

Lightning: A Hard Real-Time, Fast, and Lightweight Low-End Wireless Sensor Election Protocol for Acoustic Event Localization

Qixin Wang*, *Student Member, IEEE*, Rong Zheng†, *Member, IEEE*, Ajay Tirumala‡, Xue Liu**, *Member, IEEE*, and Lui Sha*, *Fellow, IEEE*

Abstract—We present **Lightning Protocol**, a hard real-time, fast, and lightweight protocol to elect the sensor closest to an impulsive sound source. This protocol can serve proximity-based localization or leader election for sensor collaboration. It utilizes the fact that electromagnetic wave propagates much faster than acoustic wave to efficiently reduce the number of contending sensors in the election. With simple RF bursts, most basic comparison operations, no need of clock synchronization, and a memory footprint as small as 5330 bytes of ROM and 187 bytes of RAM, the protocol incurs $O(1)$ transmissions irrespective of the sensor density and guarantees hard real-time ($O(1)$) localization time cost. Experiment results using UC Berkeley Motes in a common office environment demonstrate that the time delay for Lightning Protocol is in the order of milliseconds. The simplicity of the protocol reduces memory cost, computation complexity, and programming difficulty, making it desirable for low-end wireless sensors.

Index Terms—J.9.f Wireless sensor networks, J.7.g Real time, J.9.a Location-dependent and sensitive, J.9.d Pervasive computing

I. INTRODUCTION

RECENT technology advancements make possible the massive deployment of low-end wireless sensors to form interconnected *Wireless Sensor Networks* (WSNs) for various military and civilian applications. Localization of acoustic events using WSN provides a low-cost solution to tracking moving objects or individual incidents. It can be an integral part of ubiquitous computing environments.

In this paper, we are interested in electing the closest sensor to an acoustic event. Such election serves the purpose of *proximity-based localization*, i.e., giving the location of the elected sensor as the approximate location of the acoustic event. For many applications where wireless sensors are densely deployed, proximity-based localization provides sufficient resolution. Electing the closest sensor to the acoustic event may also serve the purpose of runtime *leader election* for collaborative sensing. For example, many fine-grain localization (triangulation) algorithms [1][2] require collaboration of sensors in the vicinity of the acoustic event, and a unique leader must be elected to start such collaboration.

* These authors are with Department of Computer Science, University of Illinois at Urbana-Champaign. E-mail: {qwang4, lrs}@uiuc.edu

† The author is with Department of Computer Science, University of Houston. E-mail: rzheng@cs.uh.edu

‡ The author is currently with Nvidia Corporation. E-mail: tirumala@uiuc.edu

** The author is with School of Computer Science at McGill University, Canada. E-mail: xueliu@cs.mcgill.ca

A protocol that quickly elects the closest sensor to the acoustic event serves this demand.

An acoustic event localization (simplified as “acoustic localization” in the following) protocol in a dense WSN shall desirably come with following properties.

- **Timeliness:** Many applications demand locating acoustic events in short and bounded time. Particularly, a hard real-time application, such as gunshot localization (see [3]), requires a short and constant ($O(1)$) localization time bound.
- **Lightweightness:** Low-end sensors have constrained computation, storage and communication capabilities. Therefore, it is desirable for acoustic localization algorithms to be simple and efficient.
- **Reasonable Accuracy:** The accuracy of localization must be acceptable, but not excessive. For many applications, proximity-based localization is sufficient.
- **Robustness:** The scheme shall work with unreliable wireless medium and irregular acoustic signals: sound intensity may be direction-dependent and sound propagation may be affected by multi-path effects.
- **Energy-Efficiency:** An energy-efficient localization algorithm should have minimal message exchanges in presence of acoustic events and conserve energy when there is none.

To achieve the above goals, we propose *Lightning Protocol*, a hard real-time, fast, and lightweight wireless sensor election protocol, for locating impulsive acoustic events using low-end wireless sensors. Lightning Protocol exploits the fact that electromagnetic wave propagates much faster than acoustic wave to localize acoustic events in a proximity-based fashion. Like many other acoustic localization protocols, it utilizes acoustic *Time-Of-Arrival* (TOA) information at sensors. However, *Lightning Protocol does not send data packets; nor does it require clock synchronization*. Instead, it allows overlapping wireless broadcasts, which greatly simplifies the implementation. Lightning Protocol exhibits the afore-mentioned desirable properties and fits indoor or open area deployment.

Lightning Protocol is implemented on U.C. Berkeley Motes[4][5] and compared with an ideal *Data Packet* (DP)-based scheme. Using Lightning protocol, the entire machine code only occupies 5330 bytes in ROM and 187 bytes in RAM. In Section VI, experimental results show that the localization time cost of Lightning Protocol is less than 3.6ms. Among 81.4% of the localization trials, there is only *one* wireless broadcast involved (with the rest being suppressed). The accuracy is comparable to or better than that of an ideal collision-free DP scheme.

