# Ultra Low Power Deep-Learning-powered Autonomous Nano Drones

Daniele Palossi, Antonio Loquercio, Francesco Conti, Eric Flamand, Davide Scaramuzza, Luca Benini

*Abstract*—Flying in dynamic, urban, highly-populated environments represents an open problem in robotics. State-of-the-art (SoA) autonomous Unmanned Aerial Vehicles (UAVs) employ advanced computer vision techniques based on computationally expensive algorithms, such as Simultaneous Localization and Mapping (SLAM) or Convolutional Neural Networks (CNNs) to navigate in such environments. In the Internet-of-Things (IoT) era, nano-size UAVs capable of autonomous navigation would be extremely desirable as self-aware mobile IoT nodes. However, autonomous flight is considered unaffordable in the context of nano-scale UAVs, where the ultra-constrained power envelopes of tiny rotor-crafts limit the on-board computational capabilities to low-power microcontrollers. In this work, we present the first vertically integrated system for fully autonomous deep neural network-based navigation on nano-size UAVs. Our system is based on GAP8, a novel parallel ultra-low-power computing platform, and deployed on a $27\,\mathrm{g}$ commercial, open-source CrazyFlie 2.0 nano-quadrotor. We discuss a methodology and software mapping tools that enable the SoA CNN presented in [1] to be fully executed on-board within a strict $12\,\mathrm{fps}$ real-time constraint with no compromise in terms of flight results, while all processing is done with only $94\,\mathrm{mW}$ on average - 1% of the power envelope of the deployed nano-aircraft.

*Index Terms*—Unmanned Autonomous Vehicles, Convolutional Neural Networks, Ultra-low-power

## I. INTRODUCTION

UNMANNED Aerial Vehicles (UAVs) have recently reached an impressive level of autonomous navigation capability and accuracy, like high-precision detection and tracking, bio-inspired navigation, reactive obstacle avoidance, aggressive maneuvers, primarily thanks to computer-vision based approaches [1], [2], [3], [4], [5], [6], [7], [8], [9]. Vision-based UAVs are increasingly being used for practical applications such as the inspection of industrial facilities or cultivated fields [10], assistance in natural disaster or
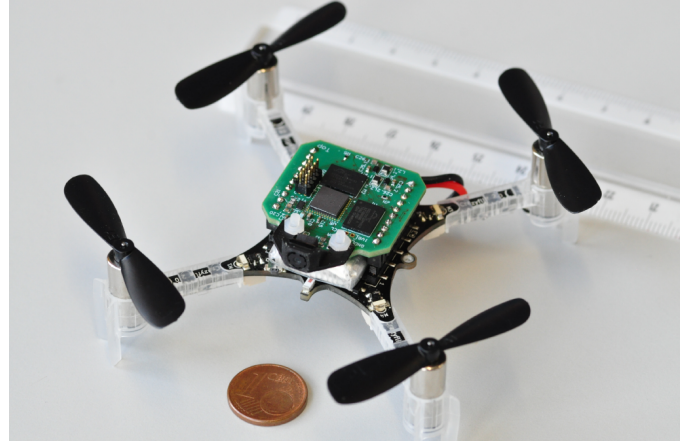
Fig. 1: Our prototype based on the COTS *Crazyflie 2.0* nano-quadrotor extended with our *PULP-Shield*. The system is able to run the *DroNet* [1] CNN for autonomous navigation up to $12\,\mathrm{fps}$ using only on-board resources.

hazardous areas [11], various surveillance and monitoring tasks [12], etc.

TABLE I: Rotorcraft UAVs taxonomy by vehicle class-size.

| Vehicle Class | ⊘ : Weight [cm:kg] | Power [W] | On-board Device |
|---|---|---|---|
| *std-size* [12] | $\sim 50 : \geq 1$ | $\geq 100$ | Desktop |
| *micro-size* [13] | $\sim 25 : \sim 0.5$ | $\sim 50$ | Embedded |
| *nano-size* [14] | $\sim 10 : \sim 0.01$ | $\sim 5$ | MCU |
| *pico-size* [15] | $\sim 2 : \sim 0.0001$ | $\sim 0.1$ | ULP |

To expand the class of activities that can be performed by UAVs, a recent trend of their evolution is their miniaturization. Commercial-Off-The-Shelf (COTS) quadrotors have already started to enter the nano-scale, featuring only few centimeters in diameter and few tens of grams in weight [16]. Commercial nano-UAVs still lack the vision-based autonomous navigation capabilities boasted by their larger counterparts, since their computational capabilities, heavily constrained by their tiny power envelopes, are totally inadequate for the execution of sophisticated workloads, as summarized in Table I. Furthermore, even the most advanced autonomous *nano-* and pico-size UAVs still perform vision-based computation on remote power-unconstrained machines, limiting the system's capability due to communication latency, bandwidth, noise, etc.

On the other hand, full autonomy of nano-scale UAVs is extremely desirable as it would make them the perfect

"smart sensors" in the Internet-of-Things (IoT) era [17]. These tiny autonomously flying robots could constitute self-aware smart IoT nodes: flying safely, they would be able to collect information both from the on-board sensors and from a plethora of devices deployed in the environment, and possibly even perform advanced on-board analytics, to pre-select essential information before transmitting to centralized servers [18]. The tiny form-factor of nano-drones is ideal both for indoor applications where they operate near humans (for surveillance, monitoring, ambient awareness, interaction with smart environments, etc.). It is also crucial for highly-populated urban areas, where autonomous robots can exploit complementary sense-act capabilities to interact with the surroundings (e.g., smart-building, smart-cities, etc.). Recent results have shown [1] that the advanced vision-based autonomous navigation capabilities required in these dynamic environments depend on computationally expensive algorithms such as Convolutional Neural Networks (CNNs), further raising upwards the request for computing performance within the allotted power budget.

To put this power budget into perspective, in Wood et al. [15] the authors estimate the power break-down for small-size UAVs; they show that the maximum power budget for on-board computation is 5% of the total, the rest being used by the propellers (86%) and the low-level control parts (9%). The problem of bringing state-of-the-art navigation capabilities on the challenging classes of nano- and pico-size UAVs is therefore strictly dependent on the development of energy-efficient computing architectures, highly optimized software and new classes of algorithms. Whereas standard-size UAVs with a power envelope of several hundreds Watts have always been able to host powerful high-end embedded computers like NVIDIA Jetson TX1 and TX2, Qualcomm Snapdragon[1], Odroid, etc., most nano-sized UAVs have been constrained by the capabilities of microcontroller devices capable to provide a few hundreds Mop/s at best. Therefore, CNN-based autonomous vision navigation was so far considered to be out of reach for this class of drones.

In this work, we present what, to the best of our knowledge, is the first deployment of a SoA, fully autonomous vision-based navigation system based on deep learning on top of a UAV compute node consuming less than $94\,\mathrm{mW}$ at peak, fully integrated within an open source COTS *CrazyFlie 2.0* UAV. Our compute node, shown in Figure 1, leverages the *GreenWaves GAP8* SoC, a high-efficiency embedded processor taking advantage of the emerging parallel ultra-low-power (PULP) computing paradigm to enable the execution of complex algorithmic flows onto power-constrained devices such as nano-scale UAVs.

The main contributions of this work can be summarized as follows:

- we adapted the CNN-based approach for autonomous navigation proposed in Loquercio et al. [1] for standard-sized UAVs to the computational requirements of a nano-sized UAV, such as fixed-point computation;

- we developed a pluggable PCB, the *PULP-shield*, to host the GAP8 SoC as well as an ultra-low-power camera directly on the $27\,\mathrm{g}$ *CrazyFlie* nano-UAV;
- we developed a methodology and tools for deploying SoA CNNs on PULP. We demonstrate this methodology for the *DroNet* CNN, achieving comparable quality-of-results in terms of UAV control with respect to the standard-sized baseline of [1] within an overall power budget of just $94\,\mathrm{mW}$

Our work demonstrates that parallel ultra-low-power computing is a viable solution to deploy autonomous navigation capabilities onboard a nano-UAVs used as smart, mobile IoT endnodes, while at the same time showcasing a complete hardware/software methodology to implement such complex workloads on heavily power- and memory-constrained device.

## II. RELATED WORK

The development of the IoT is fueling a trend toward edge computing, to improve scalability, robustness, security [17]. While today's IoT edge nodes are usually stationary, autonomous nano-UAVs can be seen as perfect examples of next-generation IoT end-nodes, with high mobility and requiring an unprecedented level of on-board intelligence.

Related work can be organized into two different categories: on one side, works focusing on the development of advanced perception algorithms running on standard-sized, power-hungry drones [1], and, on the other side, works aiming at miniaturizing the size of the vehicle while relying on off-board computation [15].

*Advanced Algorithms:* On the standard-size class of vehicles a wide variety of advanced techniques for drone navigation and obstacle avoidance is available in the literature. At high level, these methods differ for the kind of sensory input and processing employed to control the flying platform. The classical approaches are based on mapping, localization and planning [19], [20]. To infer the system state, usually the platform is provided with GPS, range and/or visual sensors [1], [2], [3], [4], [5], [6], [7], [8], [9]. However, the impossibility to use GPS in indoor environments, and the inherent difficulties of state estimation [21], make those traditional approaches prone to fail under real world conditions.

