

# Fully On-board Low-Power Localization with Multizone Time-of-Flight Sensors on Nano-UAVs

Hanna Müller<sup>\*‡</sup>, Nicky Zimmerman<sup>†‡</sup>, Tommaso Polonelli<sup>\*</sup>  
Michele Magno<sup>\*</sup>, Jens Behley<sup>†</sup>, Cyrill Stachniss<sup>†</sup>, Luca Benini<sup>\*</sup>

<sup>\*</sup>Integrated Systems Laboratory / Center for Project-Based Learning - ETH Zürich, Switzerland

<sup>†</sup>Institute of Geodesy and Geoinformation - University of Bonn

<sup>‡</sup> The authors contributed equally to the paper.

Email: {hanmuell, tpolonelli, mmagno, lbenini}@ethz.ch, {nicky.zimmerman, jens.behley, cyrill.stachniss}@igg.uni-bonn.de

**Abstract**—Nano-size unmanned aerial vehicles (UAVs) hold enormous potential to perform autonomous operations in complex environments, such as inspection, monitoring or data collection. Moreover, their small size allows safe operation close to humans and agile flight. An important part of autonomous flight is localization, a computationally intensive task, especially on a nano-UAV that usually has strong constraints in sensing, processing and memory. This work presents a real-time localization approach with low-element-count multizone range sensors for resource-constrained nano-UAVs. The proposed approach is based on a novel miniature 64-zone time-of-flight sensor from STMicroelectronics and a RISC-V-based parallel ultra-low-power processor to enable accurate and low latency Monte Carlo Localization on-board. Experimental evaluation using a nano-UAV open platform demonstrated that the proposed solution is capable of localizing on a 31.2 m<sup>2</sup> map with 0.15 m accuracy and an above 95% success rate. The achieved accuracy is sufficient for localization in common indoor environments. We analyze trade-offs in using full and half-precision floating point numbers as well as a quantized map and evaluate the accuracy and memory footprint across the design space. Experimental evaluation shows that parallelizing the execution for 8 RISC-V cores brings a 7x speedup and allows us to execute the algorithm onboard in real-time with a latency of 0.2-30 ms (depending on the number of particles), while only increasing the overall drone power consumption by 3-7%. Finally, we provide an open-source implementation of our approach.

**Index Terms**—UAV, Localization, Autonomous navigation, nano-UAVs, Perception, ToF Array

## I. INTRODUCTION

Nano-size unmanned aerial vehicles (UAVs) fit in the palm of a hand, weigh only a few tens of grams, and therefore are agile, able to pass through narrow gaps, and safe to operate in proximity of humans [1]. Achieving autonomous flight is one of the most promising and difficult challenges for nano-UAVs, as they need to execute key tasks for autonomous robot navigation such as obstacle avoidance, localization, mapping, and path planning [2]. However, the restrictions in payload and power consumption pose severe challenges in reaching the autonomy of standard-size drones, as sensing and processing are strongly limited [3].

Focusing on the localization task, RTK-GPS is commonly used for outdoor scenarios [4]. Instead, in GPS-denied environments, such as indoor scenarios, onboard localization is challenging, especially on nano-UAVs, as they can only afford to spend around 10-15% for sensing and processing [5]. Enabling indoor localization on such constrained platforms is pushing researchers to design approaches that are lightweight




Fig. 1. Top: the nano-UAV localizing in the maze. Bottom: The ground truth pose and the predicted pose for 4096 particles. The initial prediction starts in the wrong maze and then converges to the correct pose when more observations become available. The color encodes the time, purple is the start, red the end.

and efficient — the most common ones are radio-based localization methods (mostly ultra-wideband (UWB) [6], [7]) or approaches that require off-board processing [8]. However, these approaches have major drawbacks, such as depending on external infrastructure and reliable communication to other nodes or a powerful base station with a reliable wireless connection for off-board processing [8].