The rest of the paper is organized as follows. In Section II, we motivate the proposed solution by closely examining the

acoustic characteristics in our application scenarios. The basic Lightning Protocol and its improved versions (energy-efficient, random layout) are described in Section III, IV and V respectively. Experimental results are presented in Section VI, followed by discussions in Section VII. Related works are discussed in Section VIII. Section IX concludes the paper.

II. ASSUMPTIONS AND DESIGN CONSIDERATIONS

We assume following properties on the acoustic signals:

- P1 **Impulsive Sound:** The acoustic signal shall come with clearly detectable onset and thus their TOAs can be measured. We call such acoustic signals *impulsive sounds*, or *beeps*.
- P2 **Bounded Directionality:** The intensity of a beep may be directional, that is, “loudest” does not necessarily mean “closest”. Meanwhile, the directionality is bounded: the closest sensor shall be able to detect the TOA of the sound (a more quantitative definition is given in Section III-C).
- P3 **Moderate Multipath:** An acoustic signal may travel along indirect paths to a sensor; but the closest sensor shall have *Line-Of-Sight* (LOS) to the sound source.

Properties P2 and P3 naturally promote the use of *densely* placed sensors, where LOS is usually available between the sound source and its closest sensor, and hence the acoustic wave would arrive at the closest sensor before echoes. Properties P1 and P2 make TOA a better parameter than acoustic signal intensity for localization: for impulsive sound sources, as long as LOS to the closest sensor is available, TOA of the closest sensor is earlier than TOAs at other sensors, independent of the directionality and loudness of the sound.

Other assumptions are as follows:

- A1 **Uniform Acoustic Medium:** The sound propagation medium (e.g., the air) is stable and ensures uniform acoustic propagation speed, denoted by v .
- A2 **Wireless Range:** Wireless transmission range is at least twice as large as the acoustic sensing range. Therefore, all sensors that overhear the same sound can directly communicate with each other¹.
- A3 **RF Device:** Each sensor is equipped with a half-duplex radio transceiver with configurable *Radio-Frequency* (RF) channel that can be changed during runtime.
- A4 **Deployment Environment:** We assume Lightning Protocol runs in open space or lightly obstructed office environment, without severe acoustic multipath effect. Although significant acoustic multipath effects can be handled with extra time cost (see Section VII-C), they are not the focus of this paper.

III. BASIC LIGHTNING PROTOCOL

A. Intuition

As mentioned in Section II, TOA measurements based on onset detection are more robust to directionality of acoustic signals. A straightforward election algorithm is to compare the TOAs at each sensor. The one with the earliest TOA reading is considered the closest to the sound source. However, this approach suffers two problems, (i) packet communication overhead that grows with increasing sensor density; and (ii) requirement of clock

synchronization as each sensor needs to time stamp the TOA based on its local clock.

Our proposed Lightning Protocol is self-synchronized and only involves $O(1)$ RF broadcasts to elect the closest sensor. The intuition comes from lightning phenomenon in nature. When a lightning bolt strikes, people *see* the lightning much earlier than they *hear* the accompanying thunder, because electromagnetic waves travel much faster than sound waves. To apply this property to our election problem, we notice that an acoustic signal arrives at the closest sensor (denoted as S_1) first (note the dense deployment of sensors empirically guarantees the availability of LOS). If S_1 can *immediately* transmit an RF signal to notify all other sensors, they can decide that they are farther away from the sound source even before the acoustic wave reaches them. Nonetheless, there is still one difficulty: wireless transmissions are usually subject to collisions. In particular, if there are multiple sensors at similar distances to the sound source, broadcasting data packets suffer from collisions that prevent an immediate notification. To solve this problem, we propose using raw RF burst to signal the arrival of sound wave. RF bursts can overlap, which obviates the random backoff based wireless MAC, and makes the protocol immune to wireless broadcast collisions. RF bursts can be detected by measuring the received radio energy [6].

Based on the above ideas, we propose Basic Lightning Protocol in the next subsection.

B. Protocol Details

For the time being, we assume sensors are placed at regular grids on the plane (Fig. 1). Without loss of generality, we assume square grid as shown in Fig. 1. Extension of Lightning Protocol to *random layouts* is presented in Section V. In a regular grid layout, each sensor is assigned a color i , as shown in Fig. 1 (for square, $i \in \{1, 2, 3, 4\}$). The color assignment guarantees that for any point in the plane, the enclosing sensors are of distinct colors. As is to be explained later, such a color assignment guarantees the uniqueness of the elected sensor.




