

FPGA-Based Neural Thrust Controller for UAVs

Sharif Azem, David Scheunert, Mengguang Li, Jonas Gehringer,
Kai Cui, Christian Hochberger and Heinz Koepl

Abstract—The advent of unmanned aerial vehicles (UAVs) has improved a variety of fields by providing a versatile, cost-effective and accessible platform for implementing state-of-the-art algorithms. To accomplish a broader range of tasks, there is a growing need for enhanced on-board computing to cope with increasing complexity and dynamic environmental conditions. Recent advances have seen the application of Deep Neural Networks (DNNs), particularly in combination with Reinforcement Learning (RL), to improve the adaptability and performance of UAVs, especially in unknown environments. However, the computational requirements of DNNs pose a challenge to the limited computing resources available on many UAVs. This work explores the use of Field Programmable Gate Arrays (FPGAs) as a viable solution to this challenge, offering flexibility, high performance, energy and time efficiency. We propose a novel hardware board equipped with an Artix-7 FPGA for a popular open-source micro-UAV platform. We successfully validate its functionality by implementing an RL-based low-level controller using real-world experiments.

I. INTRODUCTION

As a platform for a relatively compact and cost-effective system capable of agile aerial performance, the development of UAVs has contributed to diverse fields such as agriculture, surveillance, disaster management and logistics. With their ability to access remote or hard-to-reach areas, UAVs have emerged as a tool for data collection [1], monitoring [2], and rapid response missions [3]. The application of UAVs often involves complex algorithms, such as collision avoidance [4], communication [5], data processing [6], low level control [7] and indoor navigation [8]. The increasing complexity of tasks assigned to UAVs requires advancements in onboard computation and processing capabilities to ensure a reliable and effective operation. On top of that, UAVs often need to react to sudden changes in the environment they operate in, necessitating adaptive algorithms that are applicable in such scenarios.

Recently, DNNs have been used for the application of algorithms involving UAVs [9], [10]. DNNs are often used in combination with RL, where an agent learns the optimal behavior through an interaction with its environment. The integration of DNNs on UAV-systems in the context of RL can enhance the performance of UAVs, allowing them to react to previously unseen situations and adapt to changing

**This work has been funded by the Distr@1 "Digital Innovation and Technology Funding" of the Hessian Ministry for Digitalisation and Innovation (project no. 493 21.0052.2A).

The authors are with the Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, 64287 Darmstadt, Germany. {sharif.azem, mengguang.li, kai.cui, heinz.koepl}@bcs.tu-darmstadt.de, {scheunert, gehringer, hochberger}@rs.tu-darmstadt.de




Fig. 1. LF-deck with 4 IR receivers (marked in yellow), an Artix-7 FPGA XC7A15T (marked in light blue), one JTAG interface (marked in red) and an interface for connecting the deck with the Crazyflie (marked in white).

environments. However, Many UAVs have computers with limited computing power on board, making the implementation of DNNs infeasible.

FPGAs offer a promising solution to the computational demands of deploying DNNs on UAVs. FPGAs provide flexibility in the implementation with the high performance and energy efficiency of hardware. Their re-configurable nature, as opposed to Application-Specific Integrated Circuits (ASICs), allows for an optimized implementation of different architectures of DNNs. Furthermore, the implementation of an algorithm can be designed to consume less power, compared to an implementation of the same algorithm on Graphics Processing Units (GPUs), which provides high computation power but also consumes more energy.

In this work we present an implementation of DNNs on an expansion deck equipped with an FPGA for implementation of algorithms, in particular we implement DNNs with the deepsets architecture proposed in [10]. Furthermore, the expansion deck is based on the Lighthouse deck suggested in [11] to allow decentralized position estimation. In this work, we refer to the deck as the Lighthouse-FPGA (LF) deck. The ability to calculate the next action and the own position on board contributes to the autonomous operation of the UAV, reducing the dependency on an external system.

II. SYSTEM

The design of the LF expansion deck is divided into two parts with two different functionalities, both implemented on the FPGA of the deck. One functionality is responsible for receiving IR signals, and the other functionality is used for

hardware acceleration, and consists of the MicroBlaze soft-core processor¹ and a hardware accelerator to run onboard computation faster. Furthermore, to reduce the resource utilization of the resulting hardware implementation, fixed point calculations are used. This leads to a simplified hardware design of the neural network as well as a much lower latency of the calculated output.

