

Design Space Exploration of Approximate Computing Techniques with a Reinforcement Learning Approach

Original

Design Space Exploration of Approximate Computing Techniques with a Reinforcement Learning Approach / Saeedi, S., Savino, A., Di Carlo, S.. - ELETTRONICO. - (2023), pp. 167-170. (2023 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W) Porto (PRT) 27-30 June 2023) [10.1109/DSN-W58399.2023.00047].

Availability:

This version is available at: 11583/2981399 since: 2023-08-30T11:24:56Z

Publisher:

IEEE

Published

DOI:10.1109/DSN-W58399.2023.00047

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Design Space Exploration of Approximate Computing Techniques with a Reinforcement Learning Approach

Sepide Saeedi, Alessandro Savino, Stefano Di Carlo
Politecnico di Torino, Control and Computer Eng. Dep., Torino, Italy
emails: {sepide.saeedi, alessandro.savino, stefano.dicarlo}@polito.it

Abstract—Approximate Computing (AxC) techniques have become increasingly popular in trading off accuracy for performance gains in various applications. Selecting the best AxC techniques for a given application is challenging. Among proposed approaches for exploring the design space, Machine Learning approaches such as Reinforcement Learning (RL) show promising results. In this paper, we proposed an RL-based multi-objective Design Space Exploration strategy to find the approximate versions of the application that balance accuracy degradation and power and computation time reduction. Our experimental results show a good trade-off between accuracy degradation and decreased power and computation time for some benchmarks.

Index Terms—Approximate Computing, Design Space Exploration, Reinforcement Learning.

I. INTRODUCTION

Approximate Computing (AxC) techniques have become increasingly popular to improve the energy efficiency of computations, power consumption, and computation time, degrading the accuracy of the computation. However, the reduction is possible by sacrificing the computation accuracy at different parts of the computing chain. Trading off accuracy for improved power consumption and performance is known as Approximate Computing (AxC) [1]. Selecting the most suitable AxC techniques for an application is challenging and many publications proposed how to explore the design space to find the most suitable AxC techniques for an application [2]. Among various approaches, such complex Design Space Exploration (DSE), including genetic algorithms and simulated annealing [3], machine learning (ML) ones such as Reinforcement Learning (RL) showed promising results in improving the DSE [4], minimizing the number of designs to evaluate while maximizing the quality of the DSE model and reducing the exploration time [5].

RL is a machine learning paradigm in which an agent learns a desired behavior by interacting with a dynamic environment. A standard RL setup involves an agent connected to its environment via observations and actions. During each interaction

This work has received funding from the APROPOS project in the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 956090 and from the project “National Center for HPC, Big Data and Quantum Computing”, CN00000013 (Bando M42C – Investimento 1.4 – Avviso Centri Nazionali” – D.D. n. 3138 of 16.12.2021, funded with MUR Decree n. 1031 of 17.06.2022).

with the environment, the agent receives observations about the current state of the environment and a reward based on its previous action. Based on this, it must select a new action to move toward an optimization goal, learning how to do this over time by systematic trial and error [6].

Wu et al. proposed a DSE framework using RL to optimize resource allocation and critical path timing [4]. The framework extracts the data flow graphs from the HLS C/C++, then the RL-based DSE engine explores the resource allocation options and finds optimized or Pareto solutions. The results show that their proposed RL-based engine outperforms genetic algorithms and simulated annealing. However, the exploration is limited to the trade-offs of different resources and critical path timing on an FPGA for a specific approximated version of an application, not considering the different approximated versions. A similar ML-based approach to DSE was proposed in [5] to reduce the exploration time of the HLS tools for a certain application.

In [7], the authors propose a new approach for the DSE of approximate applications to minimize the error rate using the data lifetime to select the approximation. Though their experimental results show their approach’s effectiveness, the paper does not cover multi-objective optimization, i.e., simultaneously considering accuracy, power consumption, and computation time.

This paper proposes an RL-based multi-objective DSE methodology based on the resource selection process proposed by [7] able to optimize an application balancing the final accuracy with the power consumption and the computation time introduced by the selective activation of approximate adders [8] and multipliers [9].