Recently, there has been an increasing research effort in directly learning control policies for UAV from raw sensory data using imitation learning. Given its relatively low sample complexity (i.e. not much sample data is required to generalize) and its implementation simplicity, supervised learning has become the predominant tool used to learn visual-motor policies [22], [23], [1], [24], [25]. The supervisory signal may come from a human expert [22], hard-coded trajectories [24], or model predictive control [25]. However, for most application scenarios, it can be both tedious and dangerous to collect a large set of expert trajectories. A possible approach is to collect data in simulation and then transfer the learned policy to the real world. To generate very basic navigation policies, however, either a lot of photo-realistic data [26], or some real world examples [27] are required. Therefore, in order to safely and efficiently acquire data, the authors

---

[1] https://developer.qualcomm.com/hardware/qualcomm-flight

of [1] proposed to use cameras mounted on cars and bicycles. Tightly coupling perception and control, the resulting visual motor policy unlocks good generalization performance on a set of environments unseen during training. Similar works in [23] trained a deep neural network from video collected by a mountain hiker to detect forest trails.

Clearly, those learning-based approaches are limited by the constraints imposed by the on-board computational resources available on the UAV. Indeed, either off-line computations [24], [26], [27] or simplification of the motion model [1], [23] are required to overcome the domain-shift between the expert providing supervision and the learning agent.

*Nano-size Vehicles:* Particularly relevant in this work is the nano-size class of vehicles, where the SoA is represented by solutions that either offload computation to some remote powerful base-station or perform on-board basic functionalities. The authors of [28] developed a visual-inertial pose estimation system for a $45\,\text{g}$ quadrotor. A modified simultaneous localization and and mapping (SLAM) algorithm was also implemented to assist the controller in trajectory tracking. The SLAM algorithm ran on a remote laptop connected via WiFi to the robot. In [29] a $25\,\text{g}$ nano-UAV with visual-inertial SLAM for stabilization was presented. Here all the computation was performed off-board, streaming video and inertial information to a power-unconstrained base-station. Main problems with this kind of solutions are latency, maximum distance communication, channels noise and high on-board power-consumption, due to the high-frequency video streaming.

Nano-size flying robots showing some autonomous navigation capability based only on the on-board computational resources have been presented. In [14] the authors developed a $4\,\text{g}$ stereo-camera and proposed a velocity estimation algorithm able to run on the MCU on-board a $40\,\text{g}$ flying robot. If on one side this solution allows the drone to avoid obstacles during the flight, it still requires favorable flight condition (e.g., low flight speed of $0.3\,\text{m/s}$). In [30] an optical-flow-based guidance system was developed for a $46\,\text{g}$ nano-size UAVs. The proposed ego-motion estimation algorithm did not rely on feature tracking, making it possible to run on the on-board MCU. Although, the target application was limited to hovering and the method did not reach the accuracy of computational-expensive techniques based on feature tracking.

COTS nano-size quadrotors, like the *Bitcraze Crazyflie 2.0* or the *Walkera QR LadyBug*, typically embed on-board low-power single core MCUs, like the *ST Microelectronics* STM32F4 [14], [29], [31]. Despite the undeniable research effort, we still need to bring the most advanced capabilities from powerful big airborne to the resource-constrained "pocket drones".

Pushing beyond the aforementioned approaches, in this work we propose and demonstrate a system capable of sophisticated workloads such as real-time learning-based autonomous navigation [1], entirely on-board within the limited power envelope of nano-scale UAVs ($\sim 0.2\,\text{W}$) – whereas approaches of this complexity have been previously limited to standard-sized UAVs with external server-based computation or power hungry processors ($\geq 10\,\text{W}$).

## III. BACKGROUND

### A. DroNet

DroNet is a lightweight residual convolutional neural network (CNN) architecture. By predicting the steering angle and the collision probability, it enables the safe autonomous flight of a quadrotor in a various indoor and outdoor environments.

The architecture, as illustrated in Figure 2, was inspired by residual networks [32] and was reduced in size to minimize the bare image processing time. The two tasks of steering and collision probability prediction share all the residual layers in order to reduce the network complexity and the frame processing time. Then, two separate fully connected layers independently infer steering and collision probabilities. Mean-squared error (MSE) and binary cross-entropy (BCE) have been used to train the two predictions, respectively. A temporal dependent weighting of the two losses ensures the training convergence despite the different gradients' magnitude produced by each loss. Eventually, to make the optimization focus on the samples that are most difficult to learn, hard negative mining was deployed in the final stages of learning. The two tasks have separate datasets to learn from. Steering angle prediction is learned through the Udacity dataset[2], while the collision probability was learned through the Zürich bicycle dataset [1].

The outputs of DroNet are used to command the UAV to move on a plane with forward velocity $v_k$ and steering angle $\theta_k$. More specifically, the low-pass filtered probability of collision is used to modulate the UAV forward velocity, while the low-pass filtered steering angle is converted to drone's yaw control.

The result is a single shallow network that processes all visual information concurrently and directly produces control commands for a flying drone. The coupling between perception and control, learned end-to-end, provides several advantages, such as a simple, lightweight system and high generalization abilities. Indeed, the method was shown to function not only in urban environments but also on set of new application spaces without any initial knowledge about them. Indeed, with neither a map of the environment nor retraining or fine-tuning, the approach generalizes to scenarios completely unseen at training time including indoor corridors, parking lots, and high altitudes.

### B. GAP8 Architecture

Our deployment target for the bulk of the DroNet computation is GAP8, a commercial embedded RISC-V processor derived from the PULP open source project[3]. At its heart, GAP8 is composed by an advanced RISC-V microcontroller unit coupled with a programmable eight-core accelerator for digital signal processing and embedded deep inference.

Figure 3 shows the architecture of GAP8 in detail. The processor is composed by two separate power and clock domains, the SoC and the CLUSTER. The SoC is an advanced microcontroller unit centered on a single core (*fabric*
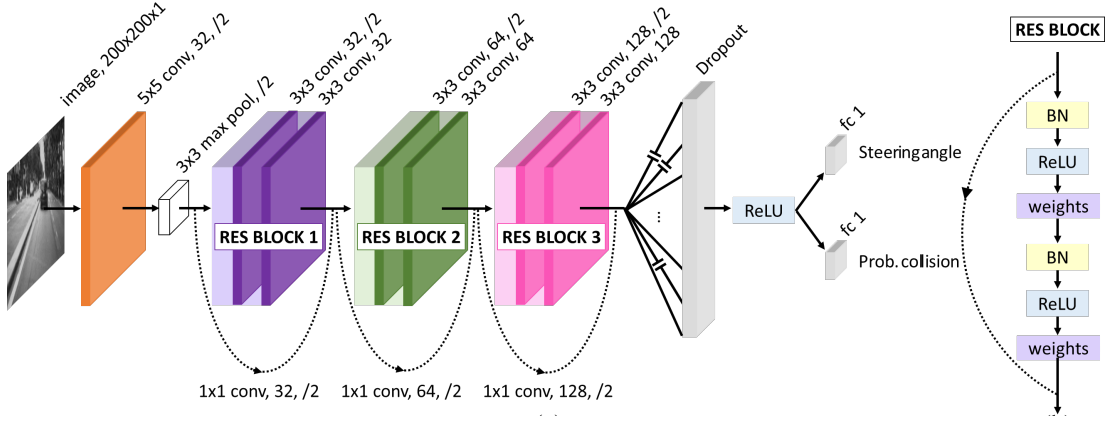
---

[2]https://www.udacity.com/self-driving-car

[3]http://pulp-platform.org
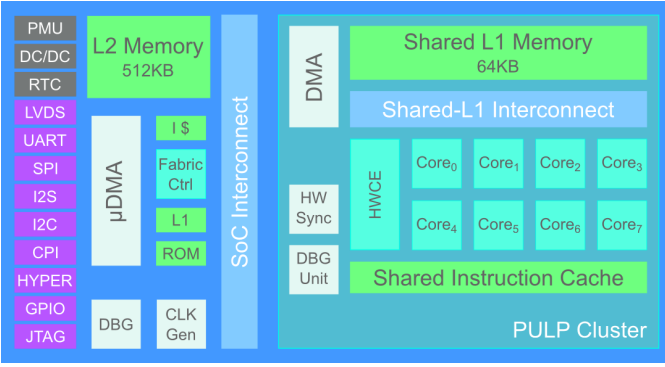
Fig. 2: *DroNet* original architecture.



Fig. 3: Architecture of the GAP8 embedded processor.

*controller*) coupled with $512\,\mathrm{kB}$ of SRAM (*L2 memory*). The fabric controller uses an in-order, DSP-extended four-stage microarchitecture implementing the RISC-V instruction set architecture [33]. The core supports the RV32IMC instruction set consisting of the standard ALU instructions plus the multiply instruction, with the possibility to execute compressed code. In addition to this, the core is extended to include a register-register multiply-accumulate instruction, vectorized DSP instructions (e.g. fixed-point dot product), bit manipulation instructions and two hardware loops. Moreover, the SoC features an autonomous multi-channel I/O DMA controller ($\mu DMA$) [34] capable of transfering data between a rich set of peripherals (QSPI, I2S, I2C, HyperBus, Camera Parallel Interface) and the L2 memory with no involvement of the fabric controller. The HyperBus and QSPI interfaces can be used to connect GAP8 with an external DRAM or Flash memory, effectively extending the memory hierarchy with an external L3 with a bandwidth of $333\,\mathrm{MB/s}$ and capacity up to $128\,\mathrm{Mbit}$.