A. Lighthouse-FPGA Expansion Deck

In this study, we employ the Crazyflie nano quadrotor [12], equipped with the STM32F405 micro-controller (MCU) together with the Lighthouse deck. For this, we attach the deck on top of the Crazyflie through the Crazyflie's 20 pin interface (Figure 1 white marks). This interface is used both for power supply (from the Crazyflie's battery) and communication between the FPGA and the MCU. The deck is equipped with four IR sensors (see Figure 1, orange marks) that can detect signals from Lighthouse base stations for on-board position estimation, and dedicated ICs to process these signals. The design and functionality are based on the Lighthouse deck proposed in [11]. However, the LF-deck is equipped with an AMD/Xilinx Artix 7 FPGA XC7A15T, as in Figure 1. The positioning module (in this work called also base function) is migrated to the LF-deck and communicates with the MCU through a Universal Asynchronous Receiver-Transmitter (UART) protocol. Furthermore, we use the MicroBlaze soft core processor alongside a hardware accelerator on the FPGA for additional calculations of computation intensive algorithms, such as the feed-forward calculation of DNNs. Utilizing the I2C protocol bidirectionally, the MCU initially transmits the necessary information to the MicroBlaze for computation. Then, the computation is executed and the result is sent back to the MCU. For programming the FPGA, we use a High-Level Synthesis (HLS) tool-chain that creates custom hardware accelerators, designed to identify parts of the implemented code to accelerate repetitive tasks, such as the execution of loops. The structure of the implementation within the FPGA is shown in Fig. 2. In order to flash the program onto the FPGA, The LF-deck is equipped with a JTAG interface (see Figure 1 red mark), which makes it possible to flash programs directly to the deck without using the Crazyflie.

B. Hardware Accelerator

The hardware accelerator is created through High-Level Synthesis (HLS). For this purpose, the tool *Plugin for Intermediate Representation ANalysis and Hardware Acceleration* (PIRANHA) [13] is used. It is a plugin for the GNU Compiler Collection (GCC), that operates on unannotated source code and requires no manual integration of the accelerator into the source code by the developer. The selection of optimization strategies and hardware architecture of the entire System-on-Chip (SoC) realized on the FPGA is created using the SpartanMC system builder [14].




Fig. 2. Block diagram of the FPGA implementation. The base function refers to the functionality provided by the LF-deck suggested in [11]. Additionally, a combination of a MicroBlaze softcore processor and a hardware accelerator is provided, where the hardware accelerator is created at compilation time using a tool called PIRANHA.




Fig. 3. Integration into GCC and workflow of PIRANHA. This process identifies segments of the code that can be accelerated using a hardware accelerator.

The working principle of PIRANHA is shown in Fig. 3. The firmware for the LF-deck is processed through PIRANHA by means of three separate compiler plugins. The first collects information about the loops in the firmware. Loops that would benefit from parallelization or hardware acceleration are automatically identified and extracted. In a second call, the `Accelerator Generation` stage, each loop is separately optimized and scheduled. If the resulting schedule is faster than a pure software implementation, PIRANHA generates the corresponding hardware accelerator and modifies the original firmware in the third and final call through the `Accelerator Integration` plugin. Here, the accelerated loop is removed from the firmware and replaced with a call to the previously generated hardware accelerator. This process is transparent to the programmer, any part of the software that is not suitable for acceleration will be executed on the MicroBlaze as expected. PIRANHA takes advantage of the various optimization strategies and static analysis methods that the GCC provides to create fast and efficient hardware accelerators. If no conclusion about independence of operations is possible based on static analysis, run-time conditions can be generated [15]. Likewise, bypasses and squashes are automatically applied to series of dependent memory accesses to further improve the accelerator throughput [16]. For the neural network evaluation, the feed-forward calculation is accelerated for both the vector-matrix multiplication as well as for the activation function.

