

SensorBFT: Fault-Tolerant Target Localization using Voronoi diagrams and Approximate Agreement

Akhil Bandarupalli*, Adithya Bhat†, Somali Chaterji*, Michael K. Reiter§¶, Aniket Kate*‡, Saurabh Bagchi*

*Purdue University {abandaru, schaterji, aniket, sbagchi}@purdue.edu

†Visa Research haxolotl.research@gmail.com

§Duke University michael.reiter@duke.edu

‡Supra Research, ¶Chainlink Labs

Abstract—The target localization primitive is used for detecting and locating an adverse event called a *target* in a geographic area. This versatile primitive is applicable in the physical security domain (e.g., detecting intruders in an area) or for disaster preemption, such as detecting ignition events of forest fires. Prior systems implemented this primitive over large areas by deploying a network of sensor devices, which detect changes in a specific physical parameter like pressure or temperature induced by a target. However, these systems are not designed for use in adverse environments where one or more sensors can behave in a faulty manner. While many algorithms in the distributed systems literature can be naively used to implement target localization in a fault-tolerant manner, these approaches are energy-intensive as they use computationally expensive cryptographic operations not appropriate for resource-constrained sensors.

We present **SENSORBFT**, an energy-efficient, fault-tolerant approach for target localization. **SENSORBFT** uses a novel asynchronous approximate agreement protocol that enables correct sensors to achieve an approximate consensus in the presence of faulty sensors. Sensors fulfill their energy budgets by tuning the precision and accuracy of localization, where precision is the difference between honest sensors’ outputs and accuracy is the difference between an honest sensor’s output and the target’s true location. In optimal scenarios, this protocol reduces communication from $\mathcal{O}(n^3)$ to $\mathcal{O}(n^2)$ messages per round, where n is the number of sensors sharing coverage over a piece of area. In a sensor testbed with $n = 19$ sensors, **SENSORBFT** consumes $\frac{2}{5}$ th the energy consumed by existing solutions for a minor 2% loss in accuracy, significantly enhancing efficiency and coverage.

I. INTRODUCTION

Recent progress in autonomous sensor systems has opened up new possibilities for use in precision agriculture [44], [46] and battlefield surveillance [7], [8], [15], [18]. Sensor-equipped CPS (Cyber-Physical Systems) are useful in detecting and neutralizing threats like pests on a farm, fire events in dry fields, or adversarial objects on a battlefield. In such applications, sensors deployed over an area must collectively sense and locate such adverse events called *targets*, to quickly and efficiently respond to them. The *target localization* primitive allows sensors to estimate and agree on the target’s location [48].

In such applications, sensors are typically resource-constrained and often deployed in adverse environments such as war zones, natural disasters, or with motivated human actors, which can make them behave in an arbitrary manner. Thus, tolerating a subset of faulty sensors is essential for operating under these conditions.

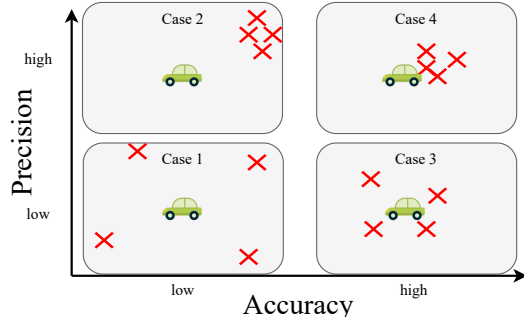


Figure 1: A pictorial description of localization precision and accuracy. The four cases describe a set of $N = 4$ sensors’ outputs on where the car (target) is centered. Case 4 is with high precision and high accuracy (low ρ and α) is desired for target localization.

More formally, we assume a system consists of N sensors, with at most T faulty sensors, deployed across an area A . In a target localization protocol, each sensor inputs a reading v_i and outputs a target location $l_i \in A \cup \{\perp\}$. Along with the standard termination guarantee that requires each honest sensor to output a location, the protocol must also be sound such that if there is no target in the area, then the output of each honest sensor must be \perp . Moreover, when the target is present in the area, the protocol should detect it with *precision* ρ , and *accuracy* α .¹ In particular, if a target is at location $\mathcal{L} \in A$, then every honest sensor’s output must be at most a distance α from the target’s location \mathcal{L} , and the outputs of any two honest sensors must be at most ρ distance from each other. See Fig. 1 towards understanding the difference between location precision and accuracy.

The conventional solution in sensor networks [14], [32], [35], [43] is to employ a central node called a *Fusion Center* (FC) or *Gateway*, a central decision maker and a trusted source of computing, to collect readings from sensors, estimate the target’s location, and send the output to all the nodes in the network. However, in our target autonomous sensor system applications, a single trusted party for the area may not be

¹We define the terms Precision ρ and Accuracy α based on literature from instrumentation, where the terms quantify the observational error.




Figure 2: The graphic represents a strawman protocol in the boolean coverage model, where each sensor has a fixed sensing range and holds the coverage responsibility for a given area (marked by dotted lines here). The orange sensor detects the car in its sensed area and sends value 1 through gossip to the network, and every other sensor sends 0. Each sensor outputs the orange sensor’s location as the target’s location.

available or acceptable, because of which we assume there is no FC available.

Early non-fault tolerant distributed approaches [16], [38] operated in the *boolean coverage model*, in which each sensor has a fixed circular sensing range with radius r [38], [47], where the sensor measures 1 if the target is within distance r from its location, and measures 0 otherwise. These systems deployed sensors such that each point in the area is within the sensing range of at least one sensor. This sensor is *responsible* to cover the area within its range. If the target is within this range, the responsible sensor sends the value 1 to every sensor in the network and gossips 0 otherwise. Upon receiving a message from every sensor, each sensor computes the average of the sensors’ locations that sent 1 and outputs it as the target’s location. This protocol achieves perfect precision ($\rho = 0$) and accuracy $\alpha \leq r$ under the assumption that all sensors are honest. We describe an example in Fig. 2.

Subsequent works identified that aggregating raw readings from sensors (rather than 1s and 0s, as in the boolean model) preserves more information about the sensors’ relative proximity to the target and leads to higher localization accuracy [4], [27], [48]. Xing *et al.* [48] requires each sensor to gossip its raw sensor reading. Each node then sorts the received readings and outputs the location of the sensor with the highest measured reading as the target’s location. This approach achieves higher localization accuracy than the boolean coverage model.

However, in the presence of a single malicious faulty sensor, both these approaches do not work. The faulty sensor i can crash and not report a target in its area, which forces other sensors to wait for i ’s response forever. Moreover, i can send different values to sensors j and k , which forces j and k to output different locations thus affecting precision. Further, i can also send an incorrect reading to every other sensor, which makes other sensors output a location far away from the target’s original location. Thus, each behavior described above violates the Termination, Precision, Accuracy, and Soundness

properties of target localization, respectively.

Some prior works [23], [28] solve localization in the presence of Byzantine faults using Byzantine Agreement (BA) protocols. Some works explored Byzantine Agreement (BA) protocols in resource-restricted CPS devices [9], [28], [34]. However, these protocols rely on bounded synchronous or partially synchronous network assumptions. However, any form of synchrony assumption is untenable in low-powered sensor networks due to potential security issues like network jamming [33], which is easier to conduct in open areas with unrestricted access. Asynchronous BA protocols are not vulnerable to these issues, but the FLP impossibility result [22] states that it is impossible to achieve deterministic asynchronous BA even with a single faulty node.

Prior works circumvent this impossibility using randomized protocols that use distributed randomness in the form of common coins. However, common coins are expensive, with the most energy-efficient implementation using threshold BLS signatures [11] requiring computationally expensive cryptography like bilinear pairings [10]. For instance, each pairing operation consumes $1000\times$ more energy than a cryptographic hash computation (such as SHA2, SHA3), which is a burden on energy-constrained sensors. Moreover, threshold BLS signatures require an Asynchronous Distributed Key Generation (ADKG) [31] setup, which is infeasible in sensor networks because of its high communication and cryptographic costs.

A. Our Approach

We build SENSORBFT, a distributed sensor system for target localization. We describe SENSORBFT in two parts: (a) We shard the area into disjoint pieces (or cells) and assign coverage responsibility of each cell A_i to a set of n sensors denoted by \mathcal{N}_i . We do this to ensure that faulty sensors do not have a monopoly over any subarea. (b) We run an Asynchronous Approximate Agreement (AAA) protocol among the n sensors of each such shard to enable them to agree on a *representative value* for their subarea, which is used in an aggregation as in Xing *et al.* [48]. Our approach ensures that each shard behaves as a single cohesive unit equivalent to a single honest sensor covering cell A_i like in Fig. 2. We show a pictorial overview in Fig. 3.

