

# LoCUS: A Multi-Robot Loss-Tolerant Algorithm for Surveying Volcanic Plumes

John Ericksen<sup>\*1</sup>, Abhinav Aggarwal<sup>1</sup>, G. Matthew Fricke<sup>1,2</sup> and Melanie E. Moses<sup>1,3,4</sup>

<sup>1</sup>Department of Computer Science, <sup>2</sup>Center for Advanced Research Computing, <sup>3</sup> Department of Biology,

University of New Mexico, Albuquerque, NM, USA, <sup>4</sup>Santa Fe Institute, Santa Fe, NM, USA

Email: \*johncarl@unm.edu

**Abstract**—Measurement of volcanic CO<sub>2</sub> flux by a drone swarm poses special challenges. Drones must be able to follow gas concentration gradients while tolerating frequent drone loss. We present the LoCUS algorithm as a solution to this problem and prove its robustness. LoCUS relies on swarm coordination and self-healing to solve the task. As a point of contrast we also implement the MoBS algorithm, derived from previously published work, which allows drones to solve the task independently. We compare the effectiveness of these algorithms using drone simulations, and find that LoCUS provides a reliable and efficient solution to the volcano survey problem. Further, the novel data-structures and algorithms underpinning LoCUS have application in other areas of fault-tolerant algorithm research.

**Index Terms**—Autonomous Drones, Fault-Tolerance, Surveying, Self-Stabilizing Systems, Consensus

## I. INTRODUCTION

More than 10% of the world's population live in the destructive zone of volcanoes, and a quarter of a million people have perished in volcanic eruptions in the last 500 years [1]. Volcanoes emit unknown amounts of CO<sub>2</sub> and other climate changing gasses, but only 10 of the approximately 300 currently active volcanoes are characterised by long-term datasets that enable any assessment of temporal CO<sub>2</sub> variability [2]. Measuring volcanic CO<sub>2</sub> flux would enable predictions of eruptions, minimizing loss of life and economic impact, as well as informing our understanding of greenhouse gas-driven climate change.

Satellite remote sensing of CO<sub>2</sub> is infeasible, so sampling is currently performed by ground based sensors or aerial surveys with piloted aircraft [3]. These techniques are costly, dangerous, and produce temporally and spatially coarse measurements. UAV present an emerging solution [4] that reduces risk to volcanologists and has the potential to markedly increase sampling resolution within volcano plumes.

An international team of research universities recently demonstrated that UAV can feasibly sample CO<sub>2</sub> from an active volcano in Papua New Guinea [5]. We developed the dragonfly drone for this task. The dragonfly is capable of measuring CO<sub>2</sub> in real time and has a flight duration of 1 h. However, drone loss was very common. Sudden and violent thermal updraughts, acidic plumes, and rugged cliffs were some of the many conditions that destroyed UAV. Further, the remoteness of many survey sites and battery life restrictions necessitate brief missions with small swarms. These hazardous and difficult conditions motivate the need for reliable perfor-

mance and surveillance algorithms that maximize the chance of completing the CO<sub>2</sub> surveillance task even with the loss of drones, short flight times, and small swarm sizes.

A key task for volcano surveillance is to locate the maximum CO<sub>2</sub> flux (max flux) in a dynamic gas plume. We propose the LoCUS algorithm to maintain a spatially dispersed swarm of drones that can simultaneously measure CO<sub>2</sub> concentrations at different locations and communicate those measurements across the entire swarm. We use deductive arguments to prove the loss-tolerance properties of LoCUS, and we test its performance and fault tolerance in simulations. In particular we show LoCUS guarantees that failed drones are replaced within flight-time proportional to the square root of swarm size, while preserving the swarm symmetry essential to efficient gradient following.

We hypothesise that maintaining a dispersed team of robots that can simultaneously measure CO<sub>2</sub> at different spatial locations will provide a better estimate of the CO<sub>2</sub> gradient, allowing fast navigation to the CO<sub>2</sub> source. We further hypothesise that the benefit of spatially dispersed measurements outweighs the increased complexity resulting from coordination and self-healing. To test these hypotheses, we develop an alternative approach which allows multiple UAV to independently search for the maximum CO<sub>2</sub> flux. We compare LoCUS to MoBS, an algorithm that combines a ballistic search algorithm for multiple agents without communication [6] [7] and a gas gradient following algorithm for robots inspired by Moth pheromone tracking [8] [9].

## II. RELATED WORK

An algorithm for reliably locating max flux using a remote-sampling robotic-platform requires the following:

- 1) *Search*: A search pattern to explore an area to make initial contact with the plume.
- 2) *Plume Gradient Following*: After plume contact is made, the platform follows the gas plume to the source.
- 3) *Failure Resistance*: The collection of robots needs to respond to failures to maintain a cohesive structure.

Schleich et. al. [10] proposes *searching* an area using a fully-connected swarm of drones and compares this against a random and pheromone-following approach. They find that a fully-connected swarm satisfies base-station connectivity requirements while achieving slightly better survey performance for larger swarm sizes. This motivates LoCUS, as keeping the




Fig. 1: MoBS simulation with 16 drones and a smooth plume. The red lines trace each drone’s independent search for the plume using golden ratio spokes from the center of the arena. After each drone contacts the plume, it switches to a Moth pheromone inspired search algorithm to find the max flux.




