	-
2^e	2	2	2	2	8	8	64	128	-

Fig. 8: Working of the proposed RI-MAB algorithm in an example setting with $K = 4$, $N = 10000$, $N_{learn} = 500$, $A = 2$ and randomly generated arm distribution

closer to SBTS-ESSR in most of the experiments except 16 and 20 where it takes more time to learn TS is better than UCB. On the other hand, RI-MAB selects SBTS-ESSR in all experiments. Despite that, the regret of RI-MAB is higher than Velcro MAB in all experiments except 16 and 30. This happens due to the initial N_{learn} period of exploration to choose between UCB and SBTS-ESSR algorithm. Still, averaged regret of RI-MAB and Velcro-MAB is nearly identical as shown in Fig. 10 (a) since regret incurred by Velcro-MAB is very high in experiments 16 and 20. Such high regret events are avoided in the RI-MAB algorithm.

In Fig. 10 (b) and (d), we repeat the above simulations by increasing the number of independent experiments to 200. It can be observed that Velcro MAB and RI-MAB offer lower regret than individual UCB and SBTS-ESSR algorithms. This is because, for experiments 31, 35, 125, 167, 174, 183, SBTS-ESSR fails to identify the optimal arm and hence, incurs huge regret. Also, as expected, UCB incurs slightly higher regret in all experiments. Though Velcro MAB avoids selection of SBTS-ESSR in experiments 31, 35, 125, 167, 174, 183, it wrongly selects UCB in few other experiments (i.e. experiments where the final regret of UCB is same as Velcro MAB). Though RI-MAB incurs higher average regret in experiments where SBTS-ESSR is optimal, it avoids the wrong selection of algorithm as well as optimal arm in the rest of the experiments. It can be observed that Velcro-MAB and RI-MAB offer lower regret than UCB and SBTS-ESSR in some experiments. This happens because learning due to switching between the two algorithms helps to avoid the wrong selection of optimal arm in SBTS-ESSR. These experiments demonstrate the need for careful algorithm switching and

selection for a given environment since a single algorithm does not offer optimal performance in all experiments.

Next, we select the arm statistics randomly in each of the experiments instead of fixed μ_3 in Fig. 10. We consider $N = 10000$, $K = 8$ and two types of arm statistics such that the minimum difference between two arm statistics is 1) 0.07 (Case 1), and 2) 0.025 (Case 2). Corresponding results are shown in Fig. 11 (a)-(d). It can be observed that SBTS-ESSR offers significantly higher regret in almost 25% of the experiments. Though UCB offers an overall higher average regret, it successfully identifies optimal arm in all experiments leading to better performance than SBTS-ESSR. Proposed RI-MAB offers similar performance as UCB but with lower average regret. On the other hand, Velcro UCB fails to identify the appropriate algorithm and hence, the optimal arm in some experiments while in other experiments it offers lower regret than RI-MAB. On average, both offer nearly identical average regret but RI-MAB eliminates high regret events.

To gain further insights into the regret performance, we study various statistical properties of the final regret over 100 experiments in Fig. 12. Using the Boxplot feature in Matlab, we consider median (central red line), and percentile (bottom and top edges of the box indicate the 25th and 75th percentiles, respectively). The outliers are plotted individually using the '+' symbol. It can be observed that the SBTS-ESSR and Velcro-MAB have large size boxes indicating large variation in regret performance and the number of outliers are more indicating wrong selection of algorithm or arm. On the other hand, UCB and RI-MAB guarantee the selection of optimal arm in all experiments and RI-MAB offers lower regret between them.

Fig. 11: Average cumulative regret at the end of 100 independent experiments with randomly chosen arm statistics in (a) Case 1, and (b) Case 2. Final regret at the end of horizon for (c) Case 1, and (d) Case 2.

Fig. 12: Statistical properties of the final regret of the different algorithms at the end of 100 independent experiments with randomly chosen arm statistics in (a) Case 1, and (b) Case 2.

Next, we compare the resource utilization and power consumption of RI-MAB and Velcro-MAB architectures in Table III. Both architectures are on-the-fly reconfigurable in terms of the number of arms and type of the algorithm. Proposed RI-MAB with SP-FP WL architecture offers around 44-65%, 36-59%, 27-56% and 15-49% savings in LUTs, FFs, DSPs, and BRAMs over Velcro-MAB architecture. Similarly, it offers 44-65% lower dynamic power consumption. The savings improve further when the WL is optimized via fixed-point representation. With the increase in the number of arms (i.e. for $K_{max} > 20$, the proposed architecture and DPR approach further improvements in resource utilization. In addition, savings will increase further with the increase in the number of candidate algorithms.