C. Fixed-Point Arithmetic

The process of converting the weights of the DNN into integers involves using fixed-point arithmetic. This conversion is achieved by multiplying each weight by 2^n and then

¹<https://www.xilinx.com/products/design-tools/microblaze.html>

rounding down the result to get an integer value, where $n \in \mathbb{N}$ represents the amount of fractional bits. For DNNs with multiple Multilayer Perceptron (MLPs) we apply a uniform transformation for conversion, meaning the same n value is used for all the MLPs within the DNN. To determine the optimal number of fractional bits, we randomly generate input samples for the DNN and run them through the network to get the output as a floating point number. The inputs are also converted into their fixed-point counterparts using the chosen n , and then they are fed into the fixed-point version of the DNN to obtain another set of outputs, this time represented as integers, which are then converted to floating point numbers by being multiplied with 2^{-n} . By comparing the outputs from the floating-point and fixed-point version of the DNN, the maximum error for each n can be assessed. This process is repeated with different samples for each n value in ascending order, starting from $n = 1$. We choose the n that results in the smallest maximum difference between the floating-point and fixed-point representations.

III. MODEL

In the following we are going to discuss the dynamic model of the quadrotor and the model of the DNN used in this work. For training, we use the simulator proposed in [17], which includes a dynamic and physical model of the quadrotor according to [7], along with a RL interface. We use the same reward function and hyper parameters described in [10], and use the RL interface in [17] to modify the deepsets architecture proposed in [10] to simplify the implementation on the LF-deck.

A. Quadrotor Dynamics

The quadrotor is modeled as a rigid body, where $\mathbf{x}$ denotes the position of its center of mass in an inertial reference frame. The rotation matrix from the quadrotor body frame to the inertial frame is denoted as $\mathbf{R}$. The equation of motion can be written as follows using the Newton-Euler equations

$$\begin{aligned}\ddot{\mathbf{x}} &= \mathbf{g} + \frac{\mathbf{R}\mathbf{f}}{m} \\ \dot{\boldsymbol{\omega}} &= \mathbf{I}^{-1}(\boldsymbol{\tau} - \boldsymbol{\omega} \times (\mathbf{I} \cdot \boldsymbol{\omega})) \\ \dot{\mathbf{R}} &= \mathbf{R}\boldsymbol{\omega}_{\times},\end{aligned}\quad (1)$$

where $\mathbf{g}$ is the gravity vector, m is the mass of the quadrotor, $\mathbf{f}$ is the total thrust vector in body frame, $\boldsymbol{\omega}$ is the angular velocity in the body frame and $\mathbf{I}$ is the inertia matrix. $\boldsymbol{\omega}_{\times}$ is the skew symmetric matrix of $\boldsymbol{\omega}$.

B. Reinforcement Learning

In the following t denotes the time step and q the observing quadrotor. The self observation vector is defined as $\mathbf{o}_t^q = (\mathbf{p}_t^{q,j}, \mathbf{v}_t^q, \mathbf{r}_t^q, \boldsymbol{\omega}_t^q)$, where $\mathbf{p}_t^{q,j} \in \mathbb{R}^3$ is the position of the quadrotor relative to its own target position, i.e. $\mathbf{p}_t^{q,j} = \mathbf{p}_t^q - \hat{\mathbf{p}}_t^j$, where $\mathbf{p}_t^q$ and $\hat{\mathbf{p}}_t^j$ are its own position and target position respectively. $\mathbf{v}_t^q \in \mathbb{R}^3$ is the velocity of the quadrotor in the global world frame, $\mathbf{r}_t^q \in \mathbb{R}^9$ is the rotation vector which we obtain by row-wise flattening the rotation matrix $\mathbf{R}_t^q \in \mathbb{R}^{3 \times 3}$ (from the quadrotor body




Fig. 4. The neural network structure implemented in the work. The network consists of the self encoder E^q and the neighbor encoder E^k . The neighbor encoder consists of the MLP B and a mean operator for calculating the mean vector e^k out of the output vectors of B . The output of both encoders are concatenated, where $[e^q, e^k]$ denotes the concatenation, and processed in the output network H .