Fig. 2: LoCUS simulation with 16 drones and a perturbed plume. The red lines trace the swarm’s Archimedes Spiral search for the plume. After contacting the plume, the swarm follows leverages its simultaneous spatially dispersed measurements to descend the gradient to the max flux.

swarm in contact provides benefits that outweigh the overhead of maintaining swarm connectivity.

Neumann et. al. [11] compares 3 algorithms for *plume gradient following*: the surge-cast algorithm, the Dung Beetle (zig-zag) algorithm, and the pseudo-gradient algorithm using a single robot agent. Through the author’s experiments, in both simulation and physical robots, they validate all three algorithms promising for micro UAVs each under different circumstances. Our approach uses multiple robots for plume gradient following, with MoBS closely resembling the surge-cast algorithm and LoCUS resembling pseudo-gradient algorithm across the swarm formation.

Chen et. al. [12] apply a Particle Swarm Optimization algorithm to follow a gas plume gradient in an indoor environment. This approach requires full swarm connectivity to communicate global arena information throughout the swarm. This motivates keeping the swarm connected with coordinated movement for gradient descent.

In [13], Cabrita et. al. investigate locating the max flux using Gaussian parameter estimation leveraging a simulated annealing error minimisation approach. They test this algorithm successfully on a swarm of 5 robots. We implement a similar model in MoBS and LoCUS, but we only use the local gradient to navigate the plume in the case of MoBS, and the gradient that spans the swarm’s full extent in LoCUS. We use their simple linear fit to determine the direction of the CO<sub>2</sub> gradient.

Flocking algorithms are effective at coordinating movement while being *failure resistant*. Souissi et. al. [14] and Yang et. al. [15] propose leader based approaches for moving a swarm flock while maintaining a given shape and detecting and recovering from failures. Their algorithms keep the swarm together during movement. LoCUS, on the other hand, makes theoretical guarantees about swarm symmetry as drones are lost, given a small collection of drones in close enough

proximity that all drones can maintain communication with each other. Our approach could be applied to heal traditional flocking algorithms like the one presented in [16].

Paliotta et. al. [17] present a *plume gradient following* agent based model for three fully networked agents [18] [19] [20]. We extend this structured plume gradient following approach with LoCUS by increasing the swarm size, tuning agent capabilities to mimic our dragonfly robotic platform, and adding a fault recovery mechanism.

### III. DERIVATION AND ANALYSIS OF LOCUS

The LoCUS algorithm ensures a fully connected swarm with efficient recovery from drone failures. A LoCUS swarm is able to be controlled as a single unit, by directing all members of the swarm at once. We first discuss the basic algorithm assuming no failures, and then discuss how the swarm recovers from drone failures and resumes its mission.

Let  $N$  be the total number of autonomous drones in the system. Each drone has a unique ID in  $\{1, \dots, N\}$  and a communication radius  $R_{\max}$  and a safety radius  $R_{\min}$ . Each drone can communicate with any drone within  $R_{\max}$  distance, but requires a minimum distance between any two drones in the swarm to be  $R_{\min}$  to avoid collisions.

#### A. Balanced Range-Limited Trees

**Definition 1.** Given  $R_{\min}, R_{\max} > 0$  and an integer  $n > 0$ , an  $(R_{\min}, R_{\max})$ -Range-Limited Tree on  $n$  nodes is a rooted tree, where the distance between any two nodes is at least  $R_{\min}$  and at most  $s$ . In particular, a maximal  $(R_{\min}, R_{\max})$ -Range-Limited Tree is one in which the distance between the parent node and any of its children is  $R_{\max}$ . The ratio  $\rho = R_{\max}/R_{\min}$  is the spread of this tree.

As with standard  $k$ -ary trees, we can define the height of a Range-Limited Tree  $\mathcal{T}$  node in terms of the heights of its children. We define the height of the root node as zero




Fig. 3: Assignment of nodes to levels based on their IDs. Since the number of nodes at each level is fixed, the assignment is deterministic and can be computed locally to determine placement in the swarm (see Section III-A1). The blue regions denote parent/child communication links.

and then, recursively, the height of  $\mathcal{T}$ , denoted  $\text{height}(\mathcal{T})$ , as  $\text{height}(\mathcal{T}) = 1 + \max_i \{\text{height}(\mathcal{T}_i)\}$ , where the maximum is over the height of all children  $\mathcal{T}_i$  of  $\mathcal{T}$ . Similarly, we define the level of a node as  $\text{level}(\mathcal{T}_i) = 1 + \text{level}(\mathcal{T}_i.\text{parent})$ , where  $\mathcal{T}_i.\text{parent}$  is the parent node of  $\mathcal{T}_i$ . For this recurrence, the root node is defined to be level zero. Thus, the root node has the largest height in the tree but is located at the lowest level.

**Definition 2.** Let  $\mathcal{T}$  be a Range-Limited Tree. We say that  $\mathcal{T}$  is Balanced if for every node in  $\mathcal{T}$ , the difference in the heights between any two of its children is at most one, i.e., for every node  $\mathcal{T}_i \in \mathcal{T}$  with children  $\mathcal{T}_i^{(1)}, \dots, \mathcal{T}_i^{(m)}$ , it must hold that  $|\text{height}(\mathcal{T}_i) - \text{height}(\mathcal{T}_j)| \leq 1$  for all  $i \neq j$ .