A successful and popular localization methodology is map-based, which does not rely on external infrastructure [9]–[11]. To estimate the pose in a given occupancy grid map [9], it is necessary to have range measurements, which can be obtained by range-sensors such as LiDARs and depth/stereo cameras [12]. However, these sensors are large, power-hungry, and in the case of stereo cameras, also computationally expensive to process, thus making them unsuitable for most nano-UAVs [3]. A promising alternative is recently-emerged multizone time-of-flight (ToF) sensors, which are also suitable for nano-UAVs and have already proved to be robust and reliable for obstacle avoidance [13]. The drawback of these sensors is their low element-count, which is sufficient for obstacle avoidance but proves challenging for localization.

This paper proposes a global localization system based on Monte Carlo localization (MCL) and exploiting novel miniaturized low-power ToF sensors, not requiring previously installed infrastructure. The proposed algorithm is designed to run online on nano-UAV processors. In particular, the paper presents an algorithm where all the computations are performed onboard on a milliwatt power RISC-V parallel system on chip (SoC), avoiding wireless communication latency and potential issues such as range limits or security risks while significantly improving the system’s robustness. To exploit the hardware architecture of the RISC-V-based SoC, this paper presents a parallelized and memory-efficient implementation tailored for the system’s computational capabilities. Finally, the proposed approach has been experimentally evaluated in the field, and an open-source implementation will be available. The nano-UAV with all sensors and processors mounted is shown in Fig. 1. In our experiments, we show that our approach can (i) accurately localize a nano-UAV in a given map, using low element-count sensors without infrastructure, (ii) reduce memory consumption with quantization and lower precision floats without a significant drop in accuracy, (iii) reduce latency by a factor of 7 through parallelized implementation and localize on-board in real-time, and (iv) operate with low power consumption, where sensing and processing only consume below 7% of the overall power.

## II. RELATED WORK

Recent literature has demonstrated that sensing and processing on nano-UAVs are strongly limited; therefore, many previous works have proposed solutions for autonomous navigation that only rely on simple state estimation techniques such as an inertial measurement unit and odometry for localization [3], [13]–[15]. The major drawback of these approaches is their inability to compensate for drift and recover from accumulated errors [13]. Most drift-correction approaches use range measurements to anchors with known locations [7].

In indoor scenarios, the most commonly used technology is UWB, but approaches with Bluetooth or WiFi are also available. They all have disadvantages — they require line-of-sight between nodes, depend on pre-installed infrastructure [7], [16] or can only estimate relative position [6]. The resulting mean localization errors are often over 20 cm (22 cm in [7], 28 cm in [6]). In contrast to previous works, this paper focuses on an infrastructure-less approach to globally navigate indoors: a map-based localization approach using particle filters, which was not explored on nano-UAVs until now.

Localization in a given map is an essential capability of most autonomous robot systems, laying the foundation for more complex tasks such as planning and manipulation. Probabilistic approaches provide robust localization and include seminal works such as the extended Kalman filter (EKF) [17], Markov localization [18] and particle filters often referred to as MCL [19]. These approaches are suitable for localization using range sensors such as 2D LiDARs, sonars, and cameras. Until now, this approach was nearly infeasible on nano-UAVs, due to bulky power-hungry sensors, and high computational demands, which are hard to satisfy on embedded systems.

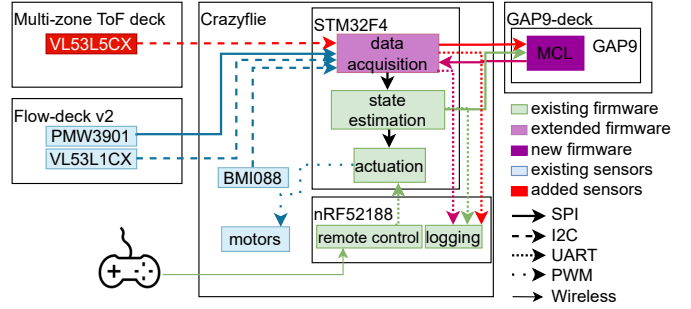


Fig. 2. System overview, showing the hardware connections and data dependencies between the Crazyflie and the three mounted extension decks.