Voronoi sharding. Our first building block is a location-based sharding approach called *Voronoi sharding*, based on higher-order Voronoi diagrams. A n -order Voronoi diagram divides the entire area into smaller disjoint pieces or *cells*, where each cell is assigned to its n closest sensors. Using this approach, we divide the entire sensed region into shards, each composed of n sensors responsible for covering their cell. We set the value n such that every shard has a supermajority, i.e., $> 2/3$, of honest sensors.

Approximate Agreement. Our second building block is an AAA protocol. The AAA primitive enables sensor nodes to output values within a ρ distance of each other within the initial range of honest sensor inputs (called Convex-Hull Validity). In contrast with asynchronous BA, AAA does not

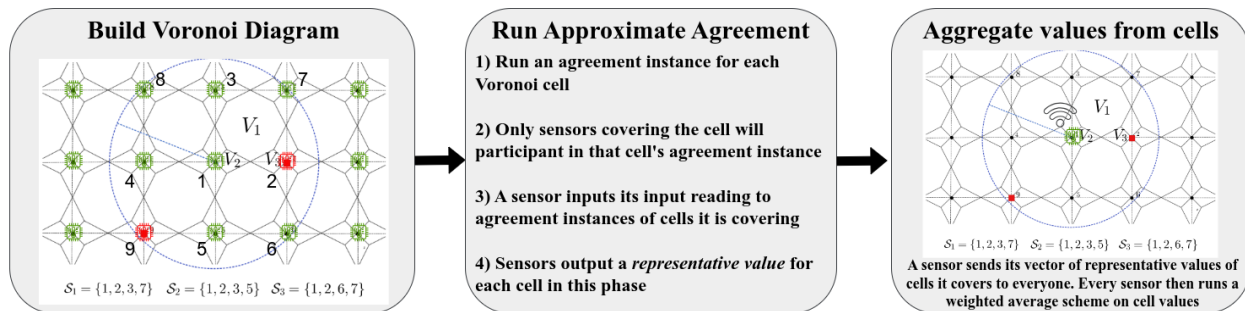


Figure 3: The three parts of SENSORBFT include (a) Dividing the area into cells based on a n -order Voronoi diagram ($n = 4$ in the figure) and assigning them to sensors, (b) Performing Asynchronous Approximate Agreement (AAA) in each cell, and (c) Aggregating representative values from all cells. This approach ensures that each set of sensors behaves like a single honest sensor covering cell A_i like in Fig. 2.

require energy-expensive public-key or threshold cryptography, which makes it friendly for resource-constrained devices like sensors. In exchange for this computational efficiency and corresponding energy savings, SENSORBFT achieves a relaxed form of localization called *approximate localization*, where sensors' estimates of the target's location are at most ρ meters away from each other.

The main challenge with using prior AAA protocols like Abraham *et al.* [1] is their high message complexity. Abraham *et al.*'s protocol requires $\mathcal{O}(n^3)$ messages per round over $\log(\frac{d \cdot \text{polylog}(N)}{\rho N})$ rounds giving a total message complexity of $\mathcal{O}(n^3 \log(\frac{d \cdot \text{polylog}(N)}{\rho N}))$, where d is the maximum distance between two points in an area (called the diameter). In circular or square-shaped areas, $d = \mathcal{O}(\sqrt{A})$. This high message complexity also results in high energy consumption, which quickly drains the sensors' energy reserves.

We address this issue using a novel AAA protocol with a reduced message complexity. We achieve this efficiency in exchange for a small loss in localization accuracy by using a novel relaxed form of Validity for AAA called δ -relaxed Validity, where honest sensors output values within the range $[\min(\mathcal{M}) - \delta, \max(\mathcal{M}) + \delta]$. \mathcal{M} is the set of honest sensors' input values and δ is a configurable parameter governing the relaxation in Validity. The accuracy α is inversely related to δ , where a higher δ and corresponding information loss in consensus implies less accurate localization. Overall, for a localization precision and accuracy of ρ and α meters, our AAA protocol requires $\mathcal{O}(n^3 \log(\frac{d \cdot \text{polylog}(N)}{\alpha N}) + n^2 \log(\frac{\alpha}{\rho}))$ messages. For an accuracy loss characterized by $\alpha = \mathcal{O}(\frac{\rho}{N} \text{dpolylog}(N))$ (which we show is practically affordable for reasonable localization), the protocol's message complexity is $\mathcal{O}(n^3 + n^2 \log(\frac{\rho}{\alpha}))$, which removes the dependence of the n^3 term from the precision ρ .

Overall, SENSORBFT provides two tuning knobs for tuning localization precision ρ and accuracy α , which sensors can use to conduct energy-efficient localization while meeting their energy budgets.

Evaluation. We implement SENSORBFT in Rust and evaluate it on a testbed of embedded nodes with acoustic sensors where each sensor can detect the sound emitted by a target. We measure and compare the energy consumed by SENSORBFT to other localization protocols. With $n = 19$ sensors, SENSORBFT achieves a practically insignificant loss in precision $\rho = 0.5\text{m}$ and consumes half the energy of a localization approach using Abraham *et al.* as the agreement protocol, and $\frac{2}{5}$ th the energy consumed by a precise ($\rho = 0$) localization approach with the state-of-the-art Asynchronous convex BA protocol FIN [21] as the agreement protocol, with only a 2% increase in accuracy parameter α .

In summary, we make the following contributions:

- 1) We devise an effective spatial sharding approach leveraging higher-order Voronoi diagrams, which guarantees comprehensive coverage of the sensed region even in the presence of malicious nodes.
- 2) We present an energy-efficient approximate agreement protocol, using precision and accuracy as tuning knobs to control the communication complexity and related energy consumption. In specific scenarios, our protocol also improves the asymptotic communication complexity of asynchronous optimally-resilient AAA by $\mathcal{O}(n)$ factor.
- 3) Using the above two contributions, we design the first computationally efficient, distributed, and approximate sensor localization scheme that can tolerate Byzantine faults in an asynchronous network.

II. DEFINITIONS AND PRELIMINARIES

A. Target Localization Problem

We assume the structure of the 2D area A that needs to be covered and the positions of the N sensors are given as input parameters to the system. We assume the network between sensors is asynchronous. This is because issues like network jamming and flooding can violate all forms of synchrony for message delivery in low-powered wireless networks.

Target. A target is an object of interest that needs to be detected and localized by the sensor network. A target is defined by its location (x, y) in the area and its innate signal strength S , which is a real number greater than or equal to a

minimum S_{\min} . This strength characterizes the disturbance or change induced by the target in a specific physical parameter like the sound level or temperature, where a higher S implies a stronger disturbance. Examples of targets include intruders or vehicles in an area that produce a sound disturbance or a forest fire that induces a temperature change. The disturbance induced by a target attenuates with increasing distance d , according to the function $f_{\text{DIST}}(d) = \frac{1}{1+d^2}$ specified in Xing *et al.* [48]. Here, $f_{\text{DIST}} : \mathbb{R}_0 \rightarrow [0, 1]$ and $f_{\text{DIST}}(0) = 1$.

We formally define the target localization problem.

Definition II.1. Target Localization: *Given N sensors, out of which at most T are faulty, deployed across an area A , the set of sensors and their locations denoted by \mathcal{N} , and at most one target present in the area, a protocol Π_{TL} where each sensor $i \in \mathcal{N}$ inputs a reading v_i and outputs $l_i \in A \cup \{\perp\}$ is a valid target localization protocol if it satisfies the following properties.*

Termination: *Every honest sensor $i \in \mathcal{N}$ must output a value l_i .*

ρ -Precision: *If a target object is present in the area, the outputs of any two honest sensors $i, j \in \mathcal{N}$ must be at most ρ distance from each other; i.e., $\|l_i - l_j\|_2 \leq \rho$.*

α -Accuracy: *If a target object is present at location $\mathcal{L} \in A$, the output of every honest sensor $i \in \mathcal{N}$ must be at most a distance α from the target's original location \mathcal{L} ; i.e., $\|l_i - \mathcal{L}\|_2 \leq \alpha$.*

Soundness: *The output of each honest sensor $i \in \mathcal{N}$ must be \perp if there is no target in the area.*

We consider two different sensor models based on which sensors measure their input values v_i : (a) Boolean coverage model, and (b) Annular disk coverage model. We describe more details about these sensor models in Section III and Section IV, respectively.

B. Preliminaries

Asynchronous Approximate Agreement. We describe the asynchronous approximate agreement primitive as defined in Dolev *et al.* [20]. We consider a *relaxed* Validity condition characterized by the parameter δ .