The CLUSTER is dedicated to the acceleration of computationally intensive tasks. It contains eight RISC-V cores (identical to the one used in the fabric controller) sharing a $64\,\mathrm{kB}$ multi-banked *shared L1 scratchpad memory* through a low-latency, high-throughput logarithmic interconnect [35]. The shared L1 memory supports single-cycle concurrent access from different cores requesting memory locations on separate banks and a starvation-free protocol in case of bank

contentions (typically $<10\%$ on intensive kernels). The eight cores are fed with instruction streams from a single shared, multi-ported cache to maximize the energy efficiency on data-parallel code. A cluster *DMA* controller is used to transfer data between the shared L1 scratchpad and the L2 memory; it is capable of 1D and 2D bulk memory transfer on the L2 side (only 1D on the L1 side). A dedicated *hardware synchronizer* is used to support fast event management and parallel thread dispatching/synchronization to enable ultra-fine grain parallelism on the cluster cores. Finally, the cluster includes a *hardware convolution engine* (HWCE) that can be used to accelerate non-strided convolutions in convolutional neural networks. CLUSTER and SoC share a single address space and communicate with one another by means of two 64-bit AXI ports, one per direction.

A software runtime resident in the fabric controller overviews all tasks offloaded to the cluster and to the $\mu DMA$. On turn, a low-overhead runtime on the cluster cores exploits the hardware synchronizer to implement shared-memory parallelism in the fashion of OpenMP.

## IV. CNN MAPPING METHODOLOGY

In this section, we discuss and characterize the main methodological aspects related to the deployment of DroNet on top of the GAP8 embedded processor.

### A. Deploying DroNet on GAP8

After a first characterization, we estimated the original convolutional neural network (CNN) to involve $\sim$41 MMAC operations per frame (accounting only for convolutional layers), yielding a baseline for the amount of computation to be performed. Note that, the given amount of multiply-accumulate operations does not account for all the additional operations required to feed the registers required by the GAP8's vectorized DSP instructions (for a detailed discussion see Section VI-B). To successfully deploy the CNN on top of GAP8, these operations have to be fit within the strict real-time constraints dictated by the target application, while respecting the bounds imposed by the on-chip and on-board resources. The primary application-related constraint is given by the minimum real-time frame-rate required to select a new

TABLE II: DroNet accuracy on PULP.

| | | Float32 | | | Fixed16 | | |
|---|---|---|---|---|---|---|---|
| Batch Normalization | | yes | no | | yes | no | |
| Max Pooling | | $3 \times 3$ | $3 \times 3$ | $2 \times 2$ | $3 \times 3$ | $3 \times 3$ | $2 \times 2$ |
| Classification | Accuracy | 0.954 | 0.955 | 0.954 | 0.945 | 0.954 | 0.953 |
| | Precision | 0.895 | 0.935 | 0.983 | 0.859 | 0.929 | 0.970 |
| | Recall | 0.895 | 0.858 | 0.810 | 0.902 | 0.857 | 0.815 |
| | F1-score | 0.895 | 0.895 | 0.888 | 0.880 | 0.892 | 0.886 |
| Regression | EVA | 0.737 | 0.695 | 0.661 | 0.716 | 0.696 | 0.670 |
| | RMSE | 0.109 | 0.120 | 0.124 | 0.117 | 0.120 | 0.123 |

trajectory on-the-fly or to detect a suspected obstacle in time to prevent a potential collision. On the other hand, the main resource-related constraint is represented by the amount of available memory on the GAP8 SoC, where as reported in Section III-B we can rely on $512\,\mathrm{kB}$ of L2 SRAM and $64\,\mathrm{kB}$ of shared L1 scratchpad (TCDM).

Therefore, it is clear that there is a strong need for a strategy aimed at reducing the memory footprint and computational load to more easily fit within the available resources, while exploiting the architectural parallelism at best to meet the real-time constraint. While applied to our specific scenario, the methodology we present in the following of this section is general in nature and could be applied also to other resource-bounded embedded systems where computationally intense tasks have to be performed under a real-time constraint on a parallel architecture.

In order to meet the tight real-time constraints of our application given the limited amount of resources on the proposed parallel ultra-low-power architecture, the original *DroNet* network [1] has been modified to ease its final deployment; we operated incrementally on the model and training flow provided by the original *DroNet*, based on Keras/TensorFlow [4]. We performed one main change, quantization, and two additional optimizations: the removal of batch normalization layers and the switch of max-pooling from $3{\times}3$ to $2{\times}2$ receptive field. Quantization, that is lowering from the full-precision *Float32* representation used in training to a *Fixed16* one, more economical for embedding, has been performed directly in training by modifying the convolution model used by the training tool. We statically targeted a *Fixed16* Q3.13 signed representation, which means that 13 bits are used to represent the fractionary part of both activations and weights (rounding down to a precision of $2^{-13}$), and values are clipped at $\pm 4$. A second change performed is the removal of batch normalization layers. We deemed the slight boost they can deal to the training process not justified by the additional complexity of either merging them with convolutional layers (which would anyway require network topology changes) or implementing them in inference (which is essentially wasted computation). Finally, we modified the receptive field of max-pooling layers from $3{\times}3$ to $2{\times}2$, which yields essentially the same final results while reducing the execution time of max-pooling layers by $2.2\times$ and simplifying their final implementation on GAP8. Table II resumes the

[4]https://github.com/uzh-rpg/rpg_public_dronet

results in terms of accuracy after all these changes.

### B. Tiling Methodology

One of the biggest constraint in ULP embedded SoC's, where the memory hierarchy is explicitly managed, is the buffer memory. To cope with it, many tiling approaches have been presented [36], commonly tiling the input, output and filter spaces. As part of this work we propose a tiling methodology that i) ensures the optimal use of the GAP8 memory ii) relieves the user from fine-grained, tedious optimizations. Data can be efficiently moved from and to L2 and L1 by the DMA or $\mu$DMA engines, but code restructuring and organization can prove to be error-prone and time-consuming. To ease development, a tool called *AutoTiler* has been developed to automate this process.

Each layer $\mathcal{L}$ in a CNN operates on a three-dimensional tensor representing a *feature space* and produces a new 3D tensor as output. Convolutional layers, in particular, are composed of i) a linear transformation that maps $n_{if}$ input feature maps into $n_{of}$ output feature maps by means of $n_{if} \times n_{of}$ convolutional kernels of size $k \times k$; ii) a pointwise non-linear activation, often a rectifier (ReLU) or a hyperbolic tangent. Figure 4 shows how a convolutional layer can be tiled in all dimensions in grids of $\mathcal{N}_i \times \mathcal{H} \times \mathcal{W}$ input tiles, $\mathcal{N}_o \times \mathcal{H} \times \mathcal{W}$ output tiles and $\mathcal{N}_i \times \mathcal{N}_o$ filter tiles.
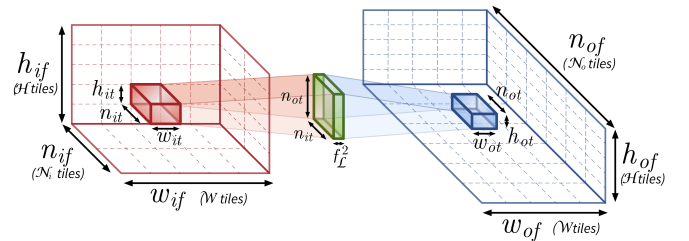


Fig. 4: Convolutional layer tiling.

With our methodology we can distinguish i) which activation dimensions are tiled, ii) in which order they are tiled (i.e. the order of the loops spanning the tiling grid), iii) which filter dimensions are tiled. Our methodology follows a *divide et impera* approach. First, we split computation in "atomic" building blocks called *basic kernels*, where we emphasize parallelization and fine grain optimizations, without addressing any specific data placement (i.e., implicitly assuming data is placed in L1). Second, we combine basic kernels into *user*

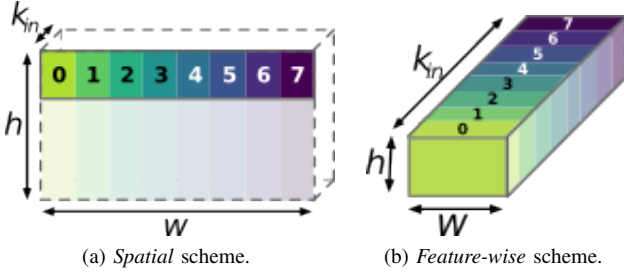(a) *Spatial* scheme.  (b) *Feature-wise* scheme.

Fig. 5: Parallelization schemes utilized in the DroNet layers for deployment on GAP8; the different colors represent allocation to a different core.

*kernels*, where we describe the multi-dimensional tiled iteration space. A user kernel contains a collection of connected basic kernels that can be inserted at pre-defined locations of the iteration space (e.g., prologue, body, epilogue, etc.).

Using this model and a L1 memory budget constraint, the tool infers the best tiling structure fitting within the given budget and satisfying the set of constraints put on the basic kernels. Once the tiling structure has been computed the *AutoTiler* generates a wrapper (as C code) orchestrating the execution of the basic kernels in a pipelined fashion as well as double-buffered DMA transactions. With this approach it is possible to build generators for specific algorithms and then combine these together.