Fig. 1. A Square Sensor Grid Layout

Let T_b be the minimum duration that an RF burst must sustain in order to be robustly detected. The Basic Lightning Protocol goes as follows:

- 1) All sensors are initially in *RF listening mode*.
- 2) While in the listening mode, if a beep is *recognized*², a sensor with color i switches immediately to *RF bursting*

²“Recognize” refers to the time instance when the arriving beep is detected by the sensor’s signal recognition module and reported to the sensor’s election module.

¹This can be achieved by properly adjust the threshold for beep detection.

mode and broadcasts *without back-off* an RF burst of $i \cdot T_{burst}$ duration (T_{burst} is an implementation-specific constant).

Immediately after the burst, the sensor switches back to RF listening mode and samples the wireless medium for T_b . If no other RF burst is recognized, it wins the election and enters the *elected mode*. Otherwise, it loses the election and enters the *suppressed mode*. In both cases, a sensor sets up a timer of length T_{reset}^{basic} .

- 3) At any time during RF listening mode, if an RF burst is recognized, the sensor fails the election and enters the suppressed mode. Meanwhile, the sensor sets up a timer of length T_{reset}^{basic} .
- 4) After the election is completed (when timer T_{reset}^{basic} expires), all sensors return to the listening mode.

T_{reset}^{basic} is a preset constant for all sensors to reset their state meanwhile ensure consistency. The setting of T_{reset}^{basic} is explained later.

As mentioned earlier in Section II, we assume sensors can switch between multiple RF channels during runtime. This allows us to use a separate RF channel for burst. When a sensor is not elected, it only listens and bursts at the RF burst channel. Once elected, a sensor may switch to data communication channel to conduct data exchanges. Fig. 2 summarizes the state transitions of the Basic Lightning Protocol.




Fig. 2. State Transition Diagram of Basic Lightning Protocol for a Sensor with Color i

C. Properties of Basic Lightning Protocol

In this section, we prove that Basic Lightning Protocol elects a unique sensor within $O(1)$ time delay, and the elected sensor is among the closest sensors enclosing the sound source.

For now, we assume there is only one acoustic event (beep). The case of multiple sound events is discussed in Section VII-A. *Without loss of generality, we consider a square sensor grid, and the sound source location $p = (x, y) \in \square ABCD$ (see Fig. 1, $\square ABCD$ refers to the shaded square area).* The distances between the sound source and the four adjacent sensors S_1, S_2, S_3 and S_4 (colored 1,2,3 and 4 respectively) are $d_1(p), d_2(p), d_3(p)$ and $d_4(p)$ respectively. $d_{other}(p)$ is the distance between the sound source location p and the closest sensor other than $S_1 \sim S_4$. We assume the sensor density is sufficiently high so that at any position $p \in \square ABCD$, a beep is recognizable to all sensors $S_1 \sim S_4$ regardless of its directionality. As RF waves travel a lot faster than acoustic signals, we ignore the propagation delay of RF burst. Let t_{recg} be units of time for a sensor to recognize a beep, $t_{recg} \in [0, \Delta_{recg}]$, i.e. Δ_{recg} is the maximum time cost to recognize a beep. The notations used in the analysis are summarized in Table I.

TABLE I
NOTATION USED IN THE ANALYSIS

t_{recg}	a random variable denoting the time cost to recognize a beep
$0, \Delta_{recg}$	the minimum and maximum possible time to recognize a beep (Δ_{recg} can be 0.12ms)
T_{burst}	the RF burst duration for a sensor colored 1 (can be 0.8ms)
iT_{burst}	the RF burst duration for a sensor colored i
T_b	minimum duration that an RF burst must sustain in order to be recognized (can be 0.4ms)
T_{elect}^*	election delay, $*$ corresponds to <i>light</i> for Lightning Protocol, or <i>data</i> for DP protocol.
T_{reset}^*	time to reset to initial mode from suppressed or elected mode. $*$ corresponds to a specific version of Lightning Protocol
T_{bound}^*	election delay bound, the time bound from the beep takes place to all sensors enter elected or suppressed mode. $*$ corresponds to a specific version of Lightning Protocol
$T_{bound}^+ + T_{reset}^*$	turn-around time
R_{beep}^{max}	the maximum audible radius of a beep
p	(x, y) location of the sound source, $p \in \square ABCD$, see Fig. 1
l	grid edge length, see Fig. 1
$d_1(p) \sim d_4(p)$	distance between the sound source ($p \in \square ABCD$) to the four adjacent sensors ($S_1 \sim S_4$) in Fig. 1.
$d_{other}(p)$	distance between the sound source and the closest sensor other than the four adjacent sensors
$d(p)$	distance between the sound source and the elected sensor
v	sound speed (1ft/ms)

Before delving into the derivation, we first present the key results. Theorem 1 states that Basic Lightning Protocol elects a single sensor and the localization error is bounded; Corollary 1 gives the upper bound of *election delay* (see Definition 1).