VIII. CONCLUSIONS AND FUTURE WORKS

In this paper, we present synthesizable and reconfigurable architecture of the Thompson Sampling (TS) multi-armed bandit (MAB) algorithm for arms with Bernoulli distribution. The functional correctness, execution time, resource, and power consumption comparisons demonstrate its superiority over the upper confidence bound (UCB) algorithm. For an environment with unknown arm distribution, a reconfigurable and intelligent MAB (RI-MAB) algorithm is proposed along with its architecture. The RI-MAB offers significant savings in resources and power consumption without compromising on the regret performance. In the future, we plan to integrate the proposed RI-MAB with wireless radio and analyze the gain in throughput using real radio signals. Other possibilities include the extension of the RI-MAB architecture for an environment

where the number of arms, as well as arm statistics, are not fixed i.e. the non-stationary environment in which an optimal arm changes over time.

REFERENCES

- [1] S. Bubeck and N. Cesa-Bianchi, *Regret analysis of stochastic and nonstochastic multi-armed bandit problems*. NOW publisher, Foundations and Trends in Machine Learning, 2012.
- [2] A. Garivier and O. Cappé, “The KL-UCB algorithm for bounded stochastic bandits and beyond,” in *Conference On Learning Theory (COLT)*, Budapest, Hungary, July 2011.
- [3] S. Agrawal and N. Goyal, “Further optimal regret bounds for thompson sampling,” in *16th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Scottsdale, USA, April 2013.
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Machine Learning*, vol. 47, no. 2, 2002.
- [5] V. Anantharam, P. Varaiya, and J. Walrand, “Asymptotically efficient allocation rules for the multi-armed bandit problem with multiple plays - part i: iid rewards,” *IEEE Transactions on Automatic Control*, vol. 32, no. 11, 1987.
- [6] A. Slivkins, “Introduction to multi-armed bandits,” in *Foundations and Trends in Machine Learning*, vol. 12, no. 1, 2019, pp. 1–286.
- [7] T. Lattimore and C. Szepesvári, “Bandit algorithms,” in *Cambridge, U.K.: Cambridge Univ. Press*, 2018.
- [8] A. Anandkumar, N. Michael, A. K. Tang, and Ananthram Swami, “Distributed algorithms for learning and cognitive medium access with logarithmic regret,” *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 4, pp. 731–745, 2011.
- [9] H. Tibrewal, S. Patchala, M. K. Hanawal, and S. J. Darak, “Distributed learning and optimal assignment in multiplayer heterogeneous networks,” in *IEEE International Conference on Computer Communications (INFOCOM)*, 2019.
- [10] S. M. Zafaruddin, I. Bistriz, A. Leshem, and D. Niyato, “Distributed learning for channel allocation over a shared spectrum,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, 2019.
- [11] M. K. Hanawal and S. J. Darak, “Multi-player multi-armed bandits for stable allocation in heterogeneous ad-hoc networks,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 10, 2019.
- [12] I. Bistriz and A. Leshem, “Distributed multi-player bandits-a game of thrones approach,” in *Advances in Neural Information Processing Systems*, 2018.
- [13] Y. Gai and B. Krishnamachari, “Distributed stochastic online learning policies for opportunistic spectrum access,” *IEEE Transactions on Signal Processing*, vol. 62, no. 23, 2014.
- [14] E. Boursier and V. Perchet, “SIC-MMAB: Synchronisation involves communication in multiplayer multi-armed bandits,” in *Neural Information Processing Systems (NeurIPS)*, 2019.
- [15] H. Claire, Y. Chen, J. Modi, M. Jung, and S. Nikolaidis, “Multi-armed bandits with fairness constraints for distributing resources to human teammates,” in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI ’20, 2020, p. 299–308.
- [16] L. Chan, D. Hadfield-Menell, S. Srinivasa, and A. Dragan, “The assistive multi-armed bandit,” in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2019, pp. 354–363.
- [17] S. V. S. Santosh and S. J. Darak, “Intelligent and reconfigurable architecture for kl divergence-based multi-armed bandit algorithms,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 68, no. 3, pp. 1008–1012, 2021.
- [18] L. Besson, E. Kaufmann, and C. Moy, “Aggregation of multi-armed bandits learning algorithms for opportunistic spectrum access,” in *2018 IEEE Wireless Communications and Networking Conference (WCNC)*, 2018, pp. 1–6.

TABLE III: Comparison of Resource Utilization and Power Consumption

Algorithm	Reconfigurability	No. of LUTs			No. of FFs			No. of DSPs			No. of BRAMs			Dynamic Power (W)		
		K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5	K=3	K=4	K=5
RI-MAB, UCB/SBTS-ESSR tunable $K \leq K_{max}$ $ \beta =20$	DPR	7664	9866	12068	7060	9093	11126	51	68	85	27	36	45	0.127	0.165	0.203
	SP-FP	-65%	-55%	-44%	-59%	-48%	-36%	-56%	-41%	-27%	-49%	-32%	-15%	-65%	-54%	-44%
	DPR	5761	7375	8989	4390	5574	6758	0	0	0	3	4	5	0.019	0.023	0.027
	WL=27	-74%	-66%	-59%	-75%	-68%	-61%	-100%	-100%	-100%	-95%	-93%	-91%	-95%	-94%	-93%
	DPR	2425	2927	3429	2443	2978	3513	0	0	0	3	4	5	0.013	0.015	0.017
	WL=11	-89%	-87%	-84%	-86%	-83%	-80%	-100%	-100%	-100%	-95%	-93%	-91%	-97%	-96%	-95%
Velcro (SP-FP)		21563			17201			115			52.5			0.358		