### C. Parallelization & Optimization

As introduced in Section III, the GAP8 SoC features 8+1 RISC-V cores with DSP-oriented extensions. To develop an optimized, high-performance and energy-efficient application for GAP8 and meet the required real-time constraint, the most important steps include *i*) parallelizing the most computationally intense kernels of the algorithm, to take advantage of the 8-core cluster; *ii*) fully using the available specialized instruction; and *iii*) orchestrating the movement of data through the various levels of the memory hierarchy. The following of this section deals with steps i) and ii), while Section IV-D reports our memory management strategy.

We parallelized all of the layers exposed in the *DroNet* architecture (Figure 2): convolutional, fully-connected, max-pooling, activation (ReLU), and sum (for the residual bypasses). To exploit the available computational/memory resources at best we use two different paralellization schemes on the basis of the layer's depth. Feature maps (FMs) have decreasing size due to the striding factor used in several convolutional layers, as well as to max pooling layers. In the first half of the network the FM size ranges from $200 \times 200$ to $25 \times 25$. In this part of the network, parallelism can be efficiently extracted with a *spatial* parallelization scheme, dividing each two-dimensional tile in 8 vertical stripes, one per core, and scheduling part of the overall 2D convolution to each core as shown in Figure 5a. For deeper layers in the network, the FM size drops so much that this scheme becomes suboptimal as the width of each stripe becomes too small. Therefore, we introduce a new scheme exploiting *feature-wise*

parallelism as shown in Figure 5b. On these deepest layers, the whole input set of a layer is divided into 8 slices on the feature map dimension, and each core is responsible for processing the multiple FMs consituting a slice. The *feature-wise* scheme has a heavier memory footprint to keep track of all the intermediate results and requires a final reduction stage to combines all the partial results from the fan-in of all input FMs into a final output pixel; however, this cost is more than offset by the possibility to achieve full utilization of all the cluster cores.

To further optimize the DroNet execution, we made use of all the optimized signal processing instructions available in GAP8. These include packed-SIMD instructions capable of exploiting sub-word parallelism, as well as bit-level manipulation and shuffling, which can be accessed by means of compiler intrinsics such as `__builtin_pulp_dotsp2` (for 16-bit dot product with 32-bit accumulation), `__builtin_shuffle` (permutation of elements within two input vectors), `__builtin_pulp_pack2` (packing two scalars into a vector).

### D. L2 Memory Management Strategy

Given i) the residual-network topology of *DroNet*, which requires to increase the lifetime of the output tensors of some of the layers, and ii) the "scarsity" of L2 memory as a resource to store all weights and temporary feature maps (we would need more than $1\,\mathrm{MB}$ in view of $512\,\mathrm{kB}$ available), an ad-hoc memory management strategy for the L2 memory is required, similar to what is done between L2 and L1 using the GAP8 automatic *AutoTiler*. Due to the high energy cost of data transfers between L3 and L2, the strategy needs to be aimed at the maximization of the L2 reuse.

At boot time, before the actual computation loop starts, i) we load all the weights, stored in the external flash memory as binary files, in the L3 DRAM memory and ii) we call from the fabric controller (FC) the *runtime* allocator to reserve two *L2 allocation stacks* where intermediate buffers will be allocated and deallocated in a linear fashion. Employing only one buffer and a linear allocation scheme we would need to keep in L2 memory up to $665\,\mathrm{kB}$ of data due to data dependencies. Our allocation strategy simply updates the pointer of the next free location in the pre-allocated L2 chunk, avoiding the *runtime* overhead of library allocation/free functions.

We differentiate our strategy between weights and FM activations: for the former, we allocate space just before their related layer and deallocate it just after the layer execution. For the latter, due to the residual network bypasses, we often have to prolongate the lifetime of a FM during the execution of the two following layers. Therefore, for each RES block there will be an amount of time where 3 FMs activation tensors have to be stored at the same time (see Figure 2).

Algorithm 1 shows the complete pseudo-code related to our solution. In the pseudo-code the second parameter of the `Alloc` and `Free` function specifies the allocation buffer (i.e., `Buffer 0` or `Buffer 1` in Table III). Note that, the $\mu$DMA copies the weights from L3 to L2 just after the destination L2 area is allocated (for the sake of readability not shown in Algorithm 1).

## Algorithm 1: DroNet on PULP execution flow

**Input:** image
**Output:** steering angle, collision probability
1: $I_1$ = image; Alloc $(O_1, 0)$; Alloc $(w_1, 0)$;
2: $O_1 \leftarrow \text{Conv}_1(I_1, w_1)$;  // Conv2d + MaxPool2d
3: Free $(w_1, 0)$; $I_2 = O_1$; $O_2 = O_1$;
RES BLOCK 1
4: $O_2 \leftarrow \text{ReLU}_1(I_2)$;
5: $I_3 = O_2$; Alloc $(O_3, 1)$; Alloc $(w_2, 1)$;
6: $O_3 \leftarrow \text{Conv}_2(I_3, w_2)$;  // Conv2d + ReLU
7: Free $(w_2, 1)$; $I_4 = O_3$; Alloc $(O_4, 0)$; Alloc $(w_3, 1)$;
8: $O_4 \leftarrow \text{Conv}_3(I_4, w_3)$;  // Conv2d
9: Free $(w_3, 1)$; Free $(O_3, 1)$; $I_5 = O_1$; Alloc $(O_5, 1)$;
10: Alloc $(w_4, 1)$;
11: $O_5 \leftarrow \text{Conv}_4(I_5, w_4)$;  // Conv2d
12: Free $(w_4, 1)$; $I_6 = O_4$; $O_6 = O_5$;
13: $O_6 \leftarrow \text{Add}_1(I_6, O_6)$;
14: Free $(O_4, 0)$; Free $(O_1, 0)$; $I_7 = O_6$; $O_7 = O_6$;
RES BLOCK 2
15: $O_7 \leftarrow \text{ReLU}_2(I_7)$;
16: $I_8 = O_7$; Alloc $(O_8, 0)$; Alloc $(w_5, 0)$;
17: $O_8 \leftarrow \text{Conv}_5(I_8, w_5)$;  // Conv2d + ReLU
18: Free $(w_5, 0)$; $I_9 = O_8$; Alloc $(O_9, 1)$; Alloc $(w_6, 0)$;
19: $O_9 \leftarrow \text{Conv}_6(I_9, w_6)$;  // Conv2d
20: Free $(w_6, 0)$; Free $(O_8, 0)$; $I_{10} = O_5$; Alloc $(O_{10}, 0)$;
21: Alloc $(w_7, 0)$;
22: $O_{10} \leftarrow \text{Conv}_7(I_{10}, w_7)$;  // Conv2d
23: Free $(w_7, 0)$; $I_{11} = O_9$; $O_{11} = O_{10}$;
24: $O_{11} \leftarrow \text{Add}_2(I_{11}, O_{11})$;
25: Free $(O_9, 1)$; Free $(O_5, 1)$; $I_{12} = O_{11}$; $O_{12} = O_{11}$;
RES BLOCK 3
26: $O_{12} \leftarrow \text{ReLU}_3(I_{12})$;
27: $I_{13} = O_{12}$; Alloc $(O_{13}, 1)$; Alloc $(w_8, 0)$;
28: $O_{13} \leftarrow \text{Conv}_8(I_{13}, w_8)$;  // Conv2d + ReLU
29: Free $(w_8, 0)$; $I_{14} = O_{13}$; Alloc $(O_{14}, 1)$; Alloc $(w_9, 0)$;
30: $O_{14} \leftarrow \text{Conv}_9(I_{14}, w_9)$;  // Conv2d
31: Free $(w_9, 0)$; $I_{15} = O_{10}$; Alloc $(O_{15}, 0)$;
32: Alloc $(w_{10}, 1)$;
33: $O_{15} \leftarrow \text{Conv}_{10}(I_{15}, w_{10})$;  // Conv2d
34: Free $(w_{10}, 1)$; $I_{16} = O_{14}$; $O_{16} = O_{15}$;
35: $O_{16} \leftarrow \text{Add}_3(I_{16}, O_{16})$;  // Add + ReLU
36: Free $(O_{14}, 1)$; Free $(O_{13}, 1)$; $I_{17} = O_{16}$;
37: Alloc $(O_{17}, 1)$; Alloc $(w_{11}, 1)$;
FULLY CONNECTED
38: $O_{17} \leftarrow \text{Dense}_1(I_{17}, w_{11})$;
39: Free $(w_{11}, 1)$; $I_{18} = O_{16}$; Alloc $(O_{18}, 1)$; Alloc $(w_{12}, 1)$;
40: $O_{18} \leftarrow \text{Dense}_2(I_{18}, w_{12})$;
41: Free $(w_{12}, 1)$; Free $(O_{15}, 0)$; Free $(O_{10}, 0)$;
42: **return** $O_{17}, O_{18}$

The buffers' memory allocation sequence is reported in Table III (from left to right) for the entire DroNet execution. The columns of the two buffers represent the data needed at each execution step, where $O_i$ and $w_j$ represent the input/output and weights, respectively. The last row of each buffer reports the total amount of memory required for each execution step. Thus, the final dimension of each buffer is given by the column with the biggest occupancy (highlighted in red in Table III), resulting in $370\,\text{kB}$ of L2 memory.