Theorem 1 (Winner Uniqueness and Error Bound):

If $T_{burst} \geq 2T_b$ and square grid edge length $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$, Basic Lightning Protocol elects *exactly one* sensor, which is among $S_1 \sim S_4$. Furthermore, the distance $d(p)$ between the sound source location p and the elected sensor satisfies $d(p) \leq d_1(p) + (\Delta_{recg} + T_b)v$.

Definition 1 (Election Delay): Election Delay refers to the time duration since the beep takes place till every sensor enters either elected or suppressed mode.

Corollary 1 (Election Delay Bound): If $T_{burst} \geq 2T_b$ and $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$, Basic Lightning Protocol incurs an election delay no greater than $T_{bound}^{basic} = \frac{\sqrt{2}l}{2v} + 2\Delta_{recg} + 4T_{burst} + 2T_b$.

To prove the above results, we first present the following lemmas.

Lemma 1: If square edge length $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$ and $T_{burst} > T_b$, then only S_1, S_2, S_3 and S_4 may burst, all other sensors are suppressed.

Proof: Recall that $d_{other}(p)$ is the distance between p and the closest sensor other than $S_1 \sim S_4$.

$\inf_{p \in \square ABCD} \{d_{other}(p)\} = l$, and this value is reached *only* when $p = A$. The sensor is either at point E or F (see Fig. 1). On the other hand, $\sup_{p \in \square ABCD} \{d_1(p)\} = \frac{\sqrt{2}l}{2}$, and this value is reached *only* when $p = C$. Hence, $\forall p \in \square ABCD$, $d_{other}(p) - d_1(p) > \frac{2-\sqrt{2}}{2}l$ (note “=” can never be achieved

because that requires p be at A and C at the same time). When $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$, we have $\forall p \in \square ABCD$, $\frac{d_{other}(p)-d_1(p)}{v} > (\Delta_{recg} + T_b)$. Hence, when the sound wave reaches any of the sensors other than $S_1 \sim S_4$, S_1 either has recognized the beep and burst for at least T_b units of time, or has been suppressed by a burst from S_2, S_3 or S_4 . In both cases, any sensor other than $S_1 \sim S_4$ is suppressed. ■

Definition 2 (Non-Deterministic Area): When the beep is at a location that is equidistant or nearly equidistant to multiple sensors, due to randomness of beep recognition time cost t_{recg} , the closest sensor may not always burst first and suppress other sensors, i.e. the elected sensor may be not the closest sensor. Such locations form a *Non-Deterministic Area*. However, we shall see the localization error incurred by non-deterministic area is *small* (Theorem 1). Specifically, we define non-deterministic area V as:

$$V \stackrel{def}{=} \left\{ p \mid p \in \square ABCD \text{ and } \frac{d_1(p)}{v} + \Delta_{recg} + T_b \geq \min \left\{ \frac{d_2(p)}{v}, \frac{d_3(p)}{v}, \frac{d_4(p)}{v} \right\} \right\}$$

Note as mentioned in the beginning of this subsection: *without loss of generality*, we only consider the cases that $p \in \square ABCD$. The above definition of V therefore only refers to the non-deterministic area *within* $\square ABCD$. For areas other than $\square ABCD$, the definition of non-deterministic area follows the same pattern, with $d_1(p)$ replaced with distance to the closest sensor, and $d_2(p), d_3(p), d_4(p)$ replaced with distances to the other three enclosing sensors.

Lemma 2: When $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$ and $T_{burst} > T_b$, if $p \in \square ABCD$ and $p \notin V$, then only S_1 would burst and thus win the election.

Proof: By Lemma 1, the only sensors that may compete with S_1 are S_2, S_3 and S_4 . $p \in \square ABCD$ and $p \notin V$ means $\frac{d_1(p)}{v} + \Delta_{recg} + T_b < \min \left\{ \frac{d_2(p)}{v}, \frac{d_3(p)}{v}, \frac{d_4(p)}{v} \right\}$. Therefore, by time $d_1(p)/v + \Delta_{recg}$, S_1 have recognized the beep and started bursting, meanwhile the beep has not yet reached S_2, S_3 and S_4 . By time $\frac{d_1(p)}{v} + \Delta_{recg} + T_b$, S_1 have already burst for T_b duration, while the beep has not yet reached any of S_2, S_3 and S_4 . Therefore, $S_2 \sim S_4$ are suppressed by S_1 (note since $S_2 \sim S_4$ are in listening mode, a burst duration of T_b from S_1 is long enough to suppress $S_2 \sim S_4$). ■