Each node maintains a pointer to its heir in the tree. This is crucial to achieve fault tolerance in LoCUS. We define the heir of a node as its successor, if it exists, or its predecessor, otherwise. If neither a successor or predecessor exists, the node is a leaf node and the heir is null. To define a successor node, we first define an inorder traversal of the tree, denoted  $\text{in}(\mathcal{T})$ . Let  $\mathcal{T}^{(1)}, \dots, \mathcal{T}^{(m)}$  be the children of the root node for  $\mathcal{T}$ . Then, the inorder traversal of  $\mathcal{T}$  prints the IDs of these nodes in the following order (here,  $\cdot$  represents the concatenation operator):  $\text{in}(\mathcal{T}^{(1)}) \dots \text{in}(\mathcal{T}^{\lfloor \frac{m}{2} \rfloor}) \cdot \text{ID}(\mathcal{T}) \cdot \text{in}(\mathcal{T}^{\lfloor \frac{m}{2} \rfloor + 1}) \dots \text{in}(\mathcal{T}^{(m)})$ . Note that the inorder traversal is unique for a given tree. We can now define the successor and predecessor of a node.

**Definition 3.** Node  $\mathcal{T}_j$  is a successor of the node  $\mathcal{T}_i$  in the tree  $\mathcal{T}$  if  $\text{ID}(\mathcal{T}_j)$  immediately follows  $\text{ID}(\mathcal{T}_i)$  in the inorder traversal of  $\mathcal{T}$ . Similarly, we say that  $\mathcal{T}_j$  is a predecessor of  $\mathcal{T}_i$  if  $\text{ID}(\mathcal{T}_j)$  immediately comes before  $\text{ID}(\mathcal{T}_i)$  in the inorder traversal of  $\mathcal{T}$ . In all cases, a node is either a leaf, or either its successor or predecessor is a leaf node of tree  $\mathcal{T}$ .

1) *Formation Algorithm:* The LoCUS algorithm swarm takes the shape of an  $(R_{\min}, R_{\max})$ -Balanced Range-Limited Tree. A balanced Range-Limited Tree layout obtains maximal

spatial coverage while maintaining a minimum separation between drones to avoid collisions, and keeps drones within communication range.

**Lemma 1** (Number of Nodes at a Given Level). Let  $\mathcal{T}$  be a maximal balanced  $(r, s)$ -Range-Limited Tree on  $N$  nodes with  $\rho = s/r$ . Then, the number of nodes at level zero is given as  $n_0(\rho) = 1$ , and for each  $k > 0$ , the number of nodes at level  $k$  is  $n_k(\rho) = \left\lfloor \frac{2\pi}{\sin^{-1}(\frac{k}{\rho})} \right\rfloor$ . This is the calculation of the whole number of nodes that fit on a circle at radius  $s \times k$  separated by distance  $r$ .

Drones deterministically compute their location in the swarm with respect to tree layout. This computation is local to the drones and can be calculated purely by the drone IDs (see Figure 3). From Lemma 1 we know that the number of drones at level  $k$  is  $n_k(\rho)$ . Thus, the space of drone IDs can be partitioned based the levels in which the drones belong. For example, the drone with ID 1 is the root node and has level zero, whereas the drones with IDs from 2 to  $n_1(\rho) + 1$  all belong to level one.

Each node (besides the root) in the LoCUS tree structure holds a parent reference and list of children. This facilitates bidirectional communication throughout the swarm, as required by the LoCUS algorithm. Parent nodes are calculated by the closest node in the previous layer.

**Lemma 2** (Number of Levels). The number of levels in a maximal balanced  $(r, s)$ -Range-Limited Tree on  $N$  nodes with  $\rho = s/r$  is  $O(\sqrt{N/\rho})$ .

This is also a bound on the maximum height of the tree and hence, the maximum number of communication hops required for any node in the tree to transmit a message to any other node in the tree. In particular, when  $\rho$  is low (i.e. when the communication radius is not too large compared to the safety radius), then the diameter of the tree is  $O(\sqrt{N})$ , however, when the communication radius is large, say with  $\rho = \Omega\left(\frac{N}{\log N}\right)$ , then the diameter of this tree becomes  $O(\log N)$ , which is similar to that of a tree with constant arity. Communication is highly efficient in this case and the latency for transmitting messages is low.

2) *Insertion of Nodes:* Always insert at the first available leaf node so that insertion cost is  $O(1)$ . Insertions do not affect the balance of the tree, since no new levels are created unless the previous level is completely full.

3) *Deletion of Nodes:* Replace the deleted node by its heir. If the deleted node is a leaf node, then there is no heir replacement for this node and hence, there is no deletion cost. However, when a node at height  $h \geq 1$  fails, then its heir is located at a communication hop distance of  $O(\sqrt{N/\rho} - h)$  from this failed node. Hence, although only  $O(1)$  link changes happen upon this replacement, the total number of messages sent is  $O(\sqrt{N/\rho} - h)$ .