Definition II.2. *A protocol Π_{AA} to which sensors input values V_i and output values w_i is a valid approximate agreement protocol if it guarantees the following properties.*

- 1) **Termination:** *All honest sensors must eventually decide and output a value.*
- 2) **ϵ -agreement:** *For a given $\epsilon > 0$, the decision values of any pair of honest sensors i and j are within ϵ of each other; i.e., $|w_i - w_j| \leq \epsilon$.*
- 3) **δ -relaxed Validity:** *The decision value of every honest sensor i must be within the δ -relaxed range of initial inputs of honest sensors \mathcal{M} ; i.e., $\min \mathcal{M} - \delta \leq w_i \leq \max \mathcal{M} + \delta$.*

Voronoi diagrams. We formally define Voronoi diagrams, as a key component of our solution. A n -order Voronoi Diagram is a geometric construct that divides the area into closed

polygonal regions or *cells*. All points within a cell have the same closest n sensors, which characterize the cell. For a given subset of n sensors \mathcal{N}_i , the Voronoi cell $A_i = f_{\text{VORO}}(A, \mathcal{N}_i, \mathcal{N})$ is defined as the locus of all points in the area whose first n closest sensors are the set \mathcal{N}_i .

$$\begin{aligned} A_i &= f_{\text{VORO}}(A, \mathcal{N}_i, \mathcal{N}) \\ &= \{p \in A \mid \forall s \in \mathcal{N}_i, \forall t \in \mathcal{N} \setminus \mathcal{N}_i : \|p - s\|_2 < \|p - t\|_2\} \end{aligned} \quad (1)$$

A is the area over which targets can appear and \mathcal{N} is the set of all deployed sensors. The Voronoi Diagram is the set of all non-empty Voronoi cells $A_i \neq \emptyset$, where each A_i is the Voronoi cell for a different subset of sensors \mathcal{N}_i . For example, Fig. 4 shows Voronoi cells of order 4.

$$\begin{aligned} f_{\text{VD}}(A, \mathcal{N}) \\ &= \{A_1, \dots, A_m\} : A_i \neq \emptyset \text{ and } \cup_{i \in \{1, \dots, m\}} A_i = A \end{aligned}$$

For a polygonal area A defined by points and edges, each Voronoi cell will also be a polygon.

III. LOCALIZATION IN THE BOOLEAN COVERAGE MODEL

For ease of understanding, we describe our solution in two stages: (a) We first solve the problem in the simpler boolean coverage model, and (b) Use this solution as a stepping stone to our solution in the annular disk model using an efficient approximate agreement protocol.

Boolean coverage model. In this model, a sensor measures 1 if a target is within distance r from the sensor and measures 0 otherwise. r is the sensing range of the sensor. In practice, this model is used only when the signal strength of the target is fixed and known to every sensor [48]. However, for ease of understanding, we first present a solution for this simpler problem where the target is assumed to be of known signal strength $S = S_{\min}$. We then discuss a method to extend this solution to the annular disk model where targets of unknown signal strengths are supported.

A. Voronoi Diagrams

Prior localization protocols [16], [38], [48] gave faulty sensors a monopoly in decision-making over multiple cells in the area, which allowed them to unilaterally create an alternate view of the state of these cells. We remedy this by assigning coverage responsibility of cells to subsets of n sensors.

We divide the entire area into disjoint pieces using a n -order Voronoi diagram (VD). As defined in Section II-B, the VD divides the area into subareas called Voronoi cells (VCs). A given cell is a collection of all points with the same n closest sensors. For each VC A_i and its corresponding set of n closest sensors \mathcal{N}_i , we create a shard $s_i = (A_i, \mathcal{N}_i)$, where we assign the responsibility of covering the area A_i to the set \mathcal{N}_i . We note that the VD generation and shard creation is done just once at the beginning by a trusted party like a network designer who knows the location of every sensor.

Each set \mathcal{N}_i of n sensors must behave as a single cohesive unit equivalent to a single honest sensor covering area A_i .




Figure 4: **4-Order VD**: Graphic for an area divided into cells A_i via a 4-order VD where every shard $s_i = (A_i, \mathcal{N}_i)$ contains $n = 4$ sensors. The blue circle denotes the sensing range of the sensor 1. The red sensors are Byzantine. A Voronoi order of $n = 4$ ensures a supermajority of honest sensors in all shards.

This requires the following two properties to be satisfied: (a) Non-equivocation: Any two external sensors accepting a value from shard s_i regarding area A_i must accept the same value, and (b) Validity: If there is a target within A_i , then shard s_i covering A_i must output 1. Further, if the target is not within the sensing range of any sensor in s_i , then s_i must output 0.

Protocol description. We achieve these conditions by first ensuring each shard has a supermajority of honest sensors. We utilize this supermajority to enable sensors to run an intra-shard Binary Byzantine Agreement protocol [2], with the n sensors in set \mathcal{N}_i as participants. Each sensor inputs 1 to the BA protocol if it detects a target in its range and 0 otherwise. Upon termination, each honest sensor in \mathcal{N}_i multicasts its output. Every other sensor in the network waits until accepting $\lfloor \frac{n}{3} \rfloor + 1$ equal outputs from sensors in \mathcal{N}_i , and accepts this value as a representative of the area A_i . Further, as part of the final aggregation phase, each sensor also collects outputs from all shards and filters shards with output 1. It then calculates the average of the *centroids* of these shards and outputs this value as the target's location.

We describe why the area A_i covered by sensors in shard s_i must be within the sensing range r of all sensors in s_i . We show that even if one sensor does not cover the entire area A_i , then the adversary can force the BA instance among sensors in s_i to always output 0.

Theorem III.1. *Consider a point p in A . Let n_r denote the number of sensors within distance r from p . Let N_x be the first x closest sensors to point p . If honest sensors do not have a supermajority in any $N_x : x \in \{4, \dots, n_r\}$, then an asynchronous BA instance involving N_x always terminates with output 0.*

Proof. We know that each set $N_x : x \in \{4, \dots, n_r\}$ comprising of the x closest sensors to the target does not have a supermajority of honest sensors. Let N_{SM} be the first set of y closest sensors in which honest sensors have a supermajority. We know that at least one honest sensor in N_{SM} cannot detect the target and measures 0 as its input. Summarizing the

scenario, N_{SM} is a set of sensors with an honest supermajority running a BA protocol. One honest sensor i in N_{SM} inputs $V_i = 0$ to the protocol and remaining honest nodes input 1. The adversary controls $< \frac{1}{3}$ fraction of nodes in N_{SM} and also the message delivery between honest nodes. Further, the adversary also knows the location of the target \mathcal{L} and by extension, the measured input of each honest node V_i .

In this scenario, we prove that it is impossible for any asynchronous agreement protocol (either exact or approximate agreement protocols) running among N_{SM} to terminate with output > 0 . By way of proving this statement, we show that The f faulty nodes in N_{SM} set their inputs to be 0. Together with honest node i , $f + 1$ nodes in N_{SM} input 0 to the BA protocol. We denote this set of nodes by $N_{SM,0}$. While executing the protocol, the adversary picks f remaining honest nodes which input 1, denoted by set $N_{SM,1}$. The combined set $N_{SM,0} \cup N_{SM,1}$ is of size $n - f = 2f + 1$. The adversary first delivers the messages between these $n - f$ nodes and delays all other messages. As the protocol is asynchronous, each honest node can only expect messages from $n - f$ nodes. Hence, the adversary forces the protocol to terminate with the input values from the $n - f$ nodes in set $N_{SM,0} \cup N_{SM,1}$. In this set, the value 0 is the only valid value that drives consensus while keeping the output within the range of honest inputs. Hence, the protocol must converge at value 0 and all honest nodes must agree on 0 as the representative value of the cell/area. \square

Therefore, from this theorem, we show that the sensors in each shard s_i must fully cover their assigned area A_i (i.e. the entire area A_i must be within distance r from sensors in s_i). **Estimating Voronoi order n .** For the described protocol to work, each shard must have a supermajority of honest sensors. As we do not know which sensors the adversary has compromised, we empirically estimate n for a static adversary who can corrupt at most f sensors in the entire network. In a given sensor deployment with N sensors, a n -order Voronoi Diagram contains $\mathcal{O}(n(N - n))$ VCs [3]. Using this information, we derive the probability that all VCs have a supermajority of non-faulty sensors. We perform asymptotic analysis for large values of N using the approach suggested in Kogias *et al.* [30] and Algorand [24], where the random variables for committee selection approach the binomial distribution. Given that shard s_i is composed of n sensors and the overall fraction of faulty sensors is $c = \frac{T}{N}$, the probability that more than $\frac{1}{3}$ rd sensors are faulty is given by Chernoff's bound.

$$\Pr[\bar{Y}_i] \leq e^{-\frac{(1-3c)^2 n}{3(1+3c)}}$$

Using this information, we derive the probability that all VCs have a supermajority of non-faulty sensors.

$$\Pr\left[\bigwedge_{i=1}^m Y_i\right] \geq 1 - N^2 e^{-\frac{(1-3c)^2 n}{3(1+3c)}}$$

For a Voronoi order of $n = \mathcal{O}(\eta + \log(N))$, this probability is $p = 1 - \frac{1}{2^{\eta}}$. For a constant $T < \frac{N}{3}$, a Voronoi order $n = \mathcal{O}(\text{polylog}(N))$ is enough to guarantee with a high

probability that all VCs in the n -order deployment will contain a supermajority of non-faulty sensors. For achieving this $n = \text{polylog}(N)$, SENSORBFT tolerates suboptimal number of faults of $T < \frac{N}{3} - \gamma$, where $\gamma > 0$ is a small constant.