For both the sensing and the processing challenges, promising hardware recently emerged. Although introduced on the market only recently, lightweight multizone ToF sensors are already working well for obstacle avoidance [13]. As for powerful and energy-efficient SoCs for processing, SoCs of the parallel ultra-low power (PULP) family have been employed on drones before. The GAP8 SoC is utilized for corridor [15] or person following [14]. These approaches use deep learning with quantized models but do not venture into float-heavy tasks such as particle-filter localization. A novel SoC, GAP9<sup>1</sup> was recently released, which with 0.33 mW per giga operation (GOP) is an order of magnitude more power efficient than GAP8 and most importantly, features increased memory and floating point support.

In this work, we combine a miniature multizone ToF sensor with a novel processor to enable on-board infrastructure-less localization in indoor environments with an accuracy that surpasses the state of the art of localization in nano-UAVs with UWB [6], [7].

## III. SYSTEM ARCHITECTURE

This section presents a complete description of the proposed infrastructure-less localization system for nano-UAVs; from the hardware design to the algorithm implementation, and the in-field evaluation. We used the commercially available Crazyflie 2.1 platform from Bitcraze, extending its functionality with custom expansion boards featuring new sensors and processors, namely the VL53L5CX from STMicroelectronics and GAP9 SoC from GreenWaves technologies as main processing unit. All used components are commercially available, and our hardware design as well as the proposed embedded algorithm implementation will be released as open-source<sup>2</sup>. Fig. 2 presents our design, composed of the Crazyflie’s integrated hardware and software parts (blue for hardware, green for software) and our own additions (red for hardware, purple for software).

### A. Hardware: Crazyflie and Extension Boards

The Crazyflie 2.1 is a commercially available open software/hardware nano-UAV. In this work, we use its inertial measurement unit (IMU), radio communication (using an nRF51822, solely to log data and steer the drone) and the main processor, an STM32F405 (168 MHz, 192 kBRAM), which is

<sup>1</sup><https://greenwaves-technologies.com>

<sup>2</sup>[https://github.com/ETH-PBL/Matrix\\_ToF\\_Drones](https://github.com/ETH-PBL/Matrix_ToF_Drones)

responsible for sensor readout, state estimation, and real-time control.

1) *Flow-deck v2*: a commercially available deck featuring a downward-facing optical flow sensor and 1D ToF sensor for odometry measurements. Those sensors improve the internal position estimate provided by the Crazyflie firmware through an extended Kalman Filter.

2) *Multizone-ToF-deck*: a custom deck featuring up to two VL53L5CX sensors (forward and backward facing), which can provide a matrix of either 8x8 or 4x4 pixels at maximally 15 Hz or 60 Hz respectively. For each zone, it provides a distance measurement coupled with an error flag, which gets raised when out-of-range measurements or interference are detected.

3) *GAP9-deck*: a custom deck, featuring GAP9. In our application, the multizone ToF sensor measurements are acquired by the STM32 via an I2C bus and then, together with the state estimation, sent via SPI to the GAP9 SoC.

### B. Processor: GAP9

GAP9's architecture is based on the open-source SoC Vega [20] and features 10 RISC-V instruction set architecture-based cores, extended with custom instructions. The compute cluster, featuring 9 cores, one for orchestration and 8 workers, delivers programmable compute power at extreme energy efficiency. GAP9 features 128 kB of shared L1 memory. The fabric controller (FC) has access to various peripherals and features 64 kB RAM, 1.5 MB interleaved memory (L2) and even 2 MB flash. The architecture employs adjustable dynamic frequency and voltage domains, allowing us to tune the energy consumption to the exact requirements at a particular point in time. At peak performance, the cores run at 400 MHz on both the cluster and the FC.

### C. Algorithm: Monte Carlo Localization

This section first provides an overview of the Monte Carlo localization algorithm and then explains our adaptations for running it onboard and in real-time on GAP9.