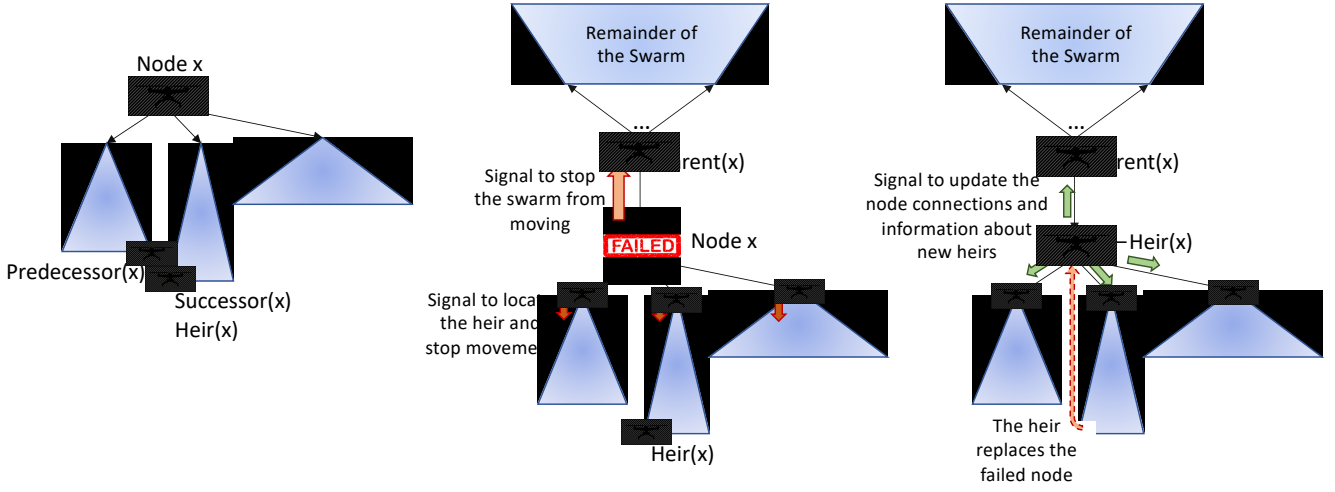


Fig. 4: A schematic of single failure recovery in the LoCUS algorithm. When a node fails, a signal is sent to its parent and children to stop the swarm movement and inform the heir. The heir node then travels to the location of the failed node and the neighboring nodes update their local information.

### B. Handling Drone Failures

The use of the Balanced Range Limited Trees data structure offers the swarm resilience against arbitrarily many crash failures, even when all but one drone remains in the system. We achieve this robustness as follows (see Figure 4) – When a node fails, a signal is sent to its parent and children to stop the swarm movement and inform the heir. This signal can be sent out when the node believes it is about to fail, for example when its battery is critically low, or by its parent and children when it fails to respond to a heartbeat. Upon receiving this signal, the heir node travels to the location of the failed node and replaces it in the swarm. Finally, because the tree structure changed, heirs are recalculated on ancestor nodes of the replacement heir node’s original leaf location.

Since each drone stores its heir information, it can directly inform the heir drone when it believes a failure is inevitable. For the case when the failure happens without any signal being sent out, the child drones use their *parentHeir* field to contact the heir drone in the swarm for recovery. If there are no child drones, then the failed drone does not require any recovery mechanism since it is already in the last layer of the swarm.

To ensure that drones do not collide with other drones while the swarm is rearranged, the heir drone descends to a distance of  $R_{\min}$  and travels at this height to its destination (see Figure 5), at which point it climbs back up to the given elevation. The swarm must stop moving during this recovery phase to avoid complicating communications and movement when the swarm is disconnected.

This heir-based recovery scheme achieves a reformation cost of  $O(\sqrt{N/\rho})$  – the bound on  $\text{height}(\mathcal{T})$  – by only inducing local adjustment in the swarm, without disturbing other drones. Note that moving one drone to replace its heir would be approximately equal to drones between the heir and the failing drone shifting up in the tree.




Fig. 5: The heir replaces the failed node by flying under the swarm at a safe distance to prevent collisions.

### C. Handling Simultaneous Drone Failures

If both a drone and its heir drone fail at the same time, and the swarm uses the algorithm above to simultaneously recover from both, then it will enter a deadlock scenario. We introduce the following algorithm for handling simultaneous failures with the caveat that it requires global knowledge of the swarm state to execute. A more advanced distributed version of this algorithm that executes without global knowledge is possible, but we leave this analysis and implementation for future work.

*Outer-Level First (OLF):* In this scheme, we use the fact the failures in outer levels of the swarm cost less to recover than failures on the inner levels. This is because the distance to the heir node is smaller in outer levels. For example, leaf nodes may be removed from the swarm outright without replacement, a node at  $\text{height}(\mathcal{T})/2$  would require its heir to move  $R_{\max} \text{height}(\mathcal{T})/2$  distance, and the root node would require its heir to move  $R_{\max} \text{height}(\mathcal{T})$  distance to replace. Thus, whenever a node gets a failure signal, it first checks to see if there is any existing failure recovery that is active in any of its children. If yes, it waits for those to finish, then proceeds to process the signal from its parent. Concretely, this is implemented by gathering a set of failures across the swarm, and processing them in descending order by  $\text{height}(\mathcal{T}_i)$ .



#### IV. THE MOBS ALGORITHM AND ITS IMPLEMENTATION

The MoBS algorithm takes a different approach to the max flux problem by allowing each UAV to navigate independently. Each UAV starts at the center of the arena and picks a uniformly random angle between  $0^\circ$  and  $360^\circ$  and sets 100 waypoints in 1 m increments from the center in that direction to produce spokes to search the arena. At each waypoint the UAV collects a gas plume sample and reacts accordingly. The UAV continues to follow the spoke waypoints if a reading of less than 0.005. Otherwise, the UAV changes strategies into the moth-pheromone chemotaxis algorithm inspired by [8]. Subsequent spokes are built by adding  $2\pi/\phi$  rad to the previous spoke angle where  $\phi$  is the golden ratio 1.618 that has been shown to search best given no communication amongst members of the swarm [6].