## V. THE PULP-SHIELD

In order to enable the development of a first prototype, we designed a lightweight, modular and configurable printed circuit board (PCB) with highly optimized layout and a form factor compatible to our nano-size quadrotor. It features a

TABLE III: L2 memory allocation sequence.

| Buffer 0 | | | | | | | |
|---|---|---|---|---|---|---|---|
| $O_1 \longrightarrow O_1$ | | $O_8 \longrightarrow O_8$ | | $O_{10} \longrightarrow O_{10}$ | | $O_{10}$ | $O_{10}$ |
| $w_1$ | $O_4$ | $w_5$ | $w_6$ | $w_7$ | $w_8$ | $w_9$ | $O_{15}$ |
| $162\,kB$ | $200\,kB$ | $59\,kB$ | $96\,kB$ | $26\,kB$ | $169\,kB$ | $308\,kB$ | $26\,kB$ |

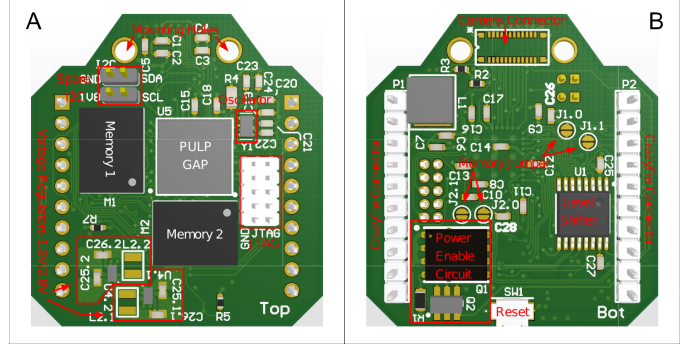| Buffer 1 | | | | | | | |
|---|---|---|---|---|---|---|---|
| $O_3 \longrightarrow O_3$ | | $O_5$ | $O_5$ | $O_{13} \longrightarrow O_{13}$ | | $O_{17} \longrightarrow O_{17}$ | |
| $w_2$ | $w_3$ | $w_4$ | $O_9$ | $O_{14} \longrightarrow O_{14}$ | | $w_{11}$ | $O_{18}$ |
| | | | | | $w_{10}$ | | $w_{12}$ |
| $58\,kB$ | $58\,kB$ | $42\,kB$ | $62\,kB$ | $26\,kB$ | $42\,kB$ | $13\,kB$ | $13\,kB$ |



Fig. 6: Schematic of the *PULP-Shield* pluggable PCB. Top view (A) and bottom view (B).

PULP-based GAP8 SoC, two Cypress *HyperBus Memories*[5] and an ultra-low power *HiMax* CMOS image sensor[6] able to run up to 60 fps with a gray-scale resolution of $320{\times}320$ pixels with just $4.5\,\text{mW}$ of power. Our pluggable PCB, named *PULP-Shield*, has been designed to be compatible with the *Crazyflie 2.0* (CF) nano-quadrotor[7]. This vehicle has been chosen due to its reduced size (i.e., $27\,\text{g}$ of weight and $10\,\text{cm}$ of diameter) and its open-source and open-hardware philosophy. The communication between the PULP chip and the main MCU on-board the nano-drone (i.e., *ST Microelectronics STM32F405*[8]) is realized via a SPI interface and two GPIO signals.

In Figure 6 the schematic of the proposed *PULP-Shield* is shown. The two BGA memory slots allow all the possible combinations of *HyperRAM*, *HyperFlash* and hybrid *Hyper-Flash/RAM* memory packages. In this way we can select the most appropriate memory configuration with respect to the target application. In our use case we soldered on one slot a $64\,\text{Mbit}$ *HyperRAM* (DRAM) chip and on the other a $128\,\text{Mbit}$ *HyperFlash* memory, embodying the system L3 and the external storage, respectively.

On the PCB (Figure 6-B) there is also a camera connector that allows the *HiMax* camera to communicate with the rest of the system through the parallel camera interface (PCI) protocol. Two mounting holes, on the side of the camera connector, allow to plug a 3D-printed camera holder that can be set either in front-looking or down-looking configuration. Those two configurations are representative of the most common visual

---

[5]http://www.cypress.com/products/hyperbus-memory
[6]http://www.himax.com.tw/products/cmos-image-sensor/image-sensors
[7]https://www.bitcraze.io/crazyflie-2
[8]http://www.st.com/en/microcontrollers/stm32f405-415.html

- ❶ Init interrupt (GPIO)
- ❷ Load binary
- ❸ Configure camera (I2C)
- ❹ Grab frames (μDMA)
- ❺ LD Weights and Exec
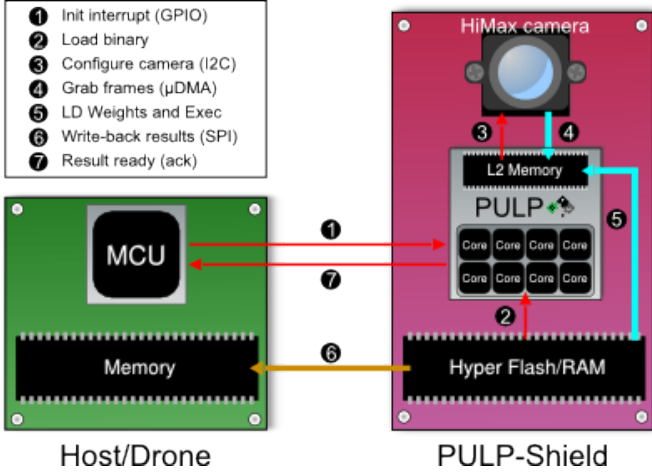- ❻ Write-back results (SPI)
- ❼ Result ready (ack)

Fig. 7: Example of interaction between the *PULP-Shield* and the drone.

sensors layouts typically embedded in any autonomous flying vehicles. In fact, the front-looking configuration can be used for many navigation tasks like path planning [37], obstacle avoidance [38], trajectory optimization [7], just to name a few. Instead, the down-looking camera configuration is usually chosen for stabilization tasks like distance estimation [39], way-point tracking and positioning [40], etc.

On the shield there are also a JTAG connector for debug purposes and an external I2C plug for future development. Two headers, located on both sides of the PCB, grant a steady physical connection with the drone and at the same time they bring the shield power supply and allow communication with the CF through the GPIOs and the SPI interface. The form factor of our final *PULP-Shield* prototype is $30 \times 28\,\mathrm{mm}$ and it weighs $\sim 5\,\mathrm{g}$ (including all components), well below the payload limit imposed by the nano-quadcopter.

Similarly to what has been presented in [41], the *PULP-Shield* embodies the *Host-Accelerator* architectural paradigm, where the CF's MCU offloads the compute-intense work to the PULP accelerator. As depicted in Figure 7 the interaction starts from the host, which wakes up the accelerator with a GPIO interrupt ❶. Then, the accelerator fetches from its external *HyperFlash* storage the kernel (stored as binary file) to be executed: *DroNet* in our case ❷. Note that, in this first part of the protocol the host can also specify which kernel should be executed, as well as a sequence of several pre-loaded kernels available on the external Flash storage. A this point, the GAP8 SoC can configure the *HiMax* camera via an internal I2C ❸ and start to transfer the frames from the sensor to the L2 shared memory through the μDMA ❹. All additional data, like the weights used in our CNN, can be loaded from the DRAM/Flash memory and parallel execution is started on the accelerator ❺. Lastly, the results of the computation are returned to the drone's MCU via SPI ❻ and the same host is acknowledged about the available results with a final interrupt over GPIO ❼.

Even if the *PULP-shield* has been developed specifically to fit the CF quadcopter, its basic concept and the functionality it provides are quite general, and portable to any drone based

on an SPI-equipped MCU. The system-level architectural template it is based on is meant for minimizing data transfers (i.e., exploiting locality of data) and communication overhead between the main MCU and the accelerator – without depending on the internal microarchitecture of either one.

## VI. EXPERIMENTAL RESULTS

In this section we present our experimental evaluation, considering three main metrics: i) the capability of respecting a given real-time deadline, ii) the capability of performing all the required computations within the allowed power budget and iii) the final accuracy of the closed-loop control, given as reaction time w.r.t. an unexpected obstacle. All the results are based on the *PULP-shield* configuration presented in Section V.

### A. Performance & Power Consumption

In Figure 8 we present the *DroNet* power traces for the execution of each single layer, measured using a bench oscilloscope. In the full *DroNet* execution, these are interposed with L3-L2 data transfers, happening with the cluster cores in clock-gated state, which account for $\sim 5\%$ of the overall execution time. The power traces are measured by powering the GAP8 SoC, with an experimental configuration at a $1.0\,\mathrm{V}$ core voltage and operating at $165\,\mathrm{MHz}$. The detailed average power consumption (including both the SOC and CLUSTER domains) is reported in Table IV. The peak power consumption of $92\,\mathrm{mW}$ is associated to the $6^{th}$ convolutional layer; we used this value to compute the overall power envelope of our node. The average power consumption, weighted over the duration of each layer, is $84\,\mathrm{mW}$, which grows to $94\,\mathrm{mW}$ if we consider also the cost of L3 memory access and the on-board ULP camera.