II. METHODOLOGY

This paper considers a CPU running software with dedicated instructions to trigger different approximate adders and multipliers. To generate approximate versions of a target application, the strategy is to select variables from the target application and approximate all sums or multiplications on those variables (as in [7]). Since several design choices exist (i.e., different sets of variables and different approximate adders and multipliers), we exploit an RL agent to explore the design space automatically. The exploration aims to find the

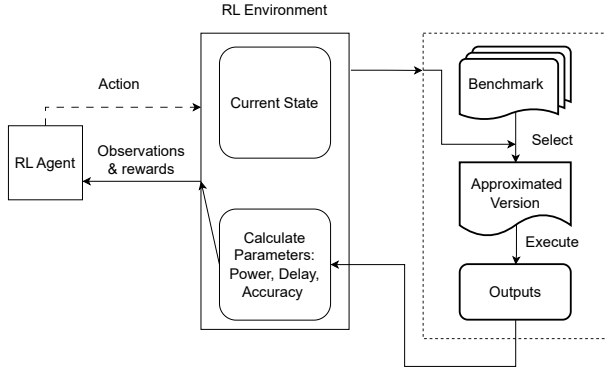


Fig. 1. RL environment: at each step selects an approximated version of the benchmark to execute and calculates the parameters to return the new observations and rewards to the RL agent.

best trade-off, or one of the best possible trade-offs, between accuracy degradation and power/computation time.

This process is depicted in Figure 1. The environment uses available approximate operators and a set of variables to create an approximate application according to agent instructions. The approximate version’s accuracy, power, and computation time are evaluated based on pre-characterized approximate operators. The agent receives observations and rewards from the environment, enabling them to take the next action to change the environment state. The RL setup mentioned above comprises four components: the environment, the agent, the state, and the action.

The *environment* formally defined in Equation 1 is the system in which the agent operates. It models all the aspect of the optimization. First, the list of variables is modeled as a boolean vector where variables are indexed and selected if the corresponding index is 1 ($variables_{approx} = \{a_0, a_1, \dots, a_{N-1} | a_i \in \{0, 1\}\}$). The approximate adders ($adder \in \{1, 2, \dots, N_{add}\}$) and the approximate multipliers ($multiplier \in \{1, 2, \dots, N_{mul}\}$) are indexed to be linked to the set of available ones. Both sets are sorted by increasing accuracy degradation. Also the outcome of the computation is modeled as the decrease in accuracy, through an evaluation function comparing the output of the precise version with the approximate one ($\Delta acc = f_{err}(output_{precise}, output_{approx})$). In the actual implementation, f_{err} evaluates the Mean Absolute Error (MAE) as the accuracy metric. MAE is calculated in Equation 2, where N is the number of the benchmark outputs. $exactOutput_i$ and $approxOutput_i$ are the exact and approximated values of each $output_i$. Eventually, the power reduction, expressed as the difference between the power consumption estimated for the precise execution and the approximated one ($\Delta power = power_{precise} - power_{approx}$) and the computational time reduction, expressed as the difference between the time used by precise execution in sums and multiplications and the one from the approximated version ($\Delta time = time_{precise} - time_{approx}$) are included.

$$environment = \{adder, multiplier, variables_{approx}, \Delta acc, \Delta power, \Delta time\} \quad (1)$$

The *state* represents the environment at a particular time (i.e., the list of approximated variables and the selected approximate adder and multiplier). This information is used to deploy the execution of an approximate application obtained through automatic code instrumentation that enables collecting observations (i.e., Δacc , $\Delta power$, and $\Delta time$). For example, the current state at a random step might look like $(4, 5, (1, 0, \dots, 0), 30, 100, 200)$, which, according to Equation 1, states that the configuration includes approximate adder number 4 and multiplier number 5, and only the first variable in the list is selected for approximation. The last three elements represent observations meaning that, i.e., accuracy was reduced by 30 units, power consumption by 100 units (e.g., mW), and computation time was reduced by 200 units (e.g., ns).

$$MAE = \frac{1}{N} \left(\sum_{i=0}^{N-1} exactOutput_i - approxOutput_i \right) \quad (2)$$

The *agent* is the learning algorithm that interacts with the environment. It observes the current state of the environment and selects an action to take based on that state. The agent’s goal is to learn a policy, which is a mapping from states to actions, that maximizes the expected cumulative reward.