Coverage. We note that all cells assigned to a given sensor must be within the sensing range r of the sensor for SENSORBFT to conduct localization. Finding the optimal deployment with the least number of sensors for this condition to be true reduces to the set-cover problem, which is NP-Hard. Hence, we assume that the created cells are within the sensing range of the sensors covering them.

Example deployment. A regular tessellation is the covering or tiling of an area without gaps or overlaps, using a regular geometric shape or tile. We choose a square tessellation of side length a . We draw a n -order Voronoi diagram and calculate the side length a of the tessellation to ensure all the n sensors in s_i detect the target if it is located in area A_i . This condition ensures the Validity of the output of BA from each shard s_i . We achieve this condition by ensuring that the entire subarea A_i is within the sensing range r of each sensor in \mathcal{N}_i . We choose lowest a where this condition is true for all shards.

B. Localization using Approximate Agreement

The described localization protocol uses Asynchronous Byzantine Agreement as a building block, which is computationally expensive because of using common coins. We overcome this bottleneck by utilizing the Approximate Agreement primitive, which does not use any energy-intensive cryptographic primitives. We employ the Binary Approximate Agreement or BINAA in [5], where honest sensors' inputs form a set of size at most two. This protocol achieves AAA in $\mathcal{O}(\frac{1}{\epsilon})$ rounds, with only $\mathcal{O}(n^2)$ communication complexity per round.

Algorithm 1 SENSORBFT localization protocol

```

1: INPUT:
2: Sensor input  $V_i$ , Precision  $\rho$ 
3: Set of all shards  $C = \{s_1, \dots, s_v\}$ 
4: OUTPUT: Target's location  $\mathcal{L}$ 

5:  $\forall y : \{1, \dots, |C|\}, cl_{i,y} \leftarrow \{\}$ 
   // Run BINAA for all shards in which sensor  $i$  is a member.
6: Run BINAA with input  $V_i \forall j : i \in \mathcal{N}_j$ 
   // Outputs of BINAA in  $\{s_1, \dots, s_v\}$ , of which  $i$  is a member
7:  $Y_i = \{w_{1,i}, w_{2,i}, \dots, w_{v,i}\}$ 
8:  $SendAll(Y_i, N)$  // Send  $Y_i$  to all sensors in  $N$ 
9: upon receiving  $Y_j$  from sensor  $j$ :
10:  $cl_{i,y} \leftarrow cl_{i,y} \cup \{w_{y,j}\}, \forall w_{y,j} \in Y_j$ 
11: upon  $\forall s_y \in C$ , if  $|cl_{i,y}| \geq \lceil \frac{2n}{3} \rceil$ :
12:  $c_y \leftarrow Centroid(A_y);$  // Shard Centroid
13:  $w_{i,y} \leftarrow Median(cl_{i,y})$ 
14: if  $\sum_{y:s_y} w_{y,i} \neq 0$  then
15: OUTPUT:  $\mathcal{L} \leftarrow \frac{\sum_{y:s_y} w_{y,i} c_y}{\sum_{y:s_y} w_{y,i}}$ 
16: else
17: OUTPUT:  $\mathcal{L} \leftarrow \perp$ 

```

The prior method of aggregating outputs from shards, where nodes filter shards with output 1 and calculate the average of

their centroids, will not work when Approximate Agreement is used for intra-shard consensus. This is because a specific sensor i might accept 1 as an output from a specific shard s_j and another sensor might accept $1 - \epsilon$ from s_j . In this case, i will include s_i , and j will not include s_i , which violates the precision property of localization. Further, agreeing on a specific order of shards by their weight with a deterministic protocol like BINAA violates the FLP impossibility result.

Weighted average. We address this limitation by proposing an aggregation method based on weighted averaging. This technique allows us to extend the approximate consensus on the sensor input readings to localization, resulting in approximate localization with a precision parameter ρ . In this protocol described in Algorithm 1, each shard member in set \mathcal{N}_i multicasts its output from BINAA and every other sensor in the network waits until receiving $\lceil \frac{2n}{3} \rceil$ messages from the shard. Then, each sensor j calculates the median of received values and assigns it as the *representative value* of the shard s_i , denoted by $w_{i,j}$. Each sensor then calculates the weighted average of shard centroids where the weight of the centroid of shard s_i is the value $w_{i,j}$, and outputs it as the target's location.

We prove the precision and accuracy properties of the protocol in Algorithm 1. We first show that shards far away from the target have zero weight in localization.

Lemma III.2. *Every shard s_y with area A_y such that every point in A_y is at a distance $d > 2r$ from the target, where r is the sensing range of each sensor, will have a representative weight $w_{y,i} = 0$, for all honest sensors i in the network.*

Proof. Each sensor covering A_y must be at a distance $> r$ from the target, which implies each such sensor will input 0 to all corresponding BINAA instances. From the Validity of BINAA, we say that the output $w_{y,j} = 0$ for all honest sensors j in the network. \square

Lemma III.3. *Given the sensor network is deployed over an area A with each sensor's sensing range r , the target localization scheme in Algorithm 1 with ϵ in BINAA, $\epsilon = \mathcal{O}(\frac{\rho A}{nr^3})$, achieves a localization precision of ρ meters.*

Proof. Every sensor in the network received at least $\lceil \frac{2k}{3} \rceil$ responses from every shard s_i in the network. We consider the weighted average calculated by sensor i .

$$\mathcal{L}_i = \frac{\sum_{y:s_y \in C} w_{y,i} c_y}{\sum_{y:s_y \in C} w_{y,i}}$$

From the ϵ -agreement property of approximate agreement primitive, we have $|w_{y,j} - w_{y,i}| \leq \epsilon \forall i, j \in N, s_y \in C$ where N is the set of all sensors and C is the set of all shards. For every pair of sensors i and j , the difference between outputs can be upper bounded by the following expression.

$$\mathcal{L}_j - \mathcal{L}_i \leq \frac{\sum_{y:s_y \in C} w_{y,j} c_y}{\sum_{y:s_y \in C} w_{y,j}} - \frac{\sum_{y:s_y \in C} w_{y,i} c_y}{\sum_{y:s_y \in C} w_{y,i}}$$

Without loss of generality, we write $w_{y,j} = w_{y,i} + \epsilon_y$, where $|\epsilon_y| \leq \epsilon$. We simplify the expression to the following.

$$\mathcal{L}_j - \mathcal{L}_i \leq \frac{\sum_{y:s_y \geq 2C} \epsilon_y (\mathcal{L}_j - c_y)}{\sum_{y:s_y \geq 2C} w_{y,i}}$$

From Lemma III.2, we know that the quantity $\mathcal{L}_j - c_y \leq 2r$. Moreover, as there exists at least one shard s_y that will always output $w_{y,i} = 1$, the quantity $\sum_{y:s_y \geq 2C} w_{y,i} \geq 1$ for all honest sensors i . Therefore, we simplify the expression to the following.

$$\mathcal{L}_j - \mathcal{L}_i \leq r \sum_{y:s_y \geq 2C} \epsilon_y$$

The sum of ϵ_y values depends on the number of Voronoi cells within an area formed by a circle of radius $2r$. We denote this quantity using $n^0 = \mathcal{O}(\frac{nr^2}{A})$. Therefore, for $\epsilon = \mathcal{O}(\frac{\rho A}{nr^3})$, the protocol achieves a localization precision of ρ meters. \square

We also show that the protocol in Algorithm 1 achieves an accuracy of $\alpha \leq \frac{3r}{2}$ meters.

Lemma III.4. *Given the sensor network is deployed over an area A with each sensor's sensing range r , the target localization scheme in Algorithm 1 achieves a localization precision of $\alpha \leq \frac{3r}{2}$ meters.*

Proof. From Lemma III.2, we know that no point at a distance greater than $2r$ can have a positive weight in localization. Consider one such point c_x with weight $w_{x,i} = 1$. We also know that the shard containing the target, s_y , must have a weight $w_{y,j} = 1$, whose centroid c_y is at most at a distance r from the target. Therefore, the difference between output location of an honest sensor i and the target's original location $|\mathcal{L}_i - \mathcal{L}| \leq \frac{c_x + c_y}{2}$, which implies $|\mathcal{L}_i - \mathcal{L}| \leq \frac{c_x + c_y}{2} \leq \frac{3r}{2}$. \square

Complexity Analysis. The BINAA protocol consumes $\mathcal{O}(n^2)$ (n is the Voronoi order) per round per shard. Each BINAA instance runs for $\mathcal{O}(\log(\frac{Nr^3}{\rho A}))$ rounds and there exist $\mathcal{O}(n(N-n))$ Voronoi cells. Further, multicasting the decisions of BINAA instances of shards is $\mathcal{O}(Nn)$ bits, as opposed to $\mathcal{O}(N^2n)$ bits. This is because each shard can summarize the weight of its neighboring shards with a single message, which is possible because of the supermajority of honest sensors within each shard. This gives us $\mathcal{O}(Nn^3 \log(\frac{Nr^3}{\rho A}))$ bits of communication overall, and $\mathcal{O}(n^3 \log(\frac{Nr^3}{\rho A}))$ bits per sensor. For large N and $n = \text{polylog}(N)$, the complexity is $\mathcal{O}(N \text{polylog}(N))$ bits and $\mathcal{O}(\text{polylog}(N))$ bits of communication per sensor.