TABLE IV: Power consumption & Execution time of DroNet on GAP8 @ $1.0\,\mathrm{V}$, $165\,\mathrm{MHz}$.

| Layer | AVG Power [mW] | Exec Time [ms] | L3-L2 Time [ms] |
|---|---|---|---|
| conv_1 + pool | 89.01 | 12.95 | 0.02 |
| ReLU | 52.07 | 0.46 | — |
| conv_2 + ReLU | 87.96 | 11.12 | 0.13 |
| conv_3 | 77.46 | 11.85 | 0.13 |
| conv_4 | 83.56 | 5.49 | 0.02 |
| add | 37.75 | 0.12 | — |
| ReLU | 36.81 | 0.12 | — |
| conv_5 + ReLU | 78.10 | 8.34 | 0.26 |
| conv_6 | 91.72 | 9.35 | 0.51 |
| conv_7 | 66.53 | 2.08 | 0.04 |
| add | 34.98 | 0.12 | — |
| ReLU | 34.01 | 0.11 | — |
| conv_8 + ReLU | 87.49 | 6.37 | 1.01 |
| conv_9 | 89.67 | 11.91 | 2.02 |
| conv_10 | 66.00 | 2.59 | 0.12 |
| add + ReLU | 34.30 | 0.11 | — |
| fully_1 | 30.76 | 0.04 | 0.09 |
| fully_2 | 36.74 | 0.04 | 0.09 |

Table V provides a complete view of the power consumption in all theoretically possible operating modes of GAP8 on the DroNet application while sweeping the clock frequency, both
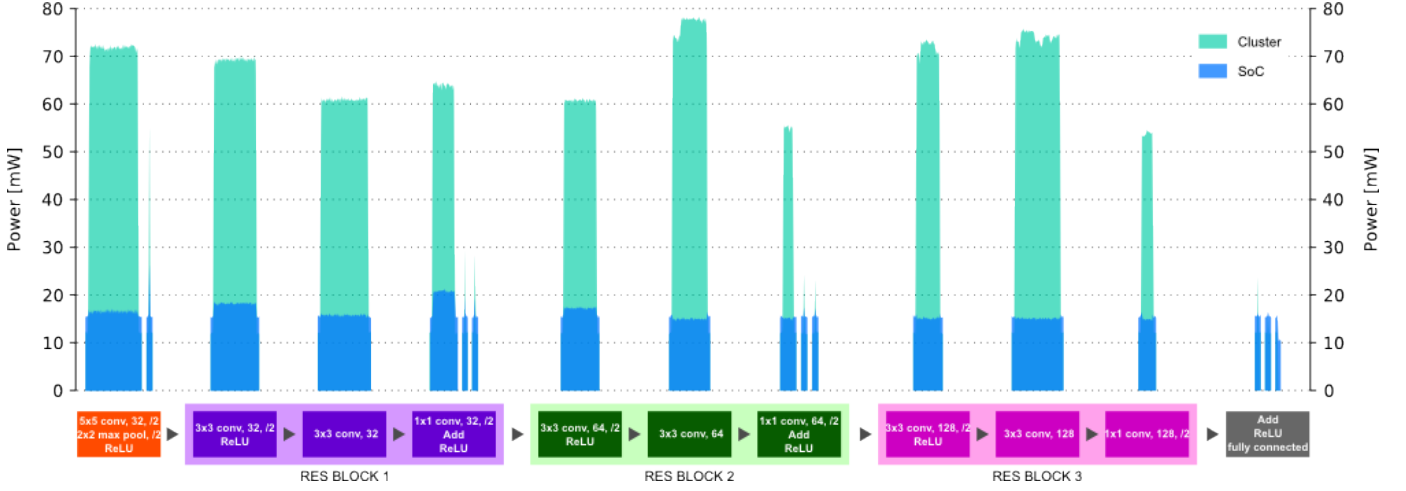
Fig. 8: Power traces per layer of *DroNet*, measured @ 1.0 V,165 MHz.

at 1.0 V and 1.2 V, and the related achievable frame rate. As the DC/DC converter utilized on the SoC is currently only able to deliver up to 100 mA, some of these operating modes are not achievable (shown in italic font). If we bypassed this limitation, selecting an operating point of 1.2 V would increase both power and performance up to 191 mW and 18 fps. However, we note that for frequencies lower than 175 MHz the 1.0 V mode has a definite advantage in terms of energy efficiency.

TABLE V: Power consumption and frame-rate of GAP8 @ 1.0 V and 1.2 V for different frequencies.

| SoC MHz | AVG mW @ 1.0 V | AVG mW @ 1.2 V | frame-rate [fps] |
|---|---|---|---|
| 50 | 25.50 | 38.25 | 4 |
| 75 | 38.25 | 57.37 | 5 |
| 100 | 50.99 | 76.49 | 7 |
| 125 | 76.49 | 95.61 | 9 |
| 150 | 84.14 | 114.74 | 11 |
| 175 | — | *133.86* | 13 |
| 200 | — | *152.98* | 15 |
| 225 | — | *172.11* | 16 |
| 250 | — | *191.23* | 18 |

In Figure 9 is reported the power break-down for the complete cyberphysical system and for proposed *PULP-Shield*. Our nano-quadcopter is equipped with a 240 mA h 3.7 V LiPo battery enabling a flight time of 7 minutes under standard conditions, which results in an average power consumption of 7.6 W. The power consumption of all the electronics on-board of the original drone amounts to 277 mW leaving 7.3 W for the 4 rotors. The electronics consumption is given by the 2 MCUs included in the quadrotor and all the additional devices (e.g., sensors, leds, etc.). In addition to that, introducing the *PULP-Shield*, we increase the peak power envelope by 94 mW (1% of the total) running the GAP8 SoC at the frequency of 165 MHz. Notice that in this case we consider the HyperRAM operating at full speed only for the time required for L3-L2 data transfers (as shown in Table IV) with an average power consumption of 84 mW. The power break-down of our pluggable PCB can be seen on the right of Figure 9, where
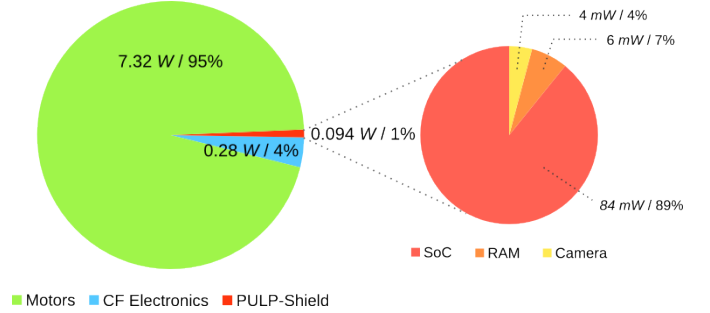


Fig. 9: Power envelope break-down of the entire cyberphysical system (on the left) and *PULP-Shield* zoom-in @ 165 MHz (on the right).

we include the computational unit, the L3 external DRAM memory and the ultra-low power camera.

As on-board computation accounts for roughly 5 % of the overall power consumption (propellers, sensors, compute and control, *cfr* Section I), our *PULP-Shield* enables the execution of the *DroNet* network (and potentially more) in the given power envelope.

### B. State-of-the-Art Comparison & Discussion

To compare and validate our experimental results with respect to the current state-of-the-art, we targeted the most efficient CNN implementation available for microcontrollers, namely CMSIS-NN [42]. At peak performance in a synthetic test, this fully optimized library is able to achieve as much as 0.69 MAC/cycle on convolutions, operating on *Fixed8* data that is internally converted to *Fixed16* in the inner loop.

By contrast, we operate directly on *Fixed16* and achieve a peak performance of 0.64 MAC/cycle/core in a similar synthetic scenario (on the $6^{th}$ layer, 3×3 convolution). The bypasses and the final layers are a bit less efficient, yielding an overall weighted peak throughput of 0.53 MAC/cycle/core on convolutional layers, which constitute the vast majority of the execution time.

TABLE VI: Cycle break-down for processing one frame on the GAP8 cluster.

|  | $\mu DMA$ L3/L2 | Marshaling | $DMA$ L2/L1 | Comput. | Total |
|---|---|---|---|---|---|
| Cycles | 0.7 M | 0.3 M | 0.8 M | 12.6 M | 14.4 M |

In Table VI we report the execution breakdown per frame for all activities performed by our CNN. We can see how the L3-L2 transfers (not overlapped to computation), the marshaling stage (required by padded convolutions) and the non-overlappable part of L2-L1 transfers account for $\sim$2 Mcycles of the overall execution time. Then, considering $\sim$41 MMAC for the original CNN, in the ideal peak-throughput case of 4.28 MAC/cycle we would need $\sim$10 Mcycles for computing one frame, instead of our measured 12.6 Mcycles. The overhead is due to inevitable non-idealities such as sub-optimal load balancing in layers exposing limited spatial parallelism as well as tiling control loops. Considering all of the aforementioned effects (i.e. computation non-idealities as well as memory transfers), we achieve a real throughput of 2.85 MAC/cycle in the *DroNet* execution – still 4$\times$ better than the CMSIS-NN peak performance.