The *action* is the decision made by the agent at a particular time step to apply to the environment and change its state. In our work, three possible actions are: changing the type of the adder, changing the type of the multiplier, or removing/adding one variable in the list of approximated variables.

In RL, the *reward* helps the agent decide which action to take and learn the optimization policy. Our approach defines the reward as shown in algorithm 1. The reward obtained in each step is accumulated with previous rewards. If the total reward reaches a maximum predefined one, the agent stops.

Algorithm 1 analyzes the current state. A further approximation can be introduced if the accuracy loss is below the tolerable accuracy loss threshold for the benchmark (line 4). This acceptable threshold is an exploration parameter and can be adapted to the case. If the most approximated adder and multiplier are in use and all variables are approximated maximum reward possible is given, and the “terminate” flag is set to true (lines 5 to 8). Otherwise (lines 9 to 16), the gain in power consumption and computation time must be above a threshold to obtain a positive reward (+1). Otherwise, the agent gets a negative reward (-1). Line 19 gives the maximum negative reward if the accuracy loss is above its threshold. Finally, the cumulative reward is updated (line 21).

We used the Q-learning as a learning algorithm along with the RL environment. It is a model-free value-based RL algorithm that lets the agent learn the value of an action in each state [10]. Q-Learning finds the optimal action selection policy

```

1 Input: R //Maximum cumulative reward ;
2 Input: Rcum // Cumulative reward ;
3 Input: state := (adder, multiplier, variables, Δacc, Δpower, Δtime);
4 if Δacc ≤ accth then
5   if (stateadder == Nadd and statemul == Nmul and variables
6     contains all ones then
7     reward = R ;
8     terminate = True;
9   end
10  else
11    if (Δpower ≥ pth) and (Δtime ≥ tth) then
12      reward = 1;
13    end
14    else
15      reward = -1;
16    end
17  end
18 else
19   reward = -R;
20 end
21 Rcum += reward;

```

Algorithm 1: RL Rewards at step i

TABLE I
SELECTED ADDERS FROM EVOAPPROXLIB

operator	Type	MRED	Power (mW)	Computation time (ns)
8-bit adder	1HG	0	0.033	0.63
8-bit adder	6PT	0.14	0.029	0.55
8-bit adder	6R6	2.93	0.012	0.27
8-bit adder	0TP	6.16	0.0095	0.24
8-bit adder	00M	14.58	0.0046	0.17
8-bit adder	02Y	24.87	0.0015	0.11
16-bit adder	1A5	0	0.072	1.28
16-bit adder	0GN	0.005	0.057	1.04
16-bit adder	0BC	0.018	0.051	0.95
16-bit adder	0HE	0.16	0.036	0.68
16-bit adder	0SL	9.54	0.011	0.27
16-bit adder	067	22.35	0.0041	0.20

using a Q function that returns the expected future reward of an action at a state. In Q-learning, an accumulated reward can be defined, and the action selection is based on the state of the previous step beside the reward. Hence, it is suitable for maximizing the accumulated reward while considering the last state. For further details on how Q-function is updated, the reader may refer to [10].

III. EXPERIMENTAL RESULTS

The RL implementation relies on the *Gymnasium* [11] library, a fork of the *OpenAI Gym* [12] library, to implement the RL engine. *Gymnasium* is an open-source toolkit for developing and comparing RL algorithms. It provides an API to define environments, or simulated tasks, where agents can learn and improve their decision-making abilities through trial and error. We implemented a Q-learning algorithm in which the maximum number of steps is 10,000. This number is selected upon trial and error. As an approximate component database, we exploited the *EvoApproxLib* [8], [9] that provides C models of 8 and 16 bits adders, and 8 and 32 bits multipliers. Tables I and II report all selected operators and their Mean Relative Error Distance (MRED), power consumption, and computation time, ordered by MRED.