IV. LOCALIZATION IN ANNULAR DISK MODEL

Sensor model. A sensor i deployed at a location $i.l = (x_i, y_i)$ measures the target's signal footprint at its location along with some ambient noise characterized by the random variable η . For our analysis, we utilize acoustic sensors for conducting

localization. The variable η in acoustic sensors is characterized as a lognormal distribution, which is a thin-tailed distribution.

$$v_i = S \times f_{\text{DIST}}(\|i.l - \mathcal{L}\|_2) + \eta$$

Sensing range. A sensor's ability to distinguish a target's signal from the noise decreases with increasing distance between them. For a given target of signal S , we assume each sensor has a sensing range $r = f_{\text{SR}}(S)$ beyond which the target's signal is indistinguishable from noise. We denote $r_{\min} = f_{\text{SR}}(S)$ as the sensing range for a target with $S = S_{\min}$.

Under this model, a target's signal strength is unknown, but $\geq S_{\min}$. A sensor can detect a target from varying distances ranging from r_{\min} (when $S = S_{\min}$). We start with a sensor network deployed with sensing range $r = r_{\min}$, using BINAA for an agreement where a sensor inputs 1 if it detects the target. This system can localize any target with $S \geq S_{\min}$ with high precision. However, when S is high, this network and the protocol localize the target with very low accuracy.

The core problem causing this low accuracy is the shards at greater distances, which measure a much lower signal footprint from the target, are contributing to localization with the same weight as a shard that is much closer to the target. This is mainly because the BINAA protocol requires the inputs to be binary, which causes information loss about measured signal strengths. Other protocols like Abraham *et al.* [1] achieve approximate agreement with convex-hull Validity (Definition II.2 with $\delta = 0$) over real-valued inputs. However, this protocol is highly energy-intensive because of its high message complexity of $\mathcal{O}(n^3 \log(\frac{1}{\epsilon}))$ over $\log(\frac{1}{\epsilon})$ rounds. Δ is the initial range of honest inputs.

A. Approximate Agreement with Relaxed Validity

We remedy this problem by proposing a hybrid protocol between Abraham *et al.*'s [1] ϵ -agreement protocol and the BINAA protocol. The protocol offers a tradeoff between communication cost (and effectively, energy consumption) and the relaxation in Validity (and effectively, accuracy loss) from discretization, characterized by δ . We incorporate a new technique based on value rounding to effectively trade energy consumption for localization accuracy, which sensors can use as a tuning knob to meet their energy budgets. We describe the protocol in Algorithm 2.

Algorithm 2 SENSORBFT-hybrid: (ϵ, δ) -agreement protocol

- 1: **INPUT:** V_i, ϵ, δ

- 2: Run Abraham *et al.* [1] with V_i and $\epsilon_{ab} = 2\delta$
- 3: Round off the output $w_{ab,i}$ to the closest multiple of 2δ , obtaining $V_{bin,i}$
- 4: Run BINAA with inputs $V_{ab,i}, \epsilon_{bin} = \epsilon, \Delta = \delta$
- 5: **OUTPUT:** $w_{bin,i}$, the output of BINAA

In this protocol, sensors run Abraham *et al.* with $\epsilon_{ab} = 2\delta$. Upon terminating this protocol, sensors round off the output $w_{ab,i}$ to the closest *checkpoint*, which is an integer multiple of




Figure 5: Description of the Validity relaxation because of rounding. The range of values before rounding and after rounding are denoted by the top and bottom intervals respectively. This rounding results in a relaxation of Validity by δ .

2δ . Then, sensors run the BINAA protocol with $\epsilon_{bin} = \epsilon$ and $\Delta = 2\delta$, and output the corresponding result $w_{bin,i}$. As the range of honest outputs $w_{ab,i}$ from Abraham *et al.* is always $\leq 2\delta$, the rounding operation ensures that the inputs to BINAA are binary, and are at most at a distance δ from the interval formed by $w_{ab,}$ values. As a result, the hybrid protocol satisfies Termination, ϵ -agreement, and δ -relaxed Validity.

Validity and Complexity analysis. The round operation results in a Validity relaxation of δ . We illustrate this tradeoff in Fig. 5. This hybrid protocol has a message complexity of $\mathcal{O}(n^3 \log(\frac{1}{\delta}) + n^2 \log(\frac{\delta}{\epsilon}))$ messages per shard. For a constant percentage fraction $\delta = \mathcal{O}(\Delta)$, protocol has a complexity of $\mathcal{O}(n^3 + n^2 \log(\frac{1}{\epsilon}))$, and has only $\mathcal{O}(n^2)$ dependence on ϵ . We show in Section V that for $\frac{1}{\delta} = 6$, the localization accuracy parameter α increases by at most 12%, which suggests that $\mathcal{O}(n^2)$ dependence on ϵ is true for practical situations.

B. Localization protocol

We use the same protocol as in Algorithm 1 for aggregating weights of cells. Say a target appears inside a cell A_i . The honest sensors composing the shard record signal strength from the target, and input it to the hybrid agreement instance. The Validity property in Definition II.2 ensures that the representative value $w_{i,}$ of shard s_i has a higher weight in the weighted average than any other shard that is farther away from the target. For example, in Fig. 4, if a target appears in A_1 , the non-faulty sensors composing s_1 will measure a higher signal than a farther away shard $s_j : \{1, 3, 7\} \notin \mathcal{N}_j$, which implies that $\min(\mathcal{M}_1) > \max(\mathcal{M}_4)$. As long as $\delta \leq \frac{\min(\mathcal{M}_1) - \max(\mathcal{M}_j)}{2}$, s_1 will have a higher weight than s_4 in the weighted average. Hence, this configurable accuracy introduced by a non-zero δ in (ϵ, δ) -agreement transforms into an increased α . By using δ as a tuning knob, the sensors can control α and the corresponding energy spent.

We can also normalize the $w_{i,}$ values and amplify the difference between weights of shards by using weight $w_{i,}^0 = e^{w_{i,}}$ in the average to give more weight to shards that are closer to the target. This normalization ensures that the weighted average is centered around the cell A_i with the largest representative value $w_{i,}$.

Evaluation setup. We evaluate SENSORBFT by setting up an embedded testbed consisting of 19 Raspberry Pi 4-B devices connected with the help of a network switch functioning as a router. We configure each device only to turn on its high-power CPU and networking module when they want to run the localization protocol. At all other times, the device only uses a low-power sensor module to collect signal readings. For comparison, the sensor module only consumes 0.3 W of power, whereas the CPU and the networking module consume an average of 4 W while running the AAA protocol. We power the devices using a power supply and measure the energy consumption of a device by integrating the power drawn by the device over the time it takes to run the protocol. We measure the power consumption of each device by multiplying the current drawn from the power supply measured using an ammeter and the supply voltage. We measure the energy consumption in Joules (J), the standard unit of energy. As a benchmark, a conventional fully charged AA battery powering a sensor device at 3V has an energy capacity of 30000 Joules.

We implement SENSORBFT in Rust², touted to be one of the most energy-efficient languages [42]. Our protocol uses MACs based on shared symmetric keys and the SHA256 Hash function to authenticate messages between every pair of sensors.

Noise distribution and targets. We use acoustic sensors that measure the intensity of the sound signal at their location for localizing targets of various signal strengths in an indoor environment. First, we estimated the ambient noise distribution η by collecting data from an acoustic sensor in an outdoor environment without any targets. We find that it follows a lognormal distribution [26], with the specific parameters being $\ln(\eta) = \text{Normal}(\mu, \sigma^2)$ with $\mu = 52.5\text{dB}$ and $\sigma = 1$.

We also calculated the ground truth strength of the signal emitted by a sample target by recording the sound intensity at a sensor at a distance of $d_0 = 1\text{m}$ from the target. We chose the sample target as an internal combustion engine in a car, whose ground truth signal strength is 84 dB. We set this quantity as $S_{\min} = 84\text{dB}$ and the maximum possible signal strength of a target as $S_{\max} = 120\text{dB}$, which is the upper bound on the signal strength caused by an accident like an explosion or fire hazard. The sound intensity falls by 6 dB when the distance from the target doubles, with intensity at distance d being $S_d = S - 20 \log_{10}(d)$, where S is the ground truth signal strength at $d_0 = 1\text{m}$.