1) *Algorithm overview*: MCL, most commonly used with occupancy grid maps [9], [10], [19], is a particle filter-based approach for estimating the posterior of the robot pose  $x_t$  given a map  $m$ , sensor readings  $z_t$  and odometry inputs  $u_t$ . As the nano-UAV flies at a fixed height and localizes in a 2D grid map, the nano-UAV's state  $x_t$  is defined by the 2D coordinates  $(x, y)^\top$  and the yaw-angle orientation  $\theta \in [0, 2\pi)$ . MCL has 3 main components: the prediction step using the motion model, the correction step using the observation model, and resampling (Fig. 3). When odometry is available, we sample from the proposal distribution  $p(x_t | x_{t-1}, u_t)$  with odometry noise  $\sigma_{\text{odom}} \in \mathbb{R}^3$ . The observation model describes the probability of observing  $z_t$  from pose  $x_t$  given a map  $m$ , where each observation  $z_t$  is composed of  $K$  elements  $z_t^k$ . As we are using a range sensor with an occupancy grid map, we chose the beam end point model [21] as our observation model, as shown in Eq. 1,

$$p(z_t^k | x_t, m) = \frac{1}{\sqrt{2\pi}\sigma_{\text{obs}}} \exp\left(-\frac{\text{EDT}(\hat{z}_t^k)^2}{2\sigma_{\text{obs}}^2}\right), \quad (1)$$

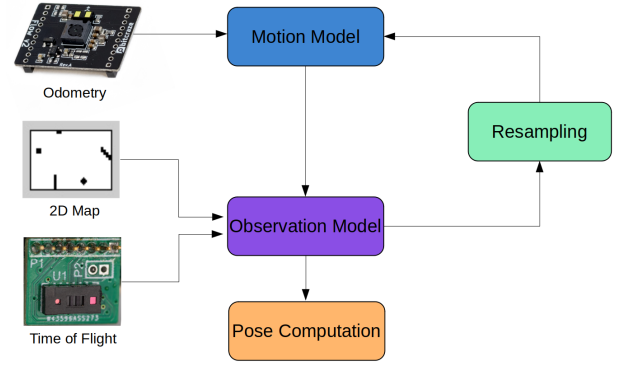


Fig. 3. The logic flow of the Monte Carlo localization algorithm.

where  $\hat{z}_t^k$  is the end point of the ToF beam in the occupancy grid map  $m$ . We estimate the distance between each cell in the occupancy grid map to an obstacle (occupied cell) using the Euclidean distance transform (EDT) [22]. The EDT is truncated at  $r_{\text{max}}$ , a predefined maximal range. In addition to the three main components (Fig. 3), we also include a fourth step, pose computation, where the pose estimation is computed as the weighted average over all particles.

2) *Adaptions for Lightweight and Parallel Embedded Implementation*: The two main constraints in the nano-UAVs hardware are memory and time - we need to use both resources efficiently to enable MCL on-board in real-time. Our implementation of MCL is asynchronous - the motion model is sampled when odometry is available, and the particles are re-weighted when new range measurements arrive. We only consider new observations if the drone moves more than  $d_{xy}$  or rotates more than  $d_\theta$ . However, we configured our sampling rates for the motion and observation update to be the same, limited by the 15 Hz maximum update rate of the ToF sensor.

The motion model, observation model and pose computation can be parallelized exploiting the GAP9 cluster by distributing the particles among the cores. The resampling step can also be parallelized; however, as it depends on all weights, we can not plan the workload distribution optimally.

The first step is weight normalization, which involves computing the sum and dividing by it - we can parallelize this step by splitting the particles evenly to all cores. We also store the partial sums, as we can use them to parallelize the main resampling step, drawing the new particles.

For drawing the new particles, we use a systematic resampling algorithm [23], which we explain with the analogy of a wheel, as shown in Fig. 4. We draw one random number, corresponding to the first *arrow* in the wheel, with the other  $N-1$  *arrows* being fixed at regular intervals from that randomly picked *arrow*. The colors show how we distribute the drawing of the next particles to the cores. The current particle weights are represented by the cone area they occupy. As we know the partial sums computed by all cores, we can directly use them to calculate which core will resample how many particles and which. In Fig. 4, the colored *arrows* represent the new particles picked by the corresponding cores.

The main components of MCL using memory space are the particles and the map. The occupancy map requires 2 bits






Fig. 4. Parallelizing the resampling wheel: Each color represents a core, the current particles are distributed evenly (here two per core) and then the new particles are chosen according to where the arrows of the resampling wheel point.