frame to the world frame), and $\boldsymbol{\omega}_t^q$ is the angular velocity of the quadrotor in body frame. The neighbor observations are $\delta_t^1, \dots, \delta_t^l, \dots, \delta_t^k$ with k neighbors, where the l th vector is related to the l th nearest neighbor, and is denoted as $\delta_t^l = (\mathbf{p}_t^{q,l}, \mathbf{v}_t^{q,l}) \in \mathbb{R}^6$, where $\mathbf{p}_t^{q,l}$ and $\mathbf{v}_t^{q,l}$ are the position and velocity of the q th quadrotor relative to the k th neighbor respectively, defined as $\mathbf{p}_t^{q,l} = \mathbf{p}_t^q - \mathbf{p}_t^l$ and $\mathbf{v}_t^{q,l} = \mathbf{v}_t^q - \mathbf{v}_t^l$, where $\mathbf{p}_t^l$ and $\mathbf{v}_t^l$ are the global position and velocity of the l th neighbor respectively. The output of the neural network is defined as $\mathbf{a} \in [-1, 1]^4$, and we obtain the thrust vector by applying an affine transformation $\hat{\mathbf{f}} = \frac{1}{2}(\text{clip}(\mathbf{a}, -1, 1) + 1)$.

The DNN model consists of the self encoder E^q , which is an MLP, that processes the self observation, and the neighbor encoder E^k that processes the neighbor observations. The DNN model we use is permutation and scale invariant. For this, each neighbor vector is inserted to the MLP B , and the mean vector is calculated out of the output vectors that are obtained from B , to obtain the output of the neighbor encoder e^k . Finally the output of the self encoder e^q and the neighbor encoder e^k are concatenated and inserted to the MLP H , to obtain the output of the neural network, denoted as $\mathbf{a}$. This structure is depicted in Figure 4. The self encoder consists of two hidden layers with 16 neurons each and ReLU activation functions after each hidden layer. The neighbor encoder has two hidden layers and 8 neurons and ReLU activation functions after each hidden layer. The MLP H has one hidden layer with 32 neurons, and a ReLU activation function, where the output layer is without an activation function.

IV. EXPERIMENTS AND RESULTS

To validate our implementation we show flight trajectories of the Crazyflie that is equipped with the LF-deck. We conduct 3 experiments in a flight area of $6.5m \times 4.5m \times 2.7m$, where the results are shown in the figure 5.

a) *4 Setpoints in different directions*: In this scenario, the Crazyflie starts at the position 1, and the setpoints 2, 3 and 4 are arbitrarily placed in different directions in order to show the ability of the learned policy to react to sudden changes in the direction of the target position. As we can see, the Crazyflie is able to arrive at the setpoints while maintaining a small distance to the desired target positions.

The overall performance shows that Crazyflie is able to fly towards target positions at different directions with a small deviation from the desired trajectory path, which means that our implementation of the DNN on the FPGA was successful.

b) *Rectangular flight trajectory*: The Crazyflie flies along a rectangular path. It starts at point 1 and lands at the end at point 5. We can observe a deviation from the desired path, in particular between the setpoints 2 to 3 and 3 to 4. However, we also see that the Crazyflie is able to fly to its target position, demonstrating the functionality of the low level control policy of the neural network.

c) *Spiral flight trajectory*: The Crazyflie is placed at the center of the room, and performs a spiral flight trajectory. The Crazyflie flies to the desired target setpoints, where the last setpoint lies above the starting position at $z = 0.8\text{m}$. Here we see a notable deviation from the desired trajectory, although the Crazyflie was able to fly according to a spiral trajectory. This performance shows the successful implementation of the DNN on the LF-deck.




Fig. 5. Flight trajectories of the Crazyflie, where the gray dashed line and blue line are the desired flight trajectory and the real flight trajectory respectively. a) 4 setpoints in different directions. b) Rectangular flight trajectory tracking. c) Spiral flight trajectory tracking.

V. CONCLUSIONS AND FUTURE WORK

In this work we have presented the successful implementation of DNN on the LF expansion deck. The trained neural network is capable of low-level motor control while avoiding collisions among the quadrotors, where in this work we demonstrated only low level control on a real quadrotor. The implementation of computation expensive task on the LF-deck can be utilized to accelerate the computation and by this also to consume less energy. We validate our implementation by showing successful flight trajectories in different flight scenarios. To the best of our knowledge, this is the first implementation of the deepsets architecture on an FPGA extension deck for the Crazyflie quadrotors, using a cost-efficient positioning system that allows on-board position estimation. In the future, we plan to implement larger DNNs on the LF-deck as well as other network structures such as the attention model [10]. A detailed comparison between the LF-deck and the STM32F405 in terms of power consumption and execution time will be conducted. In addition, more LF-decks will be manufactured to demonstrate swarming behavior while maintaining collision-free flight among a fleet of quadrotors.