TABLE II
SELECTED MULTIPLIERS FROM EVOAPPROXLIB

operator	Type	MRED	Power (mW)	Computation time (ns)
8-bit multiplier	1JQ	0	0.391	1.43
8-bit multiplier	4X5	0.033	0.380	1.40
8-bit multiplier	GTR	1.23	0.303	1.46
8-bit multiplier	L93	4.52	0.178	1.11
8-bit multiplier	18UH	17.98	0.062	0.90
8-bit multiplier	17MJ	53.17	0.0041	0.11
32-bit multiplier	precise	0	10.76	4.565
32-bit multiplier	000	0.00	10.46	4.470
32-bit multiplier	018	0.01	4.32	3.220
32-bit multiplier	043	1.45	1.63	2.440
32-bit multiplier	053	10.59	1.05	2.030
32-bit multiplier	067	41.25	0.51	1.750

We tested our approach with two applications: Matrix Multiplication with two different matrices sizes (10 by 10 and 50 by 50) and FIR (with 100 and 200 samples, all white noise signals with Low Pass Filter functionality). The thresholds were set after executing the precise versions: the power (p_{th}) and computation time (t_{th}) thresholds were set to 50% of their value for the precise version. Also, the precise outputs were averaged, and the accuracy threshold (acc_{th}) was set as 0.4 times the average output. Based on the rules defined in algorithm 1, the agent stopped after 2,000 and 4,000 steps for the Matrix Multiplication benchmarks and after 500 and 1,240 for the FIR benchmarks. All the results are reported in Table III.

TABLE III
EXPLORATIONS RESULTS FOR POWER, COMPUTATION TIME, AND ACCURACY

Benchmarks	Matrix Mult.		FIR	
	10x10	50x50	100	200
Δ Power Consumption (mW)				
min	15	0.55	529.515	1059.345
solution	415.3	753.72	10850.855	1237.247
max	418.4	1552.017	17344.390	34699.1
Δ Computation time (ns)				
min	50	-90	563.135	1126.605
solution	1780	1460.8	2664.385	3951.525
max	1840	5707.6	6547.495	13098.89
Accuracy degradation				
min	0.02	0	1096.03	395.74
solution	19.95	0.736	1096.03	27580.345
max	204.71	26.7964	31671.43	27580.35
Configuration				
Adder Type	00M	6R6	0GN	067
Multiplier Type	17MJ	L93	043	018

In Table III, power consumption and computation time are expressed as the difference (Δ) between the parameter obtained from the precise version and the one obtained from the approximation run of the last step. The min and max rows correspond to the minimum, and maximum Δ found in the exploration. The accuracy degradation shows how much computation accuracy is reduced for the selected approximate version, using Equation 2. The Adder and Multiplier types show which operators were selected for this approximate

version, as named in Tables I and II.

For all cases, the agent identified approximations respecting the given constraints confirming the capability of exploring the space. Comparing the selected solution and the maximum observed values for the different parameters can indicate how the agent could push the approximation without violating the imposed constraints.

To better analyze the results, Figure 2 and Figure 3 show the power and computation time alongside accuracy for all the exploration steps for the Matrix multiplication, 10 by 10, and FIR with 100 samples, respectively. The "Power" at each step represents the difference between the approximate and precise versions' power consumption. The same applies to the "Comp. Time" that shows the computation time. The trend lines help better visualize the exploration results. The different trends in the two figures highlight how the algorithm learns properly and move toward an optimization in the Matrix Multiplication benchmark, while it struggles in the FIR one.

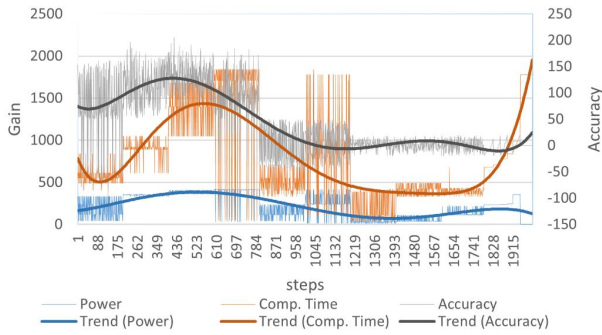


Fig. 2. Exploration outcomes evolution through for Matrix Multiplication (10x10).