The moth-pheromone chemotaxis algorithm compares the gas reading at the current time step against the previous time step and determines if the signal has increased or decreased. If the signal increased then the drone continues moving in the same direction. If the signal stays the same or decreases then the drone moves in a new uniformly random direction. A zero signal detected for greater than 4 time steps reverts the drone back to continue the golden ratio driven spoke search algorithm. Because there is never any communication among UAV in MoBS, failed UAV stop collecting samples but have no impact on other UAV.

#### V. EXPERIMENTAL METHODS

We measure performance of both algorithms for a range of swarm sizes and failure scenarios in simulation. Given the practical limitation of battery life on flight time, our primary interest is minimizing the time to find the max  $\text{CO}_2$  flux. We halt the simulation when the max flux is found (a drone samples within 1 m of the max flux location), if the entire swarm is in a failed state, or when 17.3 h of simulation time has passed ( $10^6$  time steps).

We implement the LoCUS and MoBS algorithms in Autonomous Robots Go Swarming (ARGoS) [21]. ARGoS is a C++ and Lua based physics multi-robot simulator and is suitable for proof-of-concept simulations, while preserving realistic physical dynamics with the DYN3D physics engine. We use ARGoS to simulate Spiri UAV (Pleiades Robotics Inc) including 3 dimensional locality (GPS) inputs and go-to coordinate capabilities. Additionally, we are able to command  $N$  drones in the simulation. These capabilities make ARGoS a natural fit for experimental investigation of LoCUS and MoBS.

The gas plume is modeled in ARGoS as a simple two dimensional slice of a Gaussian plume [22] with a source max flux location ( $x$  and  $y$ ), stack height ( $H = 10$  m), wind speed ( $u = 50$  m/s), emission rate ( $Q = 2$  kg/s), and diffusion rate ( $K = 1$  kg/s):

$$\text{UNPERTURBED}(x, y) = \frac{Q}{2\pi Kx} \exp\left(-\frac{u(y^2 + H^2)}{4Kx}\right) \quad (1)$$

<sup>1</sup>Implementation source code can be found at <https://tinyurl.com/tne7tzu>

The source of the plume is located at a uniformly random location in the simulation within 100 m of the UAV take-off location. Each UAV may detect the gas concentration at its given coordinate as a floating point value between 0 (low) and 1 (high) gas concentration signals. To limit the experimental variance, we only vary the location of the plume and not the shape, intensity, or rotation of the plume. We test the algorithms against the smooth plume described in [1] and a perturbed version of the plume designed to make following the gradient more realistic and challenging:

$$\text{PERTURBED}(x, y) = (0.8 + 0.2 \sin(4x)) \text{UNPERTURBED}(x, y) \quad (2)$$

Our two failure models in these experiments are motivated by flying a swarm of UAV to gather volcano gas  $\text{CO}_2$  emission data.

*Generic Failures:* To represent a UAV battery failure, crashes, or other miscellaneous failures that increase in likelihood as flight time increases, we use a uniform failure probability per drone per time step given by  $p_f > 0$ , which depends on the number of drones existing in the system at time  $t$ .

*In-Plume Failures:* We use the gas plume emissions reading  $r$  at time  $t$  to drive the probability of failure on each drone given by  $p_{fr} > 0$ . This models the higher probability of failure as corrosive gases or temperatures associated with more concentrated volcano gas emissions are encountered.

Drone failure is represented by a boolean flag on the drone controller that, if enabled, stops the drone from moving or receiving further waypoints from its parent. Once a drone fails, it is never recovered.

##### A. Implementation of LoCUS

The LoCUS algorithm arranges members of the swarm by distributing each drone through space using specified  $R_{\min}$  and  $R_{\max}$ . The unique ID of each member of the swarm allows a unique 2D location offset from the central root node to be calculated. The drones are distributed in a plane by each offset using a constant height of 10 m. Each drone's parent is assigned by finding the closest drone in the previous shell of the swarm. This parent/child relationship constructs the data structure pivotal to maintaining communication throughout the swarm.

We implement a recursive algorithm to distribute navigation waypoints by communicating them from the root drone down through its children, to its children's children, and so on. When waypoints are distributed, the swarm offset location for each drone is added to the waypoint to ensure that the swarm maintains  $R_{\min}$  and  $R_{\max}$ .

At takeoff, the root LoCUS drone is given the initial starting position. Using the recursive waypoint distribution, this initial starting position waypoint directs the swarm to assemble the shell structure exhibited by Figure 3.

To make initial contact with the plume, the swarm is directed from the root to follow the Archimedes' spiral. For

coordinated swarm search, the Archimedes' spiral has been shown to find targets faster than a spoke algorithm [23]. This search pattern is created by building waypoints along the spiral. Each waypoint is calculated to space the arms of the spiral by the radius of the largest full swarm shell and an incremented angle. Using the radius of the swarm ensures that we have full coverage of the simulation arena.