REFERENCES

- [1] S. Liu, Z. Wei, Z. Guo, X. Yuan, and Z. Feng, "Performance analysis of UAVs assisted data collection in wireless sensor network," in *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, 2018, pp. 1–5.
- [2] F. R. D. Giordan, A. Manconi and F. Nex, "Use of unmanned aerial vehicles in monitoring application and management of natural hazards," *Geomatics, Natural Hazards and Risk*, vol. 8, no. 1, pp. 1–4, 2017.
- [3] H. Hildmann and E. Kovacs, "Review: Using unmanned aerial vehicles (UAVs) as mobile sensing platforms (MSPs) for disaster response, civil security and public safety," *Drones*, vol. 3, no. 3, 2019.
- [4] J. N. Yasin, S. A. S. Mohamed, M.-H. Hagbayan, J. Heikkonen, H. Tenhunen, and J. Plosila, "Unmanned aerial vehicles (UAVs): Collision avoidance systems and approaches," *IEEE Access*, vol. 8, pp. 105 139–105 155, 2020.
- [5] S. Azem, A. Tahir, and H. Koepl, "Dynamic time slot allocation algorithm for quadcopter swarms," in *2022 IEEE 19th Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2022, pp. 1–6.
- [6] M. W. Mueller, M. Hamer, and R. D'Andrea, "Fusing ultra-wideband range measurements with accelerometers and rate gyroscopes for quadcopter state estimation," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1730–1736.
- [7] A. Molchanov, T. Chen, W. Hönig, J. A. Preiss, N. Ayanian, and G. S. Sukhatme, "Sim-to-(multi)-real: Transfer of low-level robust control policies to multiple quadrotors," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 59–66.
- [8] K. Kefferpütz and K. McGuire, "Error-state unscented kalman-filter for UAV indoor navigation," in *2022 25th International Conference on Information Fusion (FUSION)*, 2022, pp. 01–08.
- [9] R. Ourari, K. Cui, A. Elshamhory, and H. Koepl, "Nearest-neighbor-based collision avoidance for quadrotors via reinforcement learning," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 293–300.
- [10] S. Batra, Z. Huang, A. Petrenko, T. Kumar, A. Molchanov, and G. S. Sukhatme, "Decentralized control of quadrotor swarms with end-to-end deep reinforcement learning," in *Conference on Robot Learning*. PMLR, 2022, pp. 576–586.
- [11] A. Taffanel, B. Rousselot, J. Danielsson, K. McGuire, K. Richardsson, M. Eliasson, T. Antonsson, and W. Hönig, "Lighthouse positioning system: Dataset, accuracy, and precision for UAV research," *arXiv preprint arXiv:2104.11523*, 2021.
- [12] W. Giernacki, M. Skwierczyński, W. Witwicki, P. Wroński, and P. Kozierski, "Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering," in *2017 22nd International Conference on Methods and Models in Automation and Robotics (MMAR)*, 2017, pp. 37–42.
- [13] G. Hempel, C. Hochberger, and M. Raitza, "Towards GCC-based automatic soft-core customization," in *22nd International Conference on Field Programmable Logic and Applications (FPL)*, 2012, pp. 687–690.
- [14] G. Hempel and C. Hochberger, "A resource optimized SoC kit for FPGAs," in *2007 International Conference on Field Programmable Logic and Applications*, 2007, pp. 761–764.
- [15] F. Dewald, J. Rohde, C. Hochberger, and H. Mantel, "Improving loop parallelization by a combination of static and dynamic analyses in HLS," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 15, no. 3, feb 2022.
- [16] J. Rohde, K. Müller, and C. Hochberger, "Improving HLS generated accelerators through relaxed memory access scheduling," in *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2020, pp. 74–81.
- [17] Z. Huang, S. Batra, T. Chen, R. Krupani, T. Kumar, A. Molchanov, A. Petrenko, J. A. Preiss, Z. Yang, and G. S. Sukhatme, "Quadswarm: A modular multi-quadrotor simulator for deep reinforcement learning with direct thrust control," *arXiv preprint arXiv:2306.09537*, 2023.