Fig. 5. Left: The drone maze. Right: The occupancy grid map used for the localization task, where the highlighted part corresponds to the drone maze

per cell (to represent the 3 possible states), to simplify the memory access we store it as 1 byte per cell. However, we also precompute the EDT values for each cell, leading to an additional floating point number being saved for every cell. To decrease the memory usage, we compare three possibilities: 32-bit floating point numbers, 16-bit floating point numbers, and quantized 8-bit unsigned integer values. For the particles, we need four numbers each, one for x position, y position, yaw angle, and weight. With a 32-bit floating point representation, this leads to 16 bytes per particle. However, as we are double-buffering the particles for executing the resampling step, we need 32 bytes per particle for the 32-bit representation, and 16 bytes for more memory-efficient 16-bit representations.

#### IV. EXPERIMENTAL EVALUATION

This section presents our experiments to demonstrate the effective capabilities of our framework. The results of our experiments also support our key claims, our system can: (i) accurately localize a nano-UAV in a given map, using low element-count sensors without infrastructure, (ii) reduce memory consumption with quantization and lower precision floats without a significant drop in accuracy, (iii) reduce latency by a factor of 7 through parallelized implementation and localize on-board in real-time, and (iv) operate with low power consumption, where sensing and processing only consume below 7% of the overall power.

##### A. Experimental Setup

To evaluate the performance of our approach, we recorded a dataset, including 6 sequences, while flying the drone in our “drone maze”(Fig. 5). The recordings include ToF measurements from two sensors, internal state estimation based on the FlowDeck’s optical flow and ground truth pose. The ground

truth is extracted using a motion capture system, Vicon Vero 2.2, with six cameras positioned around the maze, covering an area of 16 m<sup>2</sup>. The map acquisition is made by manually measuring the maze objects, which introduces some inaccuracy and increases the localization challenge. For all experiments, we use a map resolution of 0.05 m by 0.05 m per cell. The algorithm parameters are  $\sigma_{odom} = (0.1 \text{ m}, 0.1 \text{ m}, 0.1 \text{ rad})$ ,  $\sigma_{obs} = 2.0$ ,  $r_{max} = 1.5 \text{ m}$ ,  $d_{xy} = 0.1 \text{ m}$  and  $d_{\theta} = 0.1 \text{ rad}$ . To challenge localization even further, we extended the map with three artificial mazes to a total of 31.2 m<sup>2</sup> of structured area.

Three aspects were considered — localization accuracy, runtime performance, and system power consumption. To evaluate the accuracy, we take into account 3 metrics: the success rate, the time to convergence, and absolute trajectory error (ATE) after convergence. We define convergence to occur when the estimated pose is close to the ground truth pose, within a distance of  $(36^\circ/0.2 \text{ m})$ . The localization is counted as successful if the pose tracking remains reliable from convergence until the end of the sequence, meaning that the ATE does not exceed 1 m.

##### B. Robust Localization with Sparse ToF

To support the claim that we are able to accurately localize with low element-count sensors, we evaluated our approach on six sequences. For each of them, we repeated the localization experiments with six different random seeds to verify robustness. All of our experiments were executed using data from both, front and rear ToF sensors unless explicitly mentioned.

Infrastructure-dependent localization approaches [6], [7] for nano-UAVs, which were evaluated in similar environments to ours, have achieved mean localization errors of 0.22 m and 0.28 m. As can be seen in Fig. 6 and Fig. 7, our approach can localize with 0.15 m accuracy and achieves above 95% success rate with a sufficient number of particles, outperforming the existing approaches. An illustration of successful localization can be seen in Fig. 1. Our experiments show that our approach is robust with respect to the number of particles, providing ATE of less than 0.2 m for a large range of particle numbers.

We perform additional experiments to confirm the contribution of the second (rear) ToF sensor. We compare the performance with a pair of ToF sensors, referred to as fp32, to that of a single ToF sensor which we refer to as fp321tof. For both configurations, the accuracy is calculated for a 32-bit representation of the floats in the EDT and a particle’s weight and pose. As shown in Fig. 7, the success rate when using two ToF sensors is significantly higher, and the accuracy is also improved slightly (Fig. 6). The convergence is slower when using only 1 ToF sensor, as illustrated in Fig. 8.