Likewise, we can prove that if $p \in V$, S_2, S_3 or S_4 may burst before the closest sensor S_1 , or even suppress S_1 . In Fig. 3, we plot non-deterministic area V for square $\square ABCD$ (edge-length equals 2 ft) under different $\Delta_{recg} + T_b$. To prove Theorem 1, note that by the color assignment shown in Fig. 1, $S_1 \sim S_4$ each has a distinct color. Without loss of generality, suppose S_1, S_2, S_3 and S_4 are of color 1, 2, 3 and 4 respectively. Then we have the following lemma:

Lemma 3: If $T_{burst} \geq 2T_b$, then after a beep, if multiple sensors burst, the one with “largest color number” always wins the election.

Proof: According to Lemma 1, if multiple sensors burst after a beep, the total number of bursting sensors can only be 2, 3 or 4.

Case 1: 2 sensors burst. We denote them as S_i and S_j , where $i, j \in \{1, 2, 3, 4\}$ and $i > j$. Suppose S_i bursts during time interval $[t_i^0, t_i^0 + iT_{burst}]$ and S_j bursts during $[t_j^0, t_j^0 + jT_{burst}]$. Then we




Fig. 3. Area lower-right to the contour is the non-deterministic area within $\square ABCD$ of Fig. 1, where $\Delta_{recg} + T_b = 0.2, 0.4, 0.6, 0.8$ and 1ms respectively. Assume sound propagation speed $v = 1\text{ft/ms}$ and edge-length of $\square ABCD$ to be 2ft .

must have $|t_i^0 - t_j^0| < T_b$, otherwise, the later burst would be suppressed. (i) If $t_i^0 < t_j^0$, after S_j stops bursting, burst from S_i will last for

$$\begin{aligned} & iT_{burst} - jT_{burst} - (t_j^0 - t_i^0) \\ &= (i - j)T_{burst} - (t_j^0 - t_i^0) \\ &> T_{burst} - T_b \geq 2T_b - T_b = T_b \end{aligned}$$

The last “ \geq ” is due to $T_{burst} \geq 2T_b$. Hence, S_j has enough time to recognize S_i ’s burst after its own burst, and to realize that it has lost the election. (ii) If $t_i^0 \geq t_j^0$, after S_j stops bursting, burst from S_i will last for $iT_{burst} - (jT_{burst} - (t_i^0 - t_j^0)) = (i - j)T_{burst} + (t_i^0 - t_j^0) \geq T_{burst} \geq 2T_b > T_b$. Hence, S_j has enough time to recognize S_i ’s burst after its own burst, and realize that it has lost the election. Therefore, according to (i)(ii), S_i would always win.

Case 2: 3 and 4 sensors burst. The same reasoning as in Case 1 can be applied here. ■

Now we are in the position to prove Theorem 1. Theorem 1 states that a unique sensor is elected in Basic Lightning Protocol and the localization error is bounded.

Proof: By Lemma 1, the only possible competing sensors are $S_1 \sim S_4$. By Lemma 2, Theorem 1 sustains when $p \notin V$.

When $p \in V$, by Lemma 3, only one sensor wins. If the winning sensor is S_1 , $d(p) = d_1(p) \leq d_1(p) + (\Delta_{recg} + T_b)v$. If it is S_2 , there must be $d(p) = d_2(p) \leq d_1(p) + (\Delta_{recg} + T_b)v$, otherwise S_1 would have burst for T_b before the sound wave ever reaches S_2 and therefore suppressed S_2 . The same reasoning applies when the winning sensor is S_3, S_4 . ■

From Theorem 1, we can prove that the election delay is bounded (Corollary 1) as follows:

Proof: The winning sensor starts the burst no later than $\frac{d(p)}{v} + \Delta_{recg}$. For the square sensor grid layout, the longest bursting time is $4T_{burst}$. Hence, by time $\frac{d(p)}{v} + \Delta_{recg} + 4T_{burst} + T_b$, the winning sensor has entered the elected mode. Since the winning sensor has already stopped bursting by that time, all other sensors

must have entered the *suppressed mode*. Therefore, the election process completes no later than

$$\begin{aligned} & \frac{d(p)}{v} + \Delta_{recg} + 4T_{burst} + T_b \\ \leq & \frac{d_1(p)}{v} + (\Delta_{recg} + T_b) + \Delta_{recg} + 4T_{burst} + T_b \\ & \text{(Due to Theorem 1)} \\ \leq & \frac{\sqrt{2}l}{2v} + 2\Delta_{recg} + 4T_{burst} + 2T_b \end{aligned}$$