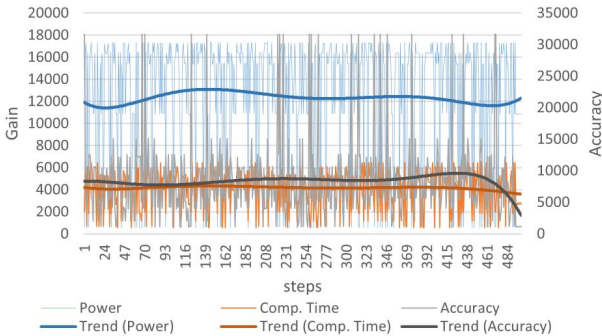


Fig. 3. Exploration outcomes evolution through for FIR (100 samples).

Figure 4 shows the average reward every 100 steps for the same Matrix multiplication and FIR benchmark configurations. For Matrix multiplication, on average, the reward continuously improves in the first 600 steps, slowing down for a while and improving again, confirming that the agent learned. On the contrary, the FIR does not follow such a continuous improvement, ensuring that the learning strategy is not entirely effective and that further investigations are required.

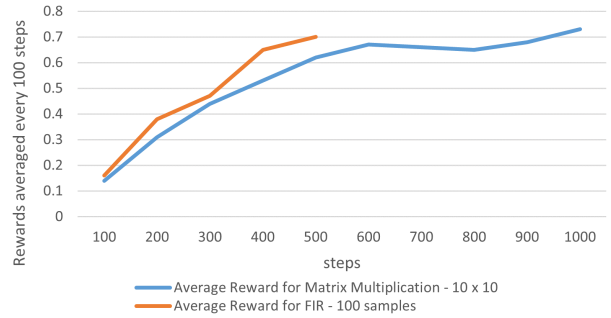


Fig. 4. Average reward evolution for the Matrix multiplication (10x10) and FIR (100 samples).

IV. CONCLUSION

In this paper, we proposed preliminary results obtained while developing a DSE strategy based on Reinforcement Learning to optimize an application, balancing the final accuracy with the power consumption and the computation time. Experimental results show that the agent can find solutions for a good trade-off between accuracy degradation and power and computation time reductions on some benchmarks. However, additional work is required to improve the learning strategy and make it general enough to work with a larger set of applications.

REFERENCES

- [1] A. Bosio, D. Ménard, and O. Sentieys, Eds., *Approximate Computing Techniques*. Springer International Publishing, 2022. [Online]. Available: <https://doi.org/10.1007/978-3-030-94705-7>
- [2] W. Hu *et al.*, "Exploring the design space of approximate arithmetic circuits using reinforcement learning," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 442–447.
- [3] M. Rizakis *et al.*, "Approximate fpga-based lstms under computation time constraints," in *Applied Reconfigurable Computing, Architectures, Tools, and Applications*, N. Voros *et al.*, Eds. Cham: Springer International Publishing, 2018, pp. 3–15.
- [4] Y. Wu *et al.*, "Ironman: Reinforcement learning based design space exploration for approximate computing," in *2021 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, 2021, pp. 1–8.
- [5] Q. Gautier *et al.*, "Sherlock: A multi-objective design space exploration framework," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 27, no. 4, pp. 1–20, 2022.
- [6] L. P. Kaelbling, M. L. Littman, and A. W. Moore, "Reinforcement learning: A survey," *Journal of artificial intelligence research*, vol. 4, pp. 237–285, 1996.
- [7] A. Savino *et al.*, "Approximate computing design exploration through data lifetime metrics," in *2019 IEEE European Test Symposium (ETS)*, 2019, pp. 1–7.
- [8] V. Mrazek *et al.*, "Evoapprox8b: Library of approximate adders and multipliers for circuit design and benchmarking of approximation methods," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*, 2017, pp. 258–261.
- [9] M. Česka *et al.*, "Approximating complex arithmetic circuits with formal error guarantees: 32-bit multipliers accomplished," in *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, 2017, pp. 416–423.
- [10] F. S. Melo, "Convergence of q-learning: A simple proof," *Institute Of Systems and Robotics, Tech. Rep.*, pp. 1–4, 2001.
- [11] S. V. Swaroop *et al.*, "Gymnasium," <https://github.com/s-vineet/gymnasium>, 2021, accessed: 2023-04-07.
- [12] G. Brockman *et al.*, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.