- [19] A. Agarwal, H. Luo, B. Neyshabur, and R. E. Schapire, "Corralling a band of bandit algorithms," in *Proceedings of the 2017 Conference on Learning Theory*, ser. *Proceedings of Machine Learning Research*, S. Kale and O. Shamir, Eds., vol. 65. PMLR, 07–10 Jul 2017, pp. 12–38. [Online]. Available: <http://proceedings.mlr.press/v65/agarwal17b.html>
- [20] "Source codes and tutorial," <https://github.com/Sai-Santosh-99/TS-Aggregator>.
- [21] J. Komiyama, J. Honda, and H. Nakagawa, "Optimal regret analysis of thompson sampling in stochastic multi-armed bandit problem with multiple plays," in *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ser. *ICML'15*. JMLR.org, 2015, p. 1152–1161.
- [22] J. Komiyama, J. Honda, and A. Takeda, "Position-based multiple-play bandit problem with unknown position bias," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 5005–5015.
- [23] S. Bubeck, R. Munos, and G. Stoltz, "Pure exploration in multi-armed bandits problems," in *Algorithmic Learning Theory*, R. Gavalda, G. Lugosi, T. Zeugmann, and S. Zilles, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 23–37.
- [24] S. Chen, T. Lin, I. King, M. R. Lyu, and W. Chen, "Combinatorial pure exploration of multi-armed bandits," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014. [Online]. Available: <https://proceedings.neurips.cc/paper/2014/file/e56954b4f6347e897f954495eab16a88-Paper.pdf>
- [25] Z. Karnin, T. Koren, and O. Somekh, "Almost optimal exploration in multi-armed bandits," in *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28*, ser. *ICML'13*. JMLR.org, 2013, p. III–1238–III–1246.
- [26] A. Grover, T. Markov, P. Attia, N. Jin, N. Perkins, B. Cheong, M. Chen, Z. Yang, S. Harris, W. Chueh, and S. Ermon, "Best arm identification in multi-armed bandits with delayed feedback," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, ser. *Proceedings of Machine Learning Research*, A. Storkey and F. Perez-Cruz, Eds., vol. 84. PMLR, 09–11 Apr 2018, pp. 833–842. [Online]. Available: <http://proceedings.mlr.press/v84/grover18b.html>
- [27] C. Vernade, A. Carpentier, T. Lattimore, G. Zappella, B. Ermiş, and M. Brückner, "Linear bandits with stochastic delayed feedback," in *Proceedings of the 37th International Conference on Machine Learning*, ser. *Proceedings of Machine Learning Research*, H. D. III and A. Singh, Eds., vol. 119. PMLR, 13–18 Jul 2020, pp. 9712–9721. [Online]. Available: <http://proceedings.mlr.press/v119/vernade20a.html>
- [28] N. Cesa-Bianchi, C. Gentile, and Y. Mansour, "Delay and cooperation in nonstochastic bandits," *Journal of Machine Learning Research*, vol. 20, no. 17, pp. 1–38, 2019. [Online]. Available: <http://jmlr.org/papers/v20/17-631.html>
- [29] S. V. S. Santosh and S. J. Darak, "Reconfigurable and computationally efficient architecture for multi-armed bandit algorithms," in *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2020, pp. 1–5.
- [30] X. Tian and K. Benkrid, "Mersenne twister random number generation on fpga, cpu and gpu," in *2009 NASA/ESA Conference on Adaptive Hardware and Systems*, 2009, pp. 460–464.
- [31] P. Beadling, A. Ferencz, and S. Banks, "Fpga implementation of pseudo random number generators for monte carlo methods in quantitative finance," in *Reconfigurable Computing and FPGAs, International Conference on*. Los Alamitos, CA, USA: IEEE Computer Society, dec 2008, pp. 271–276. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/ReConFig.2008.38>
- [32] M. Matsumoto and T. Nishimura, "Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Trans. Model. Comput. Simul.*, vol. 8, no. 1, p. 3–30, Jan. 1998. [Online]. Available: <https://doi.org/10.1145/272991.272995>
- [33] "Partial reconfiguration user guide," https://www.xilinx.com/support/documentation/sw_manuals/xilinx14_5/ug702.pdf.
- [34] "Vivado design suite tutorial: Partial reconfiguration," https://www.xilinx.com/support/documentation/sw_manuals/xilinx2019_1/ug947-vivado-partial-reconfiguration-tutorial.pdf.