That is, the election delay is upper bounded by $T_{bound}^{basic} = \frac{\sqrt{2}l}{2v} + 2\Delta_{recg} + 4T_{burst} + 2T_b$. ■

Now, let us consider how to choose an appropriate T_{reset}^{basic} . If a beep takes place at time 0, and suppose the maximal audible range of the beep is R_{beep}^{max} , by time $\frac{R_{beep}^{max}}{v} + \Delta_{recg}$, it has been recognized by the farthest away sensor. On the other hand, according to Corollary 1, every hearing node (note according to assumption A2 in Section II, wireless broadcast range is at least $2R_{beep}^{max}$) either enters suppressed mode or elected mode by time T_{bound}^{basic} . Assuming a beep is recognized only when an onset is detected preceded by a period of silence, we can reset every sensor to initial listening mode at time $\max\{\frac{R_{beep}^{max}}{v} + \Delta_{recg}, T_{bound}^{basic}\}$. Therefore, as a conservative approach, it is safe to let $T_{reset}^{basic} = \max\{\frac{R_{beep}^{max}}{v} + \Delta_{recg}, T_{bound}^{basic}\}$. Hence, by time $T_{bound}^{basic} + T_{reset}^{basic}$ (called *turn-around time*) all sensors are reset to RF listening mode.

IV. ENERGY-EFFICIENT LIGHTNING PROTOCOL

Under Basic Lightning Protocol, a sensor keeps its radio active all the time either in RF listening mode or RF bursting mode. This can be costly in energy consumption. To save energy, a naive approach would be that a sensor turns on its RF module only when an acoustic beep is recognized. It turns off the radio after election is completed (with results sent back to the sink if necessary). However, consider the case when the acoustic beep arrives at the closest sensor S^* earlier than sensor S , such that S^* recognizes the beep, say, 11ms earlier than S . If the burst of S^* lasts for only 10ms, then S would not have turned on its RF module before S^* finishes its burst. In this case, the burst from S^* cannot suppress S . Consequently, S considers itself elected as well.

To handle the above problem and conserve energy, we propose an Energy-Efficient Lightning Protocol.

A. Protocol Details

Energy-Efficient Lightning Protocol runs the following idea: upon recognition of a beep, a sensor turns on RF and listens for a period of time before it starts to burst. Let the maximal audible radius of a beep be R_{beep}^{max} . If a sensor recognizes a beep, after $\Delta_{defer} \stackrel{def}{=} \frac{R_{beep}^{max}}{v} + \Delta_{recg}$, the beep must have reached and been recognized by the farthest sensor within the R_{beep}^{max} radius. That is, after Δ_{defer} , any sensor that can hear the beep has turned on its RF module and switched to RF listening mode. And all the rest can be the same as Basic Lightning Protocol.

The formal description of Energy-Efficient Lightning Protocol is as follows:

- 1) All sensors are initially in RF *sleeping mode*.

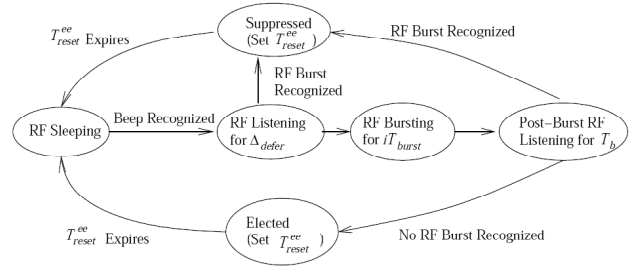


Fig. 4. State Transition Diagram of Energy-Efficient Lightning Protocol

- 2) While in RF sleeping mode, if a beep is recognized, a sensor turns on its RF module to RF *listening mode* immediately for a duration of $\Delta_{defer} = \frac{R_{beep}^{max}}{v} + \Delta_{recg}$.
- 3) If no RF burst is recognized in the Δ_{defer} period, a sensor with color i enters RF *bursting mode* and transmits *without back-off* a burst of $i \cdot T_{burst}$ duration. Immediately after the burst, the sensor switches back to RF listening mode and samples the wireless medium for T_b . If no other RF burst is recognized, the sensor decides that it wins the election and enters the *elected mode*. Otherwise, it loses the election and enters the *suppressed mode*. In both cases, a sensor sets up a timer of length T_{reset}^{ee} .
- 4) At any time during RF listening mode, if an RF burst is recognized, the sensor fails the election and enters the suppressed mode. Furthermore, the sensor sets up a timer of length T_{reset}^{ee} .
- 5) After the election is completed (i.e., after timer T_{reset}^{ee} expires), sensors return to the RF sleeping mode.

T_{reset}^{ee} is a preset constant for all sensors to reset to initial state. The setting of T_{reset}^{ee} is explained later.