Comparison with prior work. We evaluate SENSORBFT against a localization approach where sensors in each shard use: (a) An asynchronous BA protocol outputting an Asynchronous Common Subset (ACS), with $\epsilon = 0$ and relaxation $\delta = 0$, (b) Abraham *et al.*'s [1] protocol with $\epsilon > 0, \delta = 0$. For the first approach, we choose FIN [21] as the ACS protocol.

Our choice of FIN is due to its low computational expense among all ACS protocols. FIN does not use any public key signatures and only uses $\mathcal{O}(\log(n))$ common coins per

²Available at <https://github.com/akhilbs/sensorbft-rs/>

instance. For comparison, BKR-style ACS protocols such as PACE [49], and HoneyBadgerBFT [39] need $\mathcal{O}(n)$ bits of randomness to terminate, whereas MVBA-style ACS protocols like Speeding Dumbo [25] use $\mathcal{O}(n^2)$ public key signature verification operations.

A. Energy consumption of approximate agreement

We benchmark our approximate agreement protocol for different values of ϵ and δ as defined in Definition II.2. We define relative measures of Agreement and Validity $\epsilon_{\text{REL}} = \frac{\epsilon}{\Delta}$, which is the measure of the relative improvement in the agreement precision, and the accuracy $\delta_{\text{REL}} = \frac{\delta}{\Delta + \delta}$ as a measure of the relative deviation of the output from the honest nodes' input interval. $\Delta = \max(\mathcal{M}) - \min(\mathcal{M})$ is the range or the initial level of difference between honest inputs. Higher values of ϵ_{REL} and δ_{REL} are better. We first show the variation in energy consumption for different ϵ and δ values (and corresponding ϵ_{REL} and δ_{REL}) in a heatmap for $n = 37$ sensors participating in the protocol.

In our testbed, comprising of 19 physical devices, we devised a method for emulating larger networks when the number of sensors participating in agreement, n , exceeds 19. To measure the energy consumption of a process, we run multiple processes on 18 devices, while dedicating one device to a single process. Given that each device communicates through a router, this method enables us to accurately estimate energy usage in the approximate agreement protocol for larger values of n . The dedicated device executes the same number of message exchanges, cryptographic operations, and CPU cycles as it would in a genuine multi-machine setup. This emulation technique, previously utilized in Narwhal [17], allows us to effectively evaluate SENSORBFT's scalability.

The heatmap in Fig. 6 shows that the energy consumption can be reduced without sacrificing agreement precision. Previously, energy savings in an agreement protocol implied compromising (moving down the y-axis with minimum Validity relaxation). In contrast, our method allows energy conservation by relaxing Validity (moving left on the x-axis), keeping the agreement precision intact. Notably, based on the color gradient in the y-axis, our protocol yields a higher level of agreement at a given level of Validity relaxation, more affordably than Abraham *et al.*. Whenever a minor Validity relaxation is permissible, our protocol emerges as the most energy-efficient solution to achieve high agreement precision among sensors. For instance, our protocol, at $\delta_{\text{REL}} = 60\%$ and $\epsilon_{\text{REL}} = 99.99\%$, consumes 4x less energy than Abraham *et al.* and 9x less energy than FIN.

Scalability results. In Fig. 7a, we plot the energy consumed by our protocol to achieve $\epsilon_{\text{REL}} = 99.93\%$, with increasing number of sensors n participating in an agreement instance. At $n = 163$, our protocol with $\delta_{\text{REL}} = 60\%$ requires just 510 Joules, compared to 925 Joules for Abraham *et al.* and 1100 Joules for FIN. This demonstrates our protocol's ability to deliver high precision with significantly reduced energy expenditure, by permitting a trade-off with accuracy.

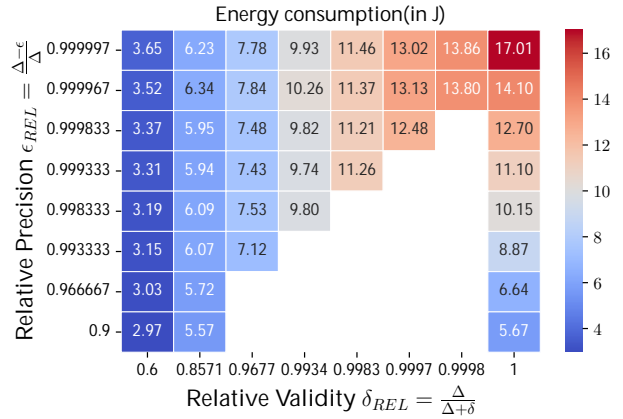


Figure 6: Energy map of Approximate Agreement: Heatmap of our agreement protocol illustrating the interplay between precision, $\epsilon_{\text{REL}} = \frac{\Delta - \epsilon}{\Delta}$, and accuracy, $\delta_{\text{REL}} = \frac{\Delta}{\Delta + \delta}$, for $n = 37$ processes. Each cell in the heatmap corresponds to a specific configuration of error tolerance and accuracy, colored according to the energy consumed as per the scale on the right. The column with maximum accuracy is Abraham *et al.* [1] protocol. Notably, the FIN protocol consumes 40.6 Joules of energy for $n = 37$ processes. The colormap ranges from blue (lower energy consumption) to red (higher energy consumption), indicating how energy efficiency varies with different protocol configurations.


Using a 30,000 J battery, a sensor device using our protocol can execute 3200 instances of approximate agreement with $n = 55$ sensors at $\delta_{\text{REL}} = 60\%$ and $\epsilon_{\text{REL}} = 99.93\%$. This contrasts with 1170 instances of ϵ -agreement for Abraham *et al.*, and only 340 instances of FIN.

We note that the relative share of energy consumed by computation is much higher than communication because of the geographic proximity of participant nodes. We observe a higher energy share is consumed by computation than communication, attributed to the close geographic proximity of participant nodes. In contrast, geo-distributed SMR applications witness higher energy consumption for communication, due to increased transmission distances. Given the low energy consumed for communication in our setup, our results also provide a practical lower bound on energy consumption of any agreement protocol offering a similar precision and accuracy from a distributed systems perspective.

B. Localization evaluation

We evaluate SENSORBFT's localization protocol and show the tradeoffs concerning the precision ρ and accuracy α with energy consumed by the protocol with the help of data from real acoustic sensors. We also compare our protocol to localization protocols using FIN and Abraham *et al.*'s protocols as agreement protocols in Algorithm 1.

We consider a rectangular area A and sensors deployed in a square tessellation with side length a . For a pre-estimated Voronoi order n , we estimate the lowest side length of the tessellation a for localizing targets of signal strengths between $S_{\text{min}} = 85\text{dB}$ and $S_{\text{max}} = 120\text{dB}$. We set a signal threshold



(a) **SENSORBFT's Scalability:** Energy consumption of our approximate agreement protocol with Abraham *et al.* (denoted by $\delta_{REL} = 100\%$) and FIN. We plot SENSORBFT for varying values of $\delta_{REL} = \frac{\Delta}{\Delta + \delta}$. Our approximate agreement protocol is significantly more energy efficient, consuming an order of magnitude less energy than FIN and Abraham *et al.*

(b) **Energy vs Precision ρ :** We plot the precision ρ and the corresponding energy consumption for systems with $n = 19, 37$ sensors per shard. SENSORBFT's energy usage remains stable even with lower ρ due to our more energy-efficient precision improvement. In return for this efficiency, SENSORBFT has a 2% higher α than Abraham *et al.* FIN, in contrast, consumes 140 J and 1500 J for $n = 19, 37$, respectively.

(c) **Energy vs Accuracy α :** The increase in localization accuracy parameter α vs. energy consumed for $\rho = 0.17m$, in comparison with Abraham *et al.* and FIN. A higher energy investment improves accuracy and reduces α . This improvement diminishes with increasing expenditure. Abraham *et al.*'s method consumes 115 J and FIN consumes 140 J for $\rho = 0.17m$. The minimum value is $\alpha_{min} = 20m$.

Figure 7: Evaluating SENSORBFT's Primitives

P such that the probability of only the noise generating a signal of P is highly improbable. Based on η 's parameters of $\mu = 52.5$ and $\sigma = 1$, we set $P = 58dB$. With a Voronoi order $n = 19$, we calculate the side length using this P value. The value for κ_c for Voronoi order $n = 19$ is $\kappa_c \leq 5$ [45], which gives us a side length of $a = 4.46m$, which implies sensors need to be placed at a distance of $a = 4.4m$ to be able to localize targets of $S_{min} = 90dB$.