To further concretize the comparison, we take as an example target a top-notch high-performance microcontroller: a STM32H7[9] sporting a Cortex-M7 core and capable of operating at up to 400 MHz. Without considering any data movement overhead, and taking into account only peak performance, this would be able to achieve up to 276 MMAC/s @ 346 mW. By comparison, our system can achieve an average performance of 470 MMAC/s @ 84 mW, i.e. 70% better within a 4.1$\times$ smaller power budget. Even if it was possible to linearly up-scale the performance of this microcontroller to the same level of our system, it would consume $\sim$580 mW, which together with the $\sim$277 mW consumed by the standard Crazyflie MCUs would constitute more than the 5% of power envelope typically dedicated to on-board computation on nano-UAV systems [15].

### C. Control Accuracy

In order to unlock the agility of a lightweight nano-quadrotor as the Crazyflie 2.0 used in our prototype, fast on-board perception is required. To evaluate the agility of our system, we performed an experiment in which our flying platform is required to react to a suddenly appearing obstacle occluding its way. The main goal of this experiment is to show that the *PULP-Shield* computational resources are enough to make full use of the nano-platform agility. More specifically, the test consists of simulating a high speed flight on a straight track of 20 m. At the beginning of the test, the track is completely free from obstacles. However, at the end of the track at T=4s after the start of the experiment, an obstacle appears. The platform needs to stop fast enough in order to avoid collision. We collected the associated dataset manually in flight-realistic conditionds, and analyze the performance of the system across a range of operational frequencies. Figure 10 and 11 illustrate the predicted collision probability of the original and quantized *DroNet* architecture as a function of

time. Exactly as in [1], the predictions of the network are low pass filtered before being passed to the controller. Indeed, reducing high-frequency noise is necessary in order to perform a safe and smooth flight. In detail, the collision probability $p_k$ is a low-pass filtered version of the collision probability $c_k$ predicted by the network ($\alpha = 0.7$):

$$p_k = (1 - \alpha)p_{k-1} + \alpha c_k, \tag{1}$$

A stop signal is given to the controller when $p_k > 0.5$.

In order to quantitatively evaluate the performance of our system at different operational frequencies, we computed the maximum time and the minimum distance from the object at which the stop command should be given to avoid collision. We deployed the Crazyflie 2.0 parameters from [43] and the classical quadrotor motion model from [44] to *analytically* compute those two quantities. From this analysis, we derived a minimum stopping time of $T = 0.40\,\text{s}$ and a braking distance of $D_b = 0.7\,\text{m}$, assuming the platform initially moves with a speed of $V = 4\,\text{m/s}$. Figure 10 illustrate the results of this evaluation. As expected, the higher the inference frequency, the quicker the platform can react to the obstacle. It is possible to notice very little change in performance between the two architectures, in particular for frequencies close to our operational one of 10 Hz. More importantly, from Figure 10 we can observe that inference at 10 Hz allows the platform to brake in time and avoid collision. Conversely, lower inference rates (e.g., 5 Hz) would not be enough to ensure reactivity of the platform, compromising its safety. This confirms that our system, processing 12 fps, can: i) make use of the agility of the Crazyflie 2.0 and ii) be deployed in the same way as the original method in [1] to perform a safe and reliable flight in a large set of application environments.

## VII. CONCLUSION

Nano- and pico-sized UAVs are ideal IoT nodes; due to their size and physical footprint, they can act as mobile IoT hubs, smart sensors and data collectors for tasks such as surveillance, inspection etc. However, to be able to perform these tasks, they must be capable of autonomous navigation of environments such as urban streets, industrial facilities and other hazardous or otherwise challenging areas. In this work, we provide the first (to the best of our knowledge) completely integrated hardware/software solution for autonomous navigation of nano-UAVs with completely on-board computation – and thus potentially able to operate in conditions in which the latency or the additional power cost of a wirelessly-connected centralized solution.

Our system, based on a GreenWaves GAP8 SoC used as an accelerator coupled with the STM32 MCU on the CrazyFlie 2.0 UAV, supports real-time (12 fps) computation of *DroNet*, an advanced CNN-based autonomous navigation algorithm, within an average power budget for computation of 94 mW. This is achieved without quality-of-results loss with respect to the baseline system on which DroNet was deployed: a COTS standard-size UAV connected with a remote PC, on which *DroNet* was running at 20 fps. Our results show that both systems are able to detect obstacles fast enough to be
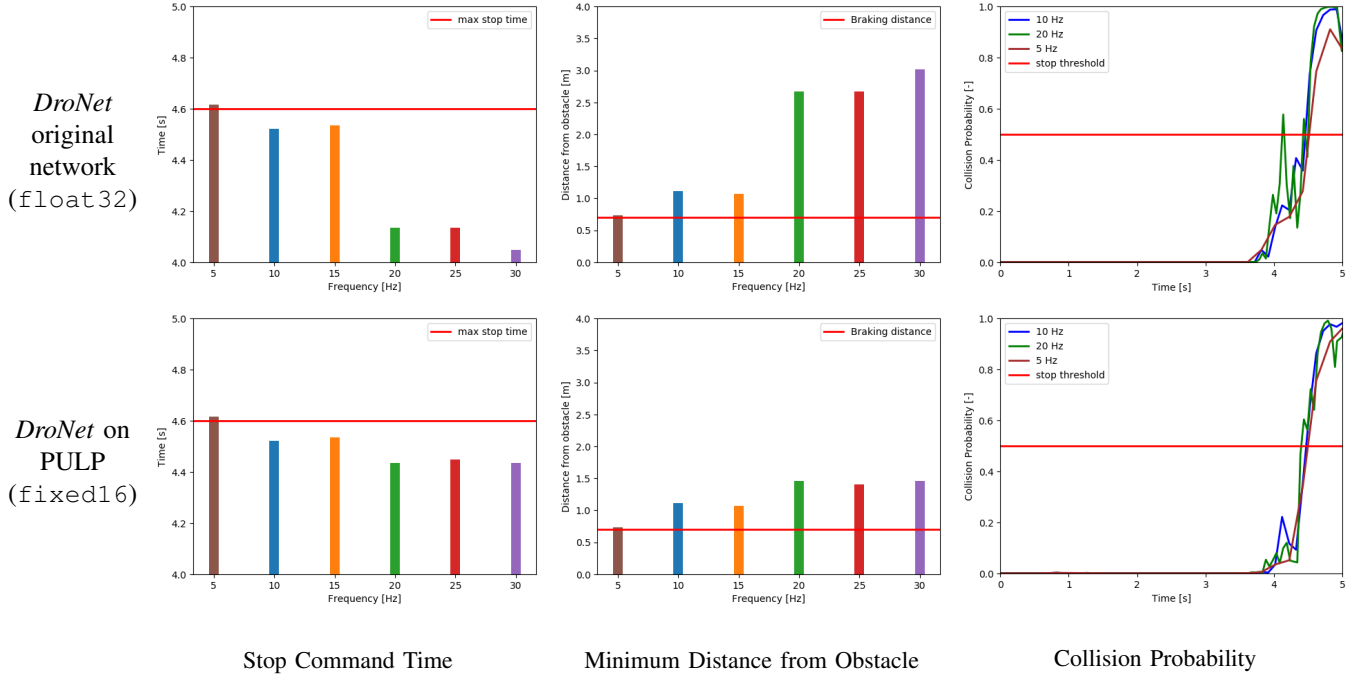
Fig. 10: Performance comparison between the original and quantized *DroNet* architectures. While the results are extremely similar for frequencies smaller than 15 Hz, small variations are observed at higher frequencies.
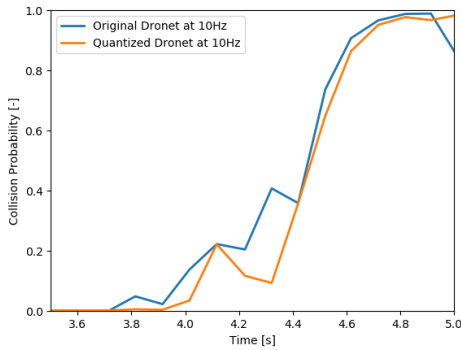


Fig. 11: Comparison of predicted probability of collision in the last 1.5 s of the experiment. The reported frequency corresponds to our system operational frequency of 10 Hz.

able to safely fly at high speed, 4 m/s in the case of the CrazyFlie. Overall, our contribution paves the way for a huge number of advanced use-cases of autonomous nano-UAVs as IoT-connected mobile smart sensors.

Our future work will be focused on advanced field tests of our prototype and on further optimization of the overall system for better performance and lower overall power consumption, by *i*) making use of the advanced functionalities available on the GAP8 SoC, as the Fabric Controller (which could in principle fully replace the CrazyFlie MCU) and the HWCE accelerator for convolutions; *ii*) further streamlining the *DroNet* algorithm for similar or better accuracy at a lower time/energy cost.

## REFERENCES

[1] A. Loquercio, A. I. Maqueda, C. R. del Blanco, and D. Scaramuzza, "Dronet: Learning to fly by driving," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, April 2018.

[2] I. Sa, M. Kamel, M. Burri, M. Bloesch, R. Khanna, M. Popovic, J. Nieto, and R. Siegwart, "Build your own visual-inertial drone: A cost-effective and open-source autonomous drone," *IEEE Robotics & Automation Magazine*, vol. 25, no. 1, pp. 89–103, mar 2018.