After a waypoint is reached, a plume gas reading is sampled from each drone and communicated via the tree structure up to the root drone where the readings ( $val$ ) and associated gps coordinates ( $x, y$ ) are aggregated into the  $uav$  array. The aggregated data is input into matrix and vector form and fit with a slope ( $b$ ) using linear regression in the form  $Ab + \epsilon = y$  by minimizing  $\epsilon$  through least squares approximation provided by the Eigen C++ library [24]:

$$\underbrace{\begin{bmatrix} 1 & uav[1].x & uav[1].y \\ 1 & uav[2].x & uav[2].y \\ \vdots & \vdots & \vdots \end{bmatrix}}_A \underbrace{\begin{bmatrix} b[0] \\ b[1] \\ b[2] \end{bmatrix}}_b + \underbrace{\begin{bmatrix} \epsilon[0] \\ \epsilon[1] \\ \vdots \end{bmatrix}}_\epsilon = \underbrace{\begin{bmatrix} uav[1].val \\ uav[2].val \\ \vdots \end{bmatrix}}_y \quad (3)$$

The slope of this linear fit ( $b[1], b[2]$ ) is used to provide a normal vector to direct the swarm to perform a gradient descent in the direction of the highest plume signal. If a zero magnitude linear slope is found, then the swarm continues to follow the Archimedes spiral.

Failures are handled as follows. First, once a waypoint is reached, failures in the swarm are queried for from the root. This is a recursive call, similar to the waypoint distribution, to gather a set of the failed status of the entire connected swarm. For these experiments and for simplicity, we implement the Outer-Level First (OLF) scheme failure recovery model. This scheme requires global knowledge of the swarm as it uses failed drones to determine the status of their children. We then proceed to heal the swarm as outlined in **III-B**. Handling Failures remove each of the failed drones and replace them by their heir in order, waiting for each heir to take the place of the failed drone before proceeding to the next failed drone. After the replacement, heirs of all ancestors up to the root are recalculated to take into account the change in the swarm. Of course, if a failed drone is found to have no heir (a leaf drone) then they are removed from the swarm without replacement.

Once all failures are processed a swarm re-balance is executed to ensure a consistent minimum radius to the swarm. This iteratively removes leaf drones from the deepest branches of the tree and inserts them into the shallowest branches of the tree. The root node executes this operation until  $height_{max} - height_{min} \leq 1$ . Once there are no more failures in the swarm and the swarm is re-balanced, then the next waypoint is calculated and the swarm movement continues.

We observed corner-case scenarios where the swarm oscillated between two points, never moving towards the max flux location. To resolve this, we added both a 0.1 m random offset and a uniformly random rotation to the swarm between  $0^\circ$  and  $45^\circ$  at each waypoint. This randomization strategy allowed to

swarm to exit these oscillating corner cases.

## B. Experimental Setup

The experimental factors we explore in simulation are swarm size, whether the plume gradient is smooth or perturbed, the failure rate, and whether the failure rate increases with gas concentration. For the LoCUS algorithm, we set  $R_{min} = 3$  m and  $R_{max} = 3$  m. The response variables are whether the max flux was found, the elapsed time before encountering the plume, and the total time taken to find the max flux.

To compare LoCUS and MoBS we perform the following four experiments. In experiment 1, we compare the time to find max flux of the smooth plume for LoCUS and MoBS in 100 trials with both failure probabilities set to 0 by varying the swarm size from 2 to 20 UAV. In experiment 2, we duplicate experiment 1 using the perturbed plume. These first two experiments are designed to compare the times to encounter the plume and navigate to the maximum flux of LoCUS and MoBS without failures. In experiment 3, we vary the generic failure probability from  $10^{-1}$  to  $10^{-6}$ , set the in-plume failure probability to 0 (so that the probability of failure is the same inside and outside of the plume), and use the smooth plume over 100 trials with a swarm size of 20. In experiment 4, we duplicate experiment 3 but vary the in-plume failure probability from  $10^{-1}$  to  $10^{-6}$  and set the generic failure probability to 0. The last two experiments measure the impact of failures on LoCUS and MoBS. In each of these experiments we compare the performance of LoCUS with healing enabled and disabled. This enables us to assess whether maintaining symmetry though healing is worth the time taken to repair the swarm.

## VI. EXPERIMENTAL RESULTS

### A. Experiment 1: Unperturbed Navigation

Experiment 1 compares the time to reach max flux of the unperturbed in swarms of 5 to 20 UAV for LoCUS and MoBS depicted in Figure 6 with the unhashed bars. We find that, for smaller swarm sizes (up to the 5 UAV shown), the average time and standard deviation to plume contact and max flux for LoCUS is significantly smaller than MoBS. For larger swarm sizes (10, 20), MoBS reaches plume contact on average faster. LoCUS navigates from plume contact to max flux in about the same time as MoBS, but LoCUS has less variance in time to achieve both plume contact and max flux.

### B. Experiment 2: Perturbed Navigation

Experiment 2 extends experiment 1 using the perturbed plume depicted in Figure 6 with the hashed bars. We find that the difference in plume dynamics significantly increases the average time to max flux for the MoBS algorithm, but the LoCUS time to max flux remains short. This increases the average time to max flux for MoBS so that is slower than LoCUS for all swarm sizes tested. Additionally, the standard deviation for MoBS is much larger than that of LoCUS which is partly driven by several outliers that lasted up to 344, 124,




Fig. 6: Time to find the max flux (top, purple) and initial plume contact (bottom, orange) for LoCUS and MoBS for swarm sizes 5, 10, and 20 in the unperturbed and perturbed plumes. Stars show the median time to max flux and error bars are one standard deviation centered at the mean.

and 31 minutes for 5, 10 and 20 UAV respectively. These times risk failure for the drones to return given the UAV battery capacity.