Energy consumption vs Precision ρ . We study the variation in consumed energy with changing localization precision ρ in Fig. 7b. We plot the maximum energy spent by any sensor on the map with ρ , for localizing a target of signal strength $S = 120dB$. For a precision of $\rho = 0.17m$, our approach consumes 49.3% lesser energy than Abraham *et al.* for $n = 19$ and 44.3% lesser energy at $n = 37$, for a modest 2% loss in localization accuracy α . Our approach also can achieve a lower ρ at a practically insignificant energy cost because of our design choice to make agreement precision cheaper at the expense of a minor increase in α . FIN consumes $2.5\times$ and $5\times$ more energy than SENSORBFT to achieve $\rho = 0$.

Accuracy α vs Energy consumption. Using the relationship between δ and α (with constants adjusted for measurements in decibels), we plot the proportional increase in α (relative to the minimum α_{min}) against the maximum energy consumed by any sensor on the map, when localizing a target with $S = 120dB$. Recall that α_{min} is the lower bound on α achieved when an agreement protocol with maximum accuracy is used. We plotted Fig. 7c by varying the δ parameter in approximate agreement and then calculating the resultant increase in α based on the target's signal strength and signal attenuation rate. For instance, at $n = 19$, when setting $\delta = 1dB$, we observe a nominal 2% increase in α , coupled with an energy reduction surpassing 50% relative to Abraham *et al.* and over 70% against FIN. The minimum α_{min} in this case is $\alpha_{min} = \kappa_s * a = 22.63m$. This can be reduced by deploying

the sensors closer together, while simultaneously increasing the threshold P .

Energy consumption map. In Fig. 8, we analyze the trend of energy consumption throughout the system in the presence of the target. The graphic shows the energy consumed by each sensor running the localization protocol with a 2% increase in α compared to Abraham *et al.* and FIN, which have perfect accuracy. As the distance from the target increases, sensors consume lesser energy for localization because of the reduced difference in their measured readings. Sensors closer to the target measure a much higher difference in their readings. Our approach consumes at least 50% less energy than Abraham *et al.*, for only a 2% increase in α . The localization approach using FIN has a uniform energy consumption of 140 J per sensor throughout the area.

The advantage and energy-efficiency of SENSORBFT stem from its approximate agreement protocol, which capitalizes on the variation between sensor reading intervals from different shards. The wide difference between the value intervals of adjacent shards allows SENSORBFT to perform an agreement with low accuracy (equivalently, high δ), yet maintain a reasonably low α . In contrast, Abraham *et al.*'s ϵ -agreement protocol solely optimizes energy using intra-shard differences between sensor readings. FIN does not exploit any contextual information from sensor readings.

VI. RELATED WORK

Works on reliable aggregation focus on connecting the sensors to the FC for reliable message transmission [12], [23]. These works propose different routing protocols to connect to the FC in the presence of faulty sensors. However, these works also assume a synchronous network between sensors and the FC, which makes them vulnerable to intermittent network failures and other malicious behavior like network jamming.




Figure 8: **SENSORBFT’s Energy Consumption:** This map displays the energy expenditure for each sensor operating SENSORBFT, as they localize a target located at (65, 65). Energy use decreases with the distance from the target due to a reduction in the variance of sensor readings. α is marginally higher by 2% compared to the baseline using Abraham *et al.*, maintaining $\rho = 0.5\text{m}$, illustrating the trade-offs between precision, accuracy, and energy efficiency inherent in SENSORBFT’s design.

Sensor Fault Tolerance. Works in sensor fault tolerance rely on an FC (Fusion Center) to either detect faulty sensors [14], [43] or to run fault-tolerant aggregations on received data. For example, specific approaches based on hypothesis testing assume non-faulty sensor inputs are from a probability distribution and utilize the FC to compute scores of honest behavior [43]. Other works use specific metrics from the sensor device, which enables the FC to identify faults [14].

Distributed Sensor Networks. Some works moved away from a centralized FC by implementing distributed consensus algorithms [28], [34], [37]. For example, SENATE [28] uses a Proof of Location algorithm to authenticate sensors and a consensus protocol to eliminate faulty sensor values. However, all works in this category assume network synchrony, making them infeasible for low-power embedded systems.

Target Localization. Many works proposed target localization protocols with information fusion from various sensors [7], [8], [29], [36], [40], [48] and specific fault-tolerant fusion methods [32]. While the first set of works cannot tolerate faults, the second set use an FC and a synchronous network. Although there exist recent advanced localization techniques that utilize time difference between signal arrival [40], frequency of the target signal, and 2D azimuth change maps [29], [36], it is not clear how these works perform amidst a single faulty sensor. Further, these works are feasible for use only in indoor controlled environments where the network is synchronous and faulty behavior is improbable.

Consensus in sensors. Prior works explored distributed consensus in resource-restricted devices with paltry infrastructure, where any form of network synchrony is a strong assumption [1], [9], [20], [41]. Bhat *et al.* [9] show that public key primitives like signatures are energy-intensive and inefficient

in embedded devices, and propose EESMR, an energy-efficient synchronous SMR protocol. However, asynchronous consensus requires common coins, whose most efficient implementation is threshold unique signatures [10], [13]. On top of being energy intensive, threshold signatures also require a threshold setup, which is very expensive to deploy in sensors [19], [31]. Although a concurrent coin protocol titled HashRand [6] overcomes this computational bottleneck using Hash functions, it has a high $\mathcal{O}(n^3)$ communication complexity in a one-time execution. These factors motivated multiple new asynchronous consensus primitives that do not require common coins.

Asynchronous Approximate Agreement (AAA) is a primitive where nodes sacrifice agreement precision in exchange for coin-free termination [1], [20]. However, even this primitive has a high communication cost [1]. A concurrent work Delphi [5] achieves AAA with $\tilde{\mathcal{O}}(n^2)$ complexity but has a minimum Validity relaxation of $\delta = \Delta$, which is too high for localization. Moniz *et al.* [41] also proposed a different primitive called k -consensus, which requires only k honest nodes to terminate the protocol. This primitive is orthogonal to the approximate agreement primitive.

VII. CONCLUSION

We presented SENSORBFT, the first asynchronous energy-efficient distributed target localization protocol that can tolerate Byzantine-faulty nodes. It has two tuning knobs, controlling precision and accuracy of localization, which enable sensors to meet their energy budgets. We introduced two building blocks: (i) Voronoi sharding, which distributes coverage responsibility amongst sensors using higher-order Voronoi diagrams to cover a larger area in the presence of Byzantine faulty nodes; (ii) approximate (ϵ, δ) -agreement, which allows the sensors to achieve high localization precision at a fractional energy expense for a practically insignificant decrease in accuracy. For only a 2% increase in localization error at $n = 19$ sensors, SENSORBFT consumes only $\frac{2}{5}$ th of the energy consumed by the FIN protocol, while incurring only a minor 2% loss in localization accuracy, thus significantly enhancing operational efficiency.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers for the helpful feedback. We thank the Purdue League of Robotic Engineers (LoRE) and its president Mythra Balakuntala for their help with assembling the Raspberry Pi testbed. We also thank Andreas Pollak for providing the Java code to calculate and visualize higher-order Voronoi diagrams in 2D. This work was supported in part by NIFA award number 2021-67021-34252, the National Science Foundation (NSF) through grants numbered CNS-2146449, CNS-1846316, CNS-2038566, and CNS-2038986, and the Army Research Lab under Contract W911NF-2020-221. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