##### C. Memory Usage

To support our claim, that we can reduce memory consumption with quantization and lower precision floats without significant loss of accuracy, we compare our optimized implementation against a full-precision implementation. We compare the localization performance for three implementations of our approach. The first implementation uses a 32-bit representation for the floats in the EDT and particles (fp32). The second

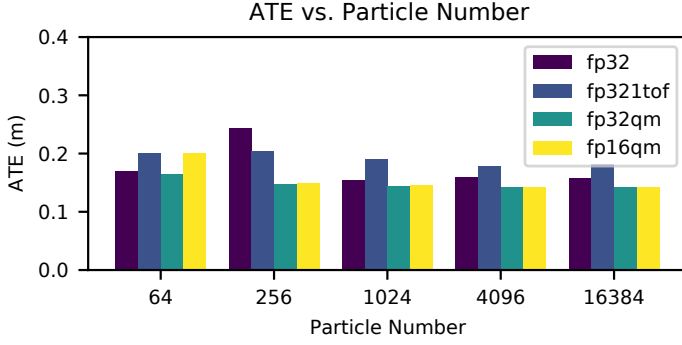


Fig. 6. The absolute trajectory error (ATE) computed over all sequences.




Fig. 7. The success rate in % of successful convergence across all sequences.

implementation, (fp32qm) uses a quantized EDT whose values are 8-bit unsigned integers. The third implementation (fp16qm) is a quantized EDT and a 16-bit float representation for a particle's weight and pose. As shown in Fig. 6 and Fig. 7, the quantized implementations maintain a high success rate and provide accuracy that surpasses the full precision implementation. As reported in Fig. 8, the convergence time is improved for the optimized implementations. We speculate that the quantization accelerates the rate at which weak particles are eliminated in the resampling step, leading to faster convergence and overall better performance.

As we do not observe a significant accuracy loss when going down to 8-bit quantized EDT values, we can reduce our memory requirements for the map from 5 bytes per cell to 2 bytes per cell. For the particles, we saw that using half-precision instead of full-precision representation does not result in a significant accuracy drop, meaning we can reduce the needed memory by a factor of 2. We visualize the memory savings in Fig. 9, where we show how many particles and square meters can be stored on the GAP9 in L1 (blue full precision, yellow quantized/FP16) respectively L2 (red full precision, green quantized/FP16) memory.

#### D. Real-time Performance

In support of our third claim, that we can reduce latency through parallelized implementation and accurately localize onboard in real-time, we compared our optimized implementation against a naive approach, and measured the latency improvement achieved by parallelization.

In Tab. I, we report the execution times per particle for the previously described steps executed sequentially on one core




Fig. 8. The probability of converging over time, across all sequences, computed for 4096 particles.




Fig. 9. The trade-off between the number of particles and the map size at a resolution of 0.05 m/cell for both L1 and L2 memory.

and in parallel on 8 cores (while using the ninth core of the cluster for orchestration). In Fig. 10, we show the achieved speedup for different numbers of particles. As expected, the resampling step scales the worst - however, for high numbers of particles, we can reach more than 5x speedup even for this step. We also report the total speedup achieved in orange - we can observe it improving until a factor of 7 at high numbers of particles. Note that the total execution time is around 40  $\mu$ s higher than the sum of the four tasks, independent of the numbers of particles and multicore usage, which are used for preprocessing the sensor data and transferring information to the tasks.

#### E. Power Consumption

We performed power measurements to support our fourth claim that our approach can operate with low power consumption. Firstly, we analyzed the power consumption for processing on GAP9. We analyze two working points, 1,024 particles, which can still fit in L1, and 16,384 particles, the maximum number of particles we considered. We measured the power consumption at the maximum possible frequency and the minimal frequency at which we can still operate in real-time, meaning processing in less than 67 ms. In Tab. II, we report our results, consisting of the average power consumption and the execution times at the respective frequencies.

Combined with the power needed for the sensors, 320 mW each, and the remaining Crazyflie electronics (besides motors), 280 mW, all sensing and processing power sums to 981 mW, i.e., around 7% of the overall power consumption of the drone.