Fig. 4 illustrates the state transition of Energy-Efficient Lightning Protocol. Compared to Fig. 2, an extra RF sleeping mode is introduced. In RF sleeping mode, a sensor can put its radio module to low power states to conserve energy.

B. Analysis of Energy-Efficient Lightning Protocol

1) First, we prove that Energy-Efficient Lightning Protocol preserves the desirable properties of the Basic Lightning Protocol, i.e., single winner, bounded error, $O(1)$ transmissions, and bounded election delay.

Theorem 2 (Winner Uniqueness and Error Bound):

Lemma 1~3 and Theorem 1 remain valid for Energy-Efficient Lightning Protocol.

Proof: According to Energy-Efficient Lightning Protocol, every sensor delays Δ_{defer} after it recognizes the beep. By the time the first RF burst starts, all sensors that can hear the beep have switched to RF listening mode. From thereon, the protocol proceeds as the Basic Lightning Protocol. In other words, it is equivalent to that the beep takes place Δ_{defer} units of time later than the actual beep with every sensor running Basic Lightning Protocol. Therefore, all proofs in Lemma 1~3 and Theorem 1 sustain. ■

Similarly, the following corollary holds.

Corollary 2 (Election Delay Bound): If $T_{burst} \geq 2T_b$ and $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$, Energy-Efficient Lightning Protocol incurs

an election delay no greater than $T_{bound}^{ee} = T_{bound}^{basic} + \Delta_{defer}^{ee} = \frac{\sqrt{2l}}{2v} + \Delta_{defer}^{ee} + 2\Delta_{recg} + 4T_{burst} + 2T_b$, where $\Delta_{defer}^{ee} = \frac{R_{beep}^{max}}{v} + \Delta_{recg}$.

Similar to the argument for T_{reset}^{basic} in Section III-C, assume the maximal audible range of the beep is R_{beep}^{max} , a conservative value for the time to reset sensors to initial sleep mode from suppressed or elected mode T_{reset}^{ee} is $T_{bound}^{ee} - \Delta_{defer}^{ee}$. Therefore, by time $T_{reset}^{ee} + T_{bound}^{ee}$ (called *turn-around time*), all sensors are reset to the initial RF sleeping mode.

2) Next, we evaluate the energy consumption. In the Energy-Efficient Lightning Protocol, sensors are active (with radio on) only during the election period. Since the election time is short, the energy consumption is small. For example, two AA batteries can sustain a U.C. Berkeley Mote with its radio active for 59 hours. In our implementation (Section VI), a burst lasts 3.2ms at the most. The RF listening time $\Delta_{defer}^{ee} + T_b = 20.5\text{ms}$ (suppose $R_{beep}^{max} = 20\text{ft}$). Approximately, with two AA batteries, over 9 million acoustic events can be localized using Energy-Efficient Lightning Protocol.

In what follows, unless otherwise specified, the term “Lightning Protocol” refers to Energy-Efficient Lightning Protocol.

C. A Quantitative Comparison with Data Packet (DP)-based Scheme

Let problem scale n be the number of sensors hearing the acoustic beep. Let T_{elec}^{light} and T_{elec}^{data} be the election delay of Lightning Protocol and a DP scheme respectively. We have $T_{elec}^{light} \leq T_{bound}^{ee}$ and $T_{bound}^{ee} \sim O(1)$ (Corollary 2), therefore $T_{elec}^{light} \sim O(1)$, which is the best possible for any protocol. $T_{bound}^{ee} \sim O(1)$ for the following reasons: by the definition of T_{bound}^{ee} given in Corollary 2, the only parameter that is relevant to total number of sensors n is l , i.e. the sensor grid edge length. However, as n increases, l becomes smaller, which makes T_{bound}^{ee} smaller. One may argue that in Theorem 1,2 and Corollary 1,2, it is required that $l > \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$. This requirement is needed to guarantee only *one* sensor is elected. In cases that $l \leq \frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$, there may be multiple winners within $\frac{2}{2-\sqrt{2}}(\Delta_{recg} + T_b)v$ radius (due to Theorem 3) to the sound source (there is no winner outside this circle). But in this case, with similar approach, it can be proved that both election delay and turn-around time ($T_{bound}^{ee} + T_{reset}^{ee}$) are still $O(1)$. The same conclusion also applies to Basic Lightning Protocol.

