Poster Abstract: Data-Parallel Versus Task-Parallel Swarms for Small Unmanned Aerial Systems

Jayson Boubin, Zichen Zhang, John Chumley and Christopher Stewart
Department of Computer Science and Engineering, The Ohio State University

I. Introduction

Small Unmanned Aerial Systems (sUAS) are small IoT devices (weighing less than 55 lbs) that fly without human pilots aboard. sUAS use sensors and auto pilot software to hover, fly and land autonomously. Without humans aboard, lightweight sUAS can fly low, only meters above the ground, to capture detailed, high resolution images. These images can inform agricultural crop management by detecting crop health issues, such as nitrogen deficiency [6], [9], overcrowding [5], and soybean defoliation [10]. Other uses for sUAS include search and rescue [3], inspection for buildings and bridges [8], and photography. A *mission* is the unit of workload for sUAS. It includes aerial maneuvers, sensing, and computation required to complete a task. For example, in digital agriculture [4], a sUAS tasked with generating a report on crop health may conduct a mission that includes flying to a targeted area, capturing images of the crop, and applying machine learning to translate the images into a crop health characterization.

sUAS support multiple operating modes for mission execution. In remote-control mode, a human operator sends flight commands to the aircraft throughout a mission via a remote controller or smartphone. The human operator decides where to fly and when the mission is complete. In waypoint mode, an operator tells the sUAS to fly to preset waypoints. In this mode, the operator still tells the drone where to fly but the aerial portion of the mission is complete after visiting all preset waypoints. Finally, in the autonomous-flight operating mode [3], [7], software running aboard the aircraft, on edge devices, or in the cloud decides where to fly during mission execution, removing the human operator from the loop. This operating mode can reduce costs substantially by (1) simplifying the role of the operator, (2) avoiding human errors, and (3) efficiently provisioning compute resources [2]. This poster paper explores research challenges when sUAS use the autonomous-flight operating mode.

Missions that cover large geographic regions present challenges for sUAS. To keep the aircraft light, sUAS carry small batteries that are drained quickly (25 minutes) during flight. To explore large regions for more than 25 minutes, sUAS operators must recharge and exchange batteries. This process is manual, time consuming, and costly. One solution is to deploy multiple sUAS in parallel, i.e., in *swarms*. In autonomous flight mode, swarm members can carry out missions independently, increasing geographic throughput in proportion to swarm size.

For this short paper, we deployed sUAS swarms for agricultural crop management. We compared two approaches for swarm design: task parallel versus data parallel. The task parallel approach splits a mission into smaller tasks. Swarm members carry out each task independently without coordination. That is, swarm members are provided all of the data needed to execute the task beforehand. In the data parallel approach, swarm members share images observed during each task, aggregate them, and use the data to update autonomous flight plans. Each swarm member processes a partition of the data represented by the whole field. The data parallel approach presents opportunities and risks. On one hand, sharing data can improve autonomous flight paths, allowing swarm members to complete tasks efficiently by visiting only waypoints that contribute to final reports. On the other hand, data parallelism increases the computational load, i.e., the cost of updating autonomous flight paths may outweigh the benefits.

II. Design and Deployment

We used 3 sUAS (DJI Mavic) to scout a 30-hectare field used to grow soybeans. The task was mapping the prevalence of soybean defoliation, a warning sign for insect infestation. If undetected, insect infestation can degrade crop yield rapidly and significantly. Using task parallelism, we could scout the whole field in 8 hours. Each sUAS autonomously maps a 0.4-hectare geographic region before its battery was recharged. For this study, we repeated the mission over 150 times across a variety of runtime conditions. Our experiments spanned three weeks in the field.

Computational Workflow: Autonomous flight mode uses edge gateways to compute the next flight action [3] given the most recent sensed image. Precisely, the data pipeline was as follows. After takeoff, the sUAS flew to an initial waypoint, hovered for 5 seconds, and captured a 4K high-definition image. The image was transmitted to an edge gateway via 5 Ghz radio through the remote control. Then, the edge gateway used a convolutional neural network to classify defoliation in the image as 'significant' or 'insignificant/none'. Reinforcement learning software on the edge gateway used recent observations of defoliation to choose the next waypoint to visit. This waypoint was then sent to the sUAS.

Data-Parallel Swarms Our task-parallel swarm design simply divides a 30-hectare mission into 75 independent 0.4-hectare geographic regions. Each sUAS maps 25 regions without sharing or reusing data between tasks. In contrast, data-parallel swarm design shares images after each mission, updating reinforcement learning software on edge gateways.

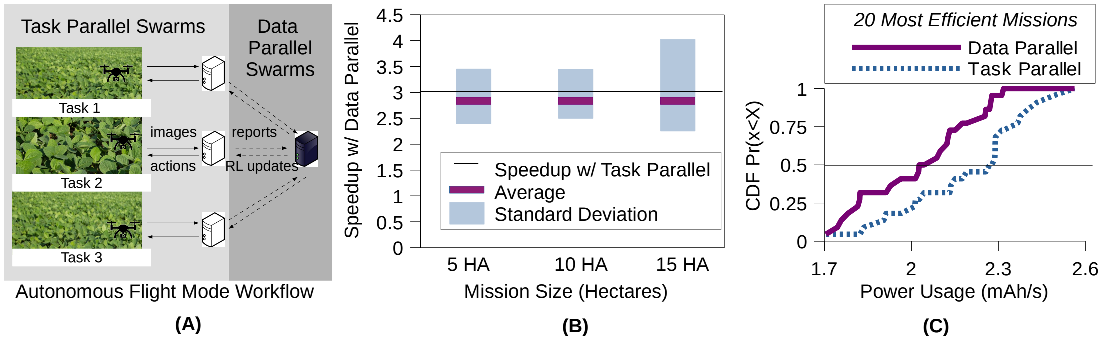

Fig. 1. (A) Work and data flow for sUAS to map soybean defoliation. (B) On average, data-parallel swarms are 6% less efficient than task-parallel swarms. (C) The most efficient data-parallel missions outpace task-parallel missions.