TABLE I  
EXECUTION TIMES BY COMPONENTS FOR DIFFERENT NUMBERS OF PARTICLES

|             | execution time per particle 1 core / 8 cores in ns, GAP9@400 MHz |           |           |                    |                     |
|-------------|--|-----------|-----------|--------------------|---------------------|
| Particles   | 64   | 256       | 1,024     | 4,096 <sup>a</sup> | 16,384 <sup>a</sup> |
| Observation | 8531/1412  | 8484/1313 | 8518/1283 | 8649/1294          | 8704/1295           |
| Motion      | 2828/500   | 2715/391  | 2689/357  | 3002/390           | 2985/386            |
| Resampling  | 313/250  | 191/121   | 161/84    | 558/108            | 556/104             |
| Pose Comp.  | 750/234  | 633/117   | 604/86    | 777/101            | 775/99              |

<sup>a</sup>Particles stored in L2.

TABLE II  
AVERAGE POWER CONSUMPTION OF THE MCL ALGORITHM ON GAP9 AT DIFFERENT OPERATING POINTS.

|  | Avg. power consumption | Execution time |
|--|------------------------|----------------|
| GAP9@400 MHz/1,024 particles               | 61 mW                  | 1.901 ms       |
| GAP9@12 MHz/1,024 particles                | 13 mW                  | 59.898 ms      |
| GAP9@400 MHz/16,384 particles <sup>a</sup> | 61 mW                  | 30.880 ms      |
| GAP9@200 MHz/16,384 particles <sup>a</sup> | 38 mW                  | 61.524 ms      |

<sup>a</sup>Particles stored in L2.

## V. CONCLUSION

This paper presents a system that includes a nano-UAV, a low-power multi-core processor, low-element-count sensors, and a hardware-specific highly-optimized MCL implementation for indoor localization. Our experiments show that this system enables nano-UAVs to accurately localize in a given map despite the sparsity of the sensors. The paper also demonstrates the benefit of the proposed parallelized and memory-efficient MCL, which runs at 15 Hz, efficiently using the available compute resources to operate onboard, in real-time, without significant loss of accuracy. Even in the most powerful configuration, sensing and processing increase the drone's power consumption by only 7%. Future works will extend the proposed system to applications such as path planning and exploration.

## ACKNOWLEDGMENT

The authors thank V. Niculescu for the picture of the drone. Moreover, this work has been partially funded by the Autonomous Robotics Research Center of the UAE Technology Innovation Institute.

## REFERENCES

- [1] E. Tsykunov, L. Labazanova, A. Tleugazy, and D. Tsetserukou, "Swarm-Touch: Tactile Interaction of Human with Impedance Controlled Swarm of Nano-Quadrotors," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2018.
- [2] A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, and D. Scaramuzza, "Deep Drone Racing: From Simulation to Reality with Domain Randomization," *IEEE Transactions on Robotics*, vol. 36, no. 1, pp. 1–14, 2019.
- [3] K. McGuire, C. De Wagter, K. Tuyls, H. Kappen, and G. C. de Croon, "Minimal Navigation Solution for a Swarm of Tiny Flying Robots to Explore an Unknown Environment," *Science Robotics*, vol. 4, no. 35, 2019.
- [4] J. Schneider, C. Eling, L. Klingbeil, H. Kuhlmann, W. Förstner, and C. Stachniss, "Fast and Effective Online Pose Estimation and Mapping for UAVs," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2016, pp. 4784–4791.
- [5] R. Wood, B. Finio, M. Karpelson, K. Ma, N. Pérez-Arancibia, P. Sreetharan, H. Tanaka, and J. Whitney, "Progress on 'Pico' Air Vehicles," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1292–1302, 2012.
- [6] S. Van Der Helm, M. Coppola, K. N. McGuire, and G. C. de Croon, "On-Board Range-Based Relative Localization for Micro Air Vehicles in Indoor Leader-Follower Flight," *Autonomous Robots*, vol. 44, no. 3, pp. 415–441, 2020.