REFERENCES

- [1] I. Abraham, Y. Amit, and D. Dolev. Optimal resilience asynchronous approximate agreement. *OPODIS'04*, page 229–239, 2004.
- [2] I. Abraham, N. Ben-David, and S. Yamamuri. Efficient and adaptively secure asynchronous binary agreement via binding crusader agreement. *PODC'22*, page 381–391, 2022.
- [3] K.P. Agarwal, M. De Berg, J. Matousek, and O. Schwarzkopf. Constructing levels in arrangements and higher order voronoi diagrams. *SIAM journal on computing*, 27(3):654–667, 1998.
- [4] S. Bagchi, T.F. Abdelzahr, R. Govindan, P. Shenoy, A. Atrey, P. Ghosh, and R. Xu. New frontiers in iot: Networking, systems, reliability, and security challenges. *IEEE IoT Journal*, 7(12):11330–11346, 2020.
- [5] A. Bandarupalli, A. Bhat, S. Bagchi, A. Kate, C. Liu-Zhang, and M. Reiter. Delphi: Efficient asynchronous approximate agreement for distributed oracles, 2024. To appear at DSN 2024.
- [6] A. Bandarupalli, A. Bhat, S. Bagchi, A. Kate, and M. Reiter. Hashrand: Efficient asynchronous random beacon without threshold cryptographic setup. *Cryptology ePrint Archive*, Paper 2023/1755, 2023. To appear at CCS 2024.
- [7] A. Bandarupalli, S. Jain, A. Melachuri, J. Pappas, and S. Chaterji. Vega: Drone-based multi-altitude target detection for autonomous surveillance. In *2023 DCOSS-IoT*, pages 209–216, 2023.
- [8] A. Bandarupalli, D. Swarup, N. Weston, and S. Chaterji. Persistent airborne surveillance using semi-autonomous drone swarms. *Dronet'21*, page 19–24, New York, NY, USA, 2021. Association for Computing Machinery.
- [9] A. Bhat, A. Bandarupalli, M. Nagaraj, S. Bagchi, A. Kate, and M. K. Reiter. Eesmr: Energy efficient bft-smr for the masses. In *MIDDLEWARE'23*, pages 1–14, 2023.
- [10] A. Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-diffie-hellman-group signature scheme. In Yvo G. Desmedt, editor, *PKC 2003*, pages 31–46, 2002.
- [11] D. Boneh, B. Lynn, and H. Shacham. Short signatures from the weil pairing. In *ASIACRYPT 2001*, pages 514–532. Springer, 2001.
- [12] J.L. Bredin, E.D. Demaine, M.T. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed fault tolerance. *IEEE/ACM Transactions on Networking*, 18(1):216–228, 2010.
- [13] C. Cachin, K. Kursawe, and V. Shoup. Random oracles in constantipole: practical asynchronous byzantine agreement using cryptography. In *PODC'00*, pages 123–132, 2000.
- [14] T. Chakraborty, A.U. Nambi, R. Chandra, R. Sharma, M. Swaminathan, Z. Kapetanovic, and J. Appavoo. Fall-curve: A novel primitive for iot fault detection and isolation. *SenSys '18*, page 95–107, 2018.
- [15] C. Cheng and C. Hsu. The deterministic sensor deployment problem for barrier coverage in wsns with irregular shape areas. *IEEE Sensors Journal*, 22(3):2899–2911, 2021.
- [16] T. Clouqueur, K.K. Saluja, and P. Ramanathan. Fault tolerance in collaborative sensor networks for target detection. *IEEE Transactions on Computers*, 53(3):320–333, 2004.
- [17] G. Danezis, L. Kokoris-Kogias, A. Sonnino, and A. Spiegelman. Narwhal and tusk: A dag-based mempool and efficient bft consensus. *EuroSys '22*, page 34–50, 2022.
- [18] K. Danilchenko, Z. Nutov, and M. Segal. Covering users with qos by a connected swarm of drones: Graph theoretical approach and experiments. *IEEE/ACM Transactions on Networking*, pages 1–16, 2022.
- [19] S. Das, T. Yurek, Z. Xiang, A. Miller, L. Kokoris-Kogias, and L. Ren. Practical asynchronous distributed key generation. In *IEEE S&P '22*, pages 2518–2534. IEEE, 2022.
- [20] D. Dolev, N.A. Lynch, S.S. Pinter, E.W. Stark, and W.E. Weihl. Reaching approximate agreement in the presence of faults. *J. ACM*, 33(3):499–516, may 1986.
- [21] S. Duan, X. Wang, and H. Zhang. Practical signature-free asynchronous common subset in constant time. In *CCS*, pages 815–829, 2023.
- [22] M.J. Fischer, N.A. Lynch, and M.S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, page 374–382, 1985.
- [23] A.A. Fröhlich, R.M. Scheffel, D. Kozhaya, and P.E. Verissimo. Byzantine resilient protocol for the iot. *IEEE IoT Journal '19*, pages 2506–2517.
- [24] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements for cryptocurrencies. *SOSP '17*, page 51–68, 2017.
- [25] B. Guo, Y. Lu, Z. Lu, Q. Tang, J. Xu, and Z. Zhang. Speeding dumbos: Pushing asynchronous BFT closer to practice. *NDSS '22*, 2022.
- [26] A.P. Hill, P. Prince, E. Piña Covarrubias, C.P. Doncaster, J.L. Snaddon, and A. Rogers. Audiomoth: Evaluation of a smart open acoustic device for monitoring biodiversity and the environment. *Methods in Ecology and Evolution*, 9(5):1199–1211, 2018.
- [27] P. Hoyingcharoen and W. Teerapabkajorndet. Expected probabilistic detection and sink connectivity in wireless sensor networks. *IEEE sensors journal*, 19(12):4480–4493, 2019.
- [28] Z. Jiang, Z. Cao, B. Krishnamachari, S. Zhou, and Z. Niu. Senate: A permissionless byzantine consensus protocol in wireless networks for real-time internet-of-things applications. *IEEE IoT Journal*, 7(7):6576–6588, 2020.
- [29] Mahmut Karakaya and Hairong Qi. Collaborative localization in visual sensor networks. *ACM Trans. Sen. Netw.*, 10(2), jan 2014.
- [30] E. Kokoris-Kogias, P. Jovanovic, L. Gasser, N. Gailly, E. Syta, and B. Ford. Omniledger: A secure, scale-out, decentralized ledger via sharding. In *2018 IEEE S&P*, pages 583–598. IEEE, 2018.
- [31] E. Kokoris-Kogias, D. Malkhi, and A. Spiegelman. Asynchronous distributed key generation for computationally-secure randomness, consensus, and threshold signatures. In *CCS '22*, pages 1751–1767, 2020.
- [32] C. Laoudias, M.P. Michaelides, and C.G. Panayiotou. fitrack: fault-tolerant target tracking in binary sensor networks. *TOSN*, 10(4):1–28, 2014.
- [33] M. Li, I. Koutsopoulos, and R. Poovendran. Optimal jamming attacks and network defense policies in wireless sensor networks. In *IEEE INFOCOM*, pages 1307–1315, 2007.
- [34] S. Li, S. Zhao, P. Yang, P. Andriotis, L. Xu, and Q. Sun. Distributed consensus algorithm for events detection in cyber-physical systems. *IEEE IoT Journal*, 6(2):2299–2308, 2019.
- [35] Zhiqiang Li, F Richard Yu, and Minyi Huang. A distributed consensus-based cooperative spectrum-sensing scheme in cognitive radios. *IEEE Transactions on Vehicular Technology*, 59(1):383–393, 2009.
- [36] X. Liu, T. Yang, S. Tang, P. Guo, and J. Niu. From relative azimuth to absolute location: Pushing the limit of pir sensor based localization. *MobiCom '20*, 2020.
- [37] Z. Liu, L. Hou, K. Zheng, Q. Zhou, and S. Mao. A dqn-based consensus mechanism for blockchain in iot networks. *IEEE IoT Journal*, pages 1–1, 2021.
- [38] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M.B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *IEEE INFOCOM 2001*, volume 3, pages 1380–1387. IEEE, 2001.
- [39] A. Miller, Y. Xia, K. Croman, E. Shi, and D. Song. The honey badger of bft protocols. *CCS '16*, pages 31–42.
- [40] A.Y. Mjaïd, V. Prasad, M. Jonker, C. Van Der Horst, L. De Groot, and S. Narayan. Ai-based simultaneous audio localization and communication for robots. *IoTDI '23*, page 172–183, 2023.
- [41] H. Moniz, N.F. Neves, and M. Correia. Turquoise: Byzantine consensus in wireless ad hoc networks. *DSN' 10*, pages 537–546. IEEE, 2010.
- [42] R. Pereira, M. Couto, F.Ribeiro, R. Rua, J. Cunha, J.P. Fernandes, and J. Saraiva. Energy efficiency across programming languages: how do energy, time, and memory relate? In *PACM SIGPLAN SLE*, pages 256–267, 2017.
- [43] A.S. Rawat, P. Anand, H. Chen, and P.K. Varshney. Collaborative spectrum sensing in the presence of byzantine attacks in cognitive radio networks. *IEEE Transactions on Signal Processing*, 59(2):774–786, 2010.
- [44] G. Rubambiza, S. Chin, M. Rehman, S. Atapattu, J. F. Martínez, and H. Weatherspoon. Comosum: An extensible, reconfigurable, and Fault-Tolerant IoT platform for digital agriculture. *ATC '23*.
- [45] K. Sakai, M. Sun, W. Ku, T.H. Lai, and A.V. Vasilakos. A framework for the optimal k -coverage deployment patterns of wireless sensors. *IEEE Sensors Journal*, 15(12):7273–7283, 2015.
- [46] D. Vasisht, Z. Kapetanovic, J. Won, X. Jin, R. Chandra, S. Sinha, A. Kapoor, M. Sudarshan, and S. Stratman. {FarmBeats}:an {IoT} platform for {Data-Driven} agriculture. *NSDI '17*, pages 515–529.
- [47] B. Wang. Coverage problems in sensor networks: A survey. *ACM Computing Surveys (CSUR)*, 43(4):1–53, 2011.
- [48] G. Xing, R. Tan, B. Liu, J. Wang, X. Jia, and C. Yi. Data fusion improves the coverage of wireless sensor networks. *MOBICOM' 09*, pages 157–168, 2009.
- [49] H. Zhang and S. Duan. Pace: Fully parallelizable bft from reproposable byzantine agreement. *CCS '22*, page 3151–3164, 2022.