[3] D. Scaramuzza, M. C. Achtelik, L. Doitsidis, F. Friedrich, E. Kosmatopoulos, A. Martinelli, M. W. Achtelik, M. Chli, S. Chatzichristofis, L. Kneip, D. Gurdan, L. Heng, G. H. Lee, S. Lynen, M. Pollefeys, A. Renzaglia, R. Siegwart, J. C. Stumpf, P. Tanskanen, C. Troiani, S. Weiss, and L. Meier, "Vision-controlled micro flying robots: From system design to autonomous navigation and mapping in GPS-denied environments," *IEEE Robotics & Automation Magazine*, vol. 21, no. 3, pp. 26–40, sep 2014.

[4] D. Floreano and R. J. Wood, "Science, technology and the future of small autonomous drones," *Nature*, vol. 521, no. 7553, pp. 460–466, may 2015.

[5] Y. Lin, F. Gao, T. Qin, W. Gao, T. Liu, W. Wu, Z. Yang, and S. Shen, "Autonomous aerial navigation using monocular visual-inertial fusion," *Journal of Field Robotics*, vol. 35, no. 1, pp. 23–51, jul 2017.

[6] G. Loianno, C. Brunner, G. McGrath, and V. Kumar, "Estimation, control, and planning for aggressive flight with a small quadrotor with a single camera and imu," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 404–411, 2017.

[7] D. Falanga, E. Mueggler, M. Faessler, and D. Scaramuzza, "Aggressive quadrotor flight through narrow gaps with onboard sensing and computing using active vision," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017.

[8] P. R. Florence, J. Carter, J. Ware, and R. Tedrake, "Nanomap: Fast, uncertainty-aware proximity queries with lazy search over local 3d data," *CoRR*, vol. abs/1802.09076, 2018. [Online]. Available: http://arxiv.org/abs/1802.09076

[9] G. Loianno, D. Scaramuzza, and V. Kumar, "Special issue on high-speed vision-based autonomous navigation of uavs," *Journal of Field Robotics*, vol. 35, no. 1, pp. 3–4, 2018.

[10] J. Nikolic, M. Burri, J. Rehder, S. Leutenegger, C. Huerzeler, and R. Siegwart, "A uav system for inspection of industrial facilities," in *2013 IEEE Aerospace Conference*, March 2013.

[11] N. Michael, S. Shen, K. Mohta, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yoshida, K. Ohno, E. Takeuchi, and S. Tadokoro, "Collaborative mapping of an earthquake damaged building via ground and aerial robots," in *Springer Tracts in Advanced Robotics*. Springer Berlin Heidelberg, dec 2013, pp. 33–47.

[12] Y. Yang, Z. Zheng, K. Bian, L. Song, and Z. Han, "Real-time profiling of fine-grained air quality index distribution using uav sensing," *IEEE Internet of Things Journal*, vol. 5, no. 1, Feb 2018.

[13] J. Conroy, G. Gremillion, B. Ranganathan, and J. S. Humbert, "Implementation of wide-field integration of optic flow for autonomous quadrotor navigation," *Autonomous robots*, vol. 27, no. 3, 2009.

[14] K. McGuire, G. de Croon, C. D. Wagter, K. Tuyls, and H. Kappen, "Efficient optical flow and stereo vision for velocity estimation and obstacle avoidance on an autonomous pocket drone," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, April 2017.

[15] R. J. Wood, B. Finio, M. Karpelson, K. Ma, N. O. Pérez-Arancibia, P. S. Sreetharan, H. Tanaka, and J. P. Whitney, *Progress on "Pico" Air Vehicles*. Cham: Springer International Publishing, 2017.

[16] M. Piccoli and M. Yim, "Piccolissimo: The smallest micro aerial vehicle," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2017.

[17] N. H. Motlagh, T. Taleb, and O. Arouk, "Low-altitude unmanned aerial vehicles-based internet of things services: Comprehensive survey and future perspectives," *IEEE Internet of Things Journal*, vol. 3, no. 6, Dec 2016.

[18] F. Conti, R. Schilling, P. D. Schiavone, A. Pullini, D. Rossi, F. K. Gürkaynak, M. Muehlberghuber, M. Gautschi, I. Loi, G. Haugou, S. Mangard, and L. Benini, "An IoT Endpoint System-on-Chip for Secure and Energy-Efficient Near-Sensor Analytics," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 64, no. 9, pp. 2481–2494, Sep. 2017.

[19] S. Scherer, J. Rehder, S. Achar, H. Cover, A. Chambers, S. Nuske, and S. Singh, "River mapping from a flying robot: state estimation, river detection, and obstacle mapping," *Autonomous Robots*, vol. 33, no. 1-2, 2012.

[20] S. Shen, Y. Mulgaonkar, N. Michael, and V. Kumar, "Multi-sensor fusion for robust autonomous flight in indoor and outdoor environments with a rotorcraft MAV," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014.

[21] S. Lynen, T. Sattler, M. Bosse, J. Hesch, M. Pollefeys, and R. Siegwart, "Get out of my lab: Large-scale, real-time visual-inertial localization," in *Robotics: Science and Systems XI*, jul 2015.

[22] S. Ross, N. Melik-Barkhudarov, K. S. Shankar, A. Wendel, D. Dey, J. A. Bagnell, and M. Hebert, "Learning monocular reactive uav control in cluttered natural environments," in *2013 IEEE International Conference on Robotics and Automation*, May 2013.

[23] N. Smolyanskiy, A. Kamenev, J. Smith, and S. Birchfield, "Toward low-flying autonomous mav trail navigation using deep neural networks for environmental awareness," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept 2017.

[24] D. Gandhi, L. Pinto, and A. Gupta, "Learning to fly by crashing," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, sep 2017.

[25] G. Kahn, T. Zhang, S. Levine, and P. Abbeel, "PLATO: Policy learning using adaptive trajectory optimization," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2017.

[26] F. Sadeghi and S. Levine, "CAD2rl: Real single-image flight without a single real image," in *Robotics: Science and Systems XIII*, jul 2017.

[27] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, may 2017.

[28] X. Zhang, B. Xian, B. Zhao, and Y. Zhang, "Autonomous flight control of a nano quadrotor helicopter in a gps-denied environment using on-board vision," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 10, Oct 2015.

[29] O. Dunkley, J. Engel, J. Sturm, and D. Cremers, "Visual-inertial navigation for a camera-equipped 25g nano-quadrotor," in *IROS2014 Aerial Open Source Robotics Workshop*, 2014.

[30] A. Briod, J.-C. Zufferey, and D. Floreano, "Optic-flow based control of a 46g quadrotor," in *Workshop on Vision-based Closed-Loop Control and Navigation of Micro Helicopters in GPS-denied Environments, IROS 2013*, no. EPFL-CONF-189879, 2013.

[31] K. McGuire, G. de Croon, C. de Wagter, B. Remes, K. Tuyls, and H. Kappen, "Local histogram matching for efficient optical flow computation applied to velocity estimation on pocket drones," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, May 2016.

[32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.

[33] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi, E. Flamand, F. K. Gürkaynak, and L. Benini, "Near-threshold risc-v core with dsp extensions for scalable iot endpoint devices," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 10, Oct 2017.

[34] A. Pullini, D. Rossi, G. Haugou, and L. Benini, "uDMA: An autonomous I/O subsystem for IoT end-nodes," in *2017 27th International Symposium on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, Sep. 2017, pp. 1–8.

[35] A. Rahimi, I. Loi, M. R. Kakoee, and L. Benini, "A Fully-Synthesizable Single-Cycle Interconnection Network for Shared-L1 Processor Clusters," in *2011 Design, Automation & Test in Europe*. IEEE, Mar. 2011, pp. 1–6.

[36] L. Cecconi, S. Smets, L. Benini, and M. Verhelst, "Optimal tiling strategy for memory bandwidth reduction for cnns," in *Advanced Concepts for Intelligent Vision Systems*, J. Blanc-Talon, R. Penne, W. Philips, D. Popescu, and P. Scheunders, Eds. Cham: Springer International Publishing, 2017.

[37] P. Kumar, S. Garg, A. Singh, S. Batra, N. Kumar, and I. You, "Mvo-based two-dimensional path planning scheme for providing quality of service in uav environment," *IEEE Internet of Things Journal*, 2018.

[38] C. Yin, Z. Xiao, X. Cao, X. Xi, P. Yang, and D. Wu, "Offline and online search: Uav multiobjective path planning under dynamic urban environment," *IEEE Internet of Things Journal*, vol. 5, no. 2, April 2018.

[39] D. Palossi, A. Marongiu, and L. Benini, "Ultra low-power visual odometry for nano-scale unmanned aerial vehicles," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2017*, March 2017.

[40] D. Falanga, A. Zanchettin, A. Simovic, J. Delmerico, and D. Scaramuzza, "Vision-based autonomous quadrotor landing on a moving platform," in *2017 IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, Oct 2017.

[41] F. Conti, D. Palossi, A. Marongiu, D. Rossi, and L. Benini, "Enabling the heterogeneous accelerator model on ultra-low power microcontroller platforms," in *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2016.

[42] L. Lai, N. Suda, and V. Chandra, "CMSIS-NN: Efficient Neural Network Kernels for Arm Cortex-M CPUs," *arXiv:1801.06601 [cs]*, Jan. 2018.

[43] J. Förster, M. Hamer, and R. D'Andrea, "System identification of the crazyflie 2.0 nano quadrocopter," B.S. thesis, 2015.

[44] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, sep 2012.