### C. Experiment 3: Generic Failure Effects

Experiment 3 includes generic failure probabilities from  $10^{-1}$  to  $10^{-6}$  including a specialized version of LoCUS that does not heal from failures as depicted in Figure 7. We find that, with generic failures, LoCUS and MoBS both respond similarly to the failure probability by beginning to unsuccessfully complete the max flux location task between the probability of failures of  $10^{-4}$  and  $10^{-3}$ . This is contrasted against the LoCUS without healing that responds much earlier to the probability of failure at about  $10^{-5}$ .

### D. Experiment 4: In-Plume Failure Effects

Experiment 4 extends experiment 3 using the in-plume failure model. With in-plume failures, LoCUS and MoBS both begin to fail to complete the max flux location task at failure rates of  $10^{-2}$ . This is contrasted against the LoCUS without healing that fails to complete the task with much lower UAV failure probabilities of about  $10^{-4}$ .

## VII. DISCUSSION

LoCUS provides a failure-tolerant structure for exploring and pinpointing the max flux location of a  $\text{CO}_2$  plume. LoCUS guarantees that a group of drones can communicate to each other simultaneous spatially dispersed measurements which can be used to calculate a gas gradient better than an individual drone. This is particularly useful for finding the location of maximum flux in a perturbed plume such as those produced by volcanos in dynamic environments. LoCUS provides a




Fig. 7: Success rates in 100 trials with generic and in-plume failures with 20 drones. The left thin solid lines are LoCUS with healing disabled and the right thick solid lines are LoCUS with healing enabled. The two left orange thin and thick solid lines are LoCUS with generic failures, while the green right solid lines are LoCUS with in-plume failures. MoBS success counts are graphed using dashed and dotted lines – the left dashed line is MoBS with generic failures and the right dotted line is MoBS with in-plume failures. This shows how significant healing is to LoCUS successfully completing the max flux location task.

way re-form the swarm given the inevitable failure of drones in hazardous conditions present when monitoring gas efflux from volcanoes. We compare LoCUS with the fully dispersed MoBS algorithm and show in experiments 1 and 2 that the LoCUS algorithm is able to find the max flux of a plume, both smooth and perturbed, at least as fast as the MoBS algorithm in expectation, but with substantially smaller variation. The better worst-case performance of LoCUS is important given time limits imposed by battery life. Additionally, the LoCUS algorithm is able to find the max flux faster than MoBS after initial plume contact, particularly in a perturbed plume simulations.

For large swarms, the MoBS algorithm makes initial contact with the plume faster than LoCUS on average. The superior performance of MoBS at finding the plume, and LoCUS of finding the source once in a plume, suggests an approach that combines the best of both algorithms. For large swarms, we may perform the initial search for the plume using the more dispersed golden spoke algorithm used in MoBS. Then, when contact is made, a LoCUS structured formation can leverage nearby drones to perform gradient descent informed by communication among drones. Future work can explore the benefits of a fully dispersed set of spokes, with the first drone that contacts the plume calling nearby UAV to join together once the plume is found. Alternatively, sufficiently many UAV could be divided into multiple small LoCUS sub-swarms to use spoke search to contact the plume and LoCUS

enabled gradient descent by each independent sub-swarm once it contacts the plume.

LoCUS relies heavily on its loss recovery model in order to maintain communication between spatially dispersed drones to perform a more robust gradient descent once a plume has been found (see the red lines within the plume in Figure 2). The loss recovery model allows the swarm to reorganise once a failure has been detected and continue to rely on receiving CO<sub>2</sub> measurements from multiple locations. In practice, we observed LoCUS successfully locating max flux with failures in a majority of the swarm, even down to a single remaining drone. We also observed that the loss recovery time is so fast that it is dominated by the time to encounter the plume and time to max flux. Thus, the time to recover the swarm formation is worth the superior gradient following performance provided by having spatially dispersed measurements. In experiments 3 and 4 we show that self-healing is critical to the success of LoCUS gradient following.

In comparison to MoBS, LoCUS is especially vulnerable to the in-plume failures. This is (ironically) because LoCUS brings the entire swarm into the plume and quickly closer to the source, putting swarm members in jeopardy due to more in-plume failures as the source is approached near the source of volcano efflux. In contrast, some of the UAV in MoBS spend more time out of the plume, making them less susceptible to in-plume failures. The results in experiment 4 show that even in the worst case for LoCUS, it can leverage the healing algorithm to mitigate this problem to complete the max flux task nearly as often as MoBS.

Being able to reliably and quickly determine the max CO<sub>2</sub> flux with drones that are limited to a maximum 1 h flight time, with practical swarm sizes for transportation to remote and hazardous regions, and are tolerant of drone loss is critical to the study of volcano behaviour. With LoCUS, we have demonstrated an algorithm that solves the CO<sub>2</sub> max flux task faster than, and approximately as reliably as a more dispersed approach.

## VIII. ACKNOWLEDGEMENTS

We thank the UNM Vice President for Research, the Department of Energy's Kansas City National Security Campus, operated by Honeywell Federal Manufacturing & Technologies, LLC under contract number DE-NA0002839, a James S. McDonnell Foundation Complex Systems Scholar Award and DARPA award #FA8650-18-C-6898 for funding. We thank the UNM Center for Advanced Research Computing, supported in part by the National Science Foundation, for providing the high performance computing resources used in this work.