We connected edge gateways for each drone with a 1 Gb/s router and connected another edge device. We used multi-agent reinforcement learning to update reinforcement learning policies based on the observations of recent tasks within the mission.

We hypothesized that data-parallel swarms would execute missions more efficiently than task-parallel swarms because of dynamic updates to reinforcement learning.

Deployment: In this paper, we implement a convolutional neural network (CNN) model, called DefoNet [10], to predict the leaf defoliation conditions in a soybean field. DefoNet is a binary CNN that classifies images of soybeans as either significantly defoliated or healthy. DefoNet uses eight convolutional layers partitioned into 3 blocks with activation and pooling between blocks. In testing, DefoNet correctly classifies 92% of expertly labeled soybean images. In our deployment, DefoNet was deployed using the Fleet Computer [1], an adaptive edge deployment model for autonomous systems. Using these technologies, we deployed our UAV swarm to analyze soybean fields in flight.

III. Results

We examined power efficiency for battery drain during mission execution. Specifically, our metric of merit is speedup from parallelism. Aggregate energy consumption divided by geographic coverage times mission duration.

$$\text{Speedup} = \frac{\text{Energy}}{\text{RegionSize} \times \text{Duration}} \tag{1}$$

Recall, we report averages over many runs in varying environmental conditions.

Figure 1(B) reports our first finding: our data parallel approach was 6% less efficient than task parallel on average. We attribute this to additional time spent idly hovering while data from missions were transferred to the edge gateway to update reinforcement learning models. RL model updates are computationally costly to update, leading to increased wait times for inference.

Figure 1(B) also reports significant variance in power efficiency of the data parallel approach. In all cases, a single standard deviation can cause a significant increase speedup.

Figure 1(C) reports the primary finding for our preliminary work: **Data-parallel swarms, under the right deployment conditions, can be 2X more efficient than task-parallel swarms.** This plot examines the 20 most efficient settings for each approach. We observe that data-parallel swarms can achieve power efficiency unattainable by task-parallel swarms. This gives rise to an important opportunity for computer systems research. Can we adaptively deploy data-parallel swarms to improve efficiency when the environment warrants? Alternatively, can we tailor the environment of sUAS missions to exploit the potential for data-parallel swarms?

References

[1] J. Boubin, C. Burley, P. Han, B. Li, B. Porter, and C. Stewart. Programming and deployment of autonomous swarms using multi-agent reinforcement learning. *arXiv preprint arXiv:2105.10605*, 2021.

[2] J. Boubin, J. Chumley, C. Stewart, and S. Khanal. Autonomic computing challenges in fully autonomous precision agriculture. In *2019 IEEE International Conference on Autonomic Computing (ICAC)*. IEEE, 2019.

[3] J. G. Boubin, N. T. Babu, C. Stewart, J. Chumley, and S. Zhang. Managing edge resources for fully autonomous aerial systems. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pages 74–87. ACM, 2019.

[4] R. Chandra and S. Collis. Digital agriculture for small-scale producers: challenges and opportunities. *Communications of the ACM*, 64(12):75–84, 2021.

[5] Y. Lan, K. Huang, C. Yang, L. Lei, J. Ye, J. Zhang, W. Zeng, Y. Zhang, and J. Deng. Real-time identification of rice weeds by uav low-altitude remote sensing based on improved semantic segmentation model. *Remote Sensing*, 13(21):4370, 2021.

[6] M. Qiao, X. He, X. Cheng, P. Li, H. Luo, L. Zhang, and Z. Tian. Crop yield prediction from multi-spectral, multi-temporal remotely sensed imagery using recurrent 3d convolutional neural networks. *International Journal of Applied Earth Observation and Geoinformation*, 102:102436, 2021.

[7] J. L. Sanchez-Lopez, R. A. S. Fernández, H. Bavle, C. Sampedro, M. Molina, J. Pestana, and P. Campoy. Aerostack: An architecture and open-source software framework for aerial robotics. In *International Conference on Unmanned Aircraft Systems*, 2016.

[8] A. Van't Hof and J. Nieh. Androne: Virtual drone computing in the cloud. In *Proceedings of the Fourteenth EuroSys Conference 2019*, pages 1–16, 2019.

[9] Z. Zhang, J. Boubin, C. Stewart, and S. Khanal. Whole-field reinforcement learning: A fully autonomous aerial scouting method for precision agriculture. *Sensors*, 20(22):6585, 2020.

[10] Z. Zhang, S. Khanal, A. Raudenbush, K. Tilmon, and C. Stewart. Assessing the efficacy of machine learning techniques to characterize soybean defoliation from unmanned aerial vehicles. *Computer and Electronics in Agriculture*, 2022.