Fig. 10. The speedup on GAP9 by parallelizing the correction, motion update, resampling, and prediction computation for different numbers of particles.

- [7] V. Niculescu, D. Palossi, M. Magno, and L. Benini, "Energy-efficient, Precise UWB-based 3-D Localization of Sensor Nodes with a Nano-UAV," *IEEE Internet of Things Journal*, 2022.
- [8] B. Şimşek and H. Ş. Bilge, "A Novel Motion Blur Resistant vSLAM Framework for Micro/Nano-UAVs," *Drones*, vol. 5, no. 4, 2021.
- [9] H. P. Moravec, "Sensor Fusion in Certainty Grids for Mobile Robots," in *Sensor Devices and Systems for Robotics (SDSR)*, 1989.
- [10] D. Meyer-Delius, M. Beinhofner, and W. Burgard, "Occupancy Grid Models for Robot Mapping in Changing Environments," *Proc. of the Conference on Advancements of Artificial Intelligence (AAAI)*, vol. 26, pp. 2024–2030, Sep. 2021.
- [11] N. Zimmerman, L. Wiesmann, T. Guadagnino, T. Låbe, J. Behley, and C. Stachniss, "Robust Onboard Localization in Changing Environments Exploiting Text Spotting," *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2022.
- [12] C. Debeunne and D. Vivet, "A Review of Visual-LiDAR Fusion based Simultaneous Localization and Mapping," *Sensors*, vol. 20, no. 7, 2020.
- [13] H. Müller, V. Niculescu, T. Polonelli, M. Magno, and L. Benini, "Robust and Efficient Depth-based Obstacle Avoidance for Autonomous Miniaturized UAVs," *arXiv preprint*, 2022. [Online]. Available: <https://arxiv.org/abs/2208.12624>
- [14] D. Palossi, N. Zimmerman, A. Burrello, F. Conti, H. Müller, L. M. Gambardella, L. Benini, A. Giusti, and J. Guzzi, "Fully Onboard AI-Powered Human-Drone Pose Estimation on Ultralow-Power Autonomous Flying Nano-UAVs," *IEEE Internet of Things Journal*, vol. 9, no. 3, pp. 1913–1929, 2022.
- [15] V. Niculescu, L. Lamberti, F. Conti, L. Benini, and D. Palossi, "Improving Autonomous Nano-Drones Performance via Automated End-to-End Optimization and Deployment of DNNs," *IEEE Journal on Emerging and Selected Topics in Circuits and Sys.*, vol. 11, no. 4, pp. 548–562, 2021.
- [16] S. Ito, F. Endres, M. Kuderer, G. Tipaldi, C. Stachniss, and W. Burgard, "W-RGB-D: Floor-Plan-Based Indoor Global Localization Using a Depth Camera and WiFi," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2014.
- [17] J. Leonard and H. Durrant-Whyte, "Mobile Robot Localization by Tracking Geometric Beacons," *IEEE Trans. on Robotics and Automation*, vol. 7, no. 3, pp. 376–382, 1991.
- [18] D. Fox, W. Burgard, and S. Thrun, "Markov Localization for Mobile Robots in Dynamic Environments," *Journal of Artificial Intelligence Research (JAIR)*, vol. 11, pp. 391–427, 1999.
- [19] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo Localization for Mobile Robots," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 1999.
- [20] D. Rossi, F. Conti, M. Eggiman, A. D. Mauro, G. Tagliavini, S. Mach, M. Guermandi, A. Pullini, I. Loi, J. Chen, E. Flamand, and L. Benini, "Vega: A Ten-Core SoC for IoT Endnodes With DNN Acceleration and Cognitive Wake-Up From MRAM-Based State-Retentive Sleep Mode," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 1, pp. 127–139, 2022.
- [21] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, 2005.
- [22] P. F. Felzenszwalb and D. P. Huttenlocher, "Distance Transforms of Sampled Functions," *Theory of Computing*, vol. 8, no. 1, pp. 415–428, 2012.
- [23] R. Douc and O. Cappe, "Comparison of Resampling Schemes for Particle Filtering," in *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*, 2005., 2005, pp. 64–69.