## REFERENCES

- [1] S. K. Brown, S. F. Jenkins, R. S. J. Sparks, H. Odbert, and M. R. Auken, "Volcanic fatalities database: analysis of volcanic threat with distance and victim classification," *Journal of Applied Volcanology*, 2017.
- [2] A. Aiuppa, T. P. Fischer, T. Plank, and P. Bani, "CO<sub>2</sub> flux emissions from the Earth's most actively degassing volcanoes, 2005–2015," *Scientific Reports*, 2019.
- [3] J. A. Diaz *et al.*, "Utilization of in situ airborne MS-based instrumentation for the study of gaseous emissions at active volcanoes," *International Journal of Mass Spectrometry*, vol. 295, no. 3, pp. 105–112, 2010.
- [4] E. J. Liu *et al.*, "Dynamics of Outgassing and Plume Transport Revealed by Proximal Unmanned Aerial System (UAS) Measurements at Volcán Villarica, Chile," *Geochemistry, Geophysics, Geosystems*, 2019.
- [5] E. J. Liu *et al.*, "Multi-species volatile fluxes from manam, papua new guinea," in prep.
- [6] A. Aggarwal, W. F. Vining, D. Gupta, J. Saia, and M. E. Moses, "A most irrational foraging algorithm," 2019.
- [7] A. Aggarwal, *Thwarting Adversaries with Randomness and Irrationality*. PhD thesis, University of New Mexico, 2019.
- [8] W. Li, J. A. Farrell, S. Pang, and R. M. Arrieta, "Moth-inspired chemical plume tracing on an autonomous underwater vehicle," *IEEE Transactions on Robotics*, 2006.
- [9] J. H. Belanger and M. A. Willis, "Biologically-inspired search algorithms for locating unseen odor sources," in *Proceedings of the 1998 IEEE International Symposium on Intelligent Control (ISIC) held jointly with IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA) Intell*, pp. 265–270, IEEE, 1998.
- [10] J. Schleich, A. Panchapakesan, G. Danoy, and P. Bouvry, "UAV fleet area coverage with network connectivity constraint," in *MobiWac 2013 - Proceedings of the 11th ACM International Symposium on Mobility Management and Wireless Access, Co-located with ACM MSWiM 2013*, 2013.
- [11] P. P. Neumann, V. H. Bennetts, A. J. Lilienthal, and M. Bartholmai, "From insects to micro air vehicles—a comparison of reactive plume tracking strategies," in *Advances in Intelligent Systems and Computing*, 2015.
- [12] Y. Chen, H. Cai, Z. Chen, and Q. Feng, "Using multi-robot active olfaction method to locate time-varying contaminant source in indoor environment," *Building and Environment*, vol. 118, pp. 101–112, 2017.
- [13] G. Cabrita and L. Marques, "Estimation of Gaussian plume model parameters using the simulated annealing algorithm," in *Advances in Intelligent Systems and Computing*, 2014.
- [14] S. Souissi, T. Izumi, and K. Wada, "Oracle-based flocking of mobile robots in crash-recovery model," *Theoretical Computer Science*, vol. 412, no. 33, pp. 4350–4360, 2011.
- [15] Y. Yang, S. Souissi, X. Défago, and M. Takizawa, "Fault-tolerant flocking for a group of autonomous mobile robots," *Journal of systems and Software*, vol. 84, no. 1, pp. 29–36, 2011.
- [16] C. Virágh, G. Vásárhelyi, N. Tarcai, T. Szörényi, G. Somorjai, T. Nepusz, and T. Vicsek, "Flocking algorithm for autonomous flying robots," *Bioinspiration & biomimetics*, vol. 9, no. 2, p. 025012, 2014.
- [17] C. Paliotta, D. J. Belleter, and K. Y. Pettersen, "Adaptive source seeking with leader-follower formation control," *IFAC-PapersOnLine*, vol. 48, no. 16, pp. 285–290, 2015.
- [18] M. Breivik, V. E. Hovstein, and T. I. Fossen, "Ship formation control: A guided leader-follower approach," *IFAC Proceedings Volumes*, vol. 41, no. 2, pp. 16008–16014, 2008.
- [19] E. Biyik and M. Arcak, "Gradient climbing in formation via extremum seeking and passivity-based coordination rules," in *2007 46th IEEE Conference on Decision and Control*, pp. 3133–3138, IEEE, 2007.
- [20] R. Bachmayer and N. E. Leonard, "Vehicle networks for gradient descent in a sampled environment," in *Proceedings of the 41st IEEE Conference on Decision and Control*, 2002., vol. 1, pp. 112–117, IEEE, 2002.
- [21] C. Pinciroli *et al.*, "Argos: a modular, multi-engine simulator for heterogeneous swarm robotics," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5027–5034, IEEE, 2011.
- [22] J. M. Stockie, "The mathematics of atmospheric dispersion modeling," *Siam Review*, vol. 53, no. 2, pp. 349–372, 2011.
- [23] A. Aggarwal, D. Gupta, W. F. Vining\*, G. M. Fricke, and M. E. Moses, "Ignorance is not bliss: An analysis of central-place foraging algorithms," in *Proceedings of the Conference on Intelligent Robots and Systems (IROS) IEEE/RSJ. Accepted and available at https://bit.ly/2QcPaX8*, 2019.
- [24] B. Jacob and G. Guennebaud, "Eigen is a c++ template library for linear algebra: matrices, vectors, numerical solvers, and related algorithms," 2012.