For DP schemes, sensors that hear the beep must contend for the wireless medium to broadcast their TOA readings. If an IEEE 802.11 like MAC protocol is adopted to resolve the channel contention, the expected number of data packet collisions before the *first* successful data packet broadcast grows exponentially with the number of contending sensors [7], i.e., $\Omega(c^n)$, where c is a constant > 1 . For wireless sensors equipped with half-duplex radio, once a transmission starts, it lasts till the entire packet is transmitted, and collisions are detected by lack of acknowledgments at the sender. Let T_{pack} be the time to transmit a data packet, the expected delay before the first successful data packet broadcast is, therefore, $\Omega(c^n \times T_{pack}) = \Omega(c^n)$. The expected election delay $\mathbb{E}[T_{elec}^{data}]$ to get the data packet with earliest TOA is strictly no less than $\Omega(c^n \times T_{pack})$. Therefore, $\mathbb{E}[T_{elec}^{data}] \sim \Omega(c^n)$.

One may argue that MAC contention can be alleviated by reducing the radio range in data communication. In this case,

assumption A2 (see Section II) that radio range is at least twice as large as the acoustic sensing range may no longer hold. Sensors overhearing the same acoustic event cannot reach each other directly. Therefore, multi-hop communications are needed to elect the closest sensors. As a result, the election delay is thus the sum of single-hop election delay and the delay incurred by multi-hop forwarding of data packets. When the sensors get denser (i.e. when n increases), the election delay increases monotonically anyway.

V. RANDOM PLACEMENT OF SENSORS

So far, regular placement of sensors (e.g. square grid) is assumed, where nodes are colored to ensure the *uniqueness* of elected sensor. If electing multiple (closest) sensors is allowed, placement of sensors can be random, and coloring of sensors is no longer necessary. The reason is that the closest sensors burst first and suppress sensors farther away. The only modification to the Basic and Energy-Efficient Lightning Protocol is that the RF burst time of each sensor is T_{burst} instead of iT_{burst} (i is the color of the sensor). Also, as before, we require $T_{burst} \geq 2T_b$.

Let $d(S)$ be the distance from sound source to sensor S . Let S^* be the closest sensor to the sound source. Then we have:

Theorem 3 (Error Bound): For both Basic Lightning Protocol and Energy-Efficient Lightning Protocol, with random sensor placement, only sensors within $d(S^*) + (\Delta_{recg} + T_b)v$ radius of the sound source may be elected. We denote this area as $\odot Beep$.

Proof: First, we consider the Basic Lightning Protocol. Suppose a beep takes place at time 0. By time $\frac{d(S^*)}{v}$, the sound wave reaches the closest sensor S^* and by time $\frac{d(S^*)}{v} + \Delta_{recg}$, S^* should have recognized the beep and started bursting. By time $\tau = \frac{d(S^*)}{v} + \Delta_{recg} + T_b$, S^* should have burst for at least duration T_b . On the other hand, for any sensor S whose $d(S) > \tau v$, by time τ , it would have not yet heard the beep. Therefore S is suppressed by S^* 's burst, i.e. any sensor whose $d(S) > \tau v = d(S^*) + (\Delta_{recg} + T_b)v$ would not win the election.

The above argument also applies to the Energy-Efficient Lightning Protocol. ■

Corollary 3 (Election Delay Bound): If a beep takes place at time 0, we have (i) the election delay of Basic Lightning Protocol with random sensor placement is bounded by $T_{bound}^{basic} = \frac{d(S^*)}{v} + 2\Delta_{recg} + T_{burst} + 2T_b$; (ii) Energy-efficient Lightning Protocol with random sensor placement elects the closest sensors in $T_{bound}^{ee} = \frac{d(S^*)}{v} + \Delta_{defer}^{ee} + 2\Delta_{recg} + T_{burst} + 2T_b$, where $\Delta_{defer}^{ee} = \frac{R_{beep}^{max}}{v} + \Delta_{recg}$.

Proof: First, we consider the Basic Lightning Protocol. By Theorem 3, only sensors within $\odot Beep$ would burst, all sensors outside of $\odot Beep$ would be suppressed by the bursts from sensors within $\odot Beep$. That is, when the sensors within $\odot Beep$ complete their Basic Lightning Protocol election, the whole election is completed.

From Theorem 3, the radius of $\odot Beep$ is $d(S^*) + (\Delta_{recg} + T_b)v$, therefore $\forall S \in \odot Beep$, the latest time that the beep would reach S is $\tau_0 = \frac{d(S^*) + (\Delta_{recg} + T_b)v}{v} = \frac{d(S^*)}{v} + \Delta_{recg} + T_b$. The latest time S would recognize the beep is $\tau_1 = \tau_0 + \Delta_{recg}$, and then it will burst for T_{burst} and perform a post-burst RF sample of T_b to see whether other RF bursts exist. Therefore, by time $\tau_2 = \tau_1 + T_{burst} + T_b = \frac{d(S^*)}{v} + 2\Delta_{recg} + T_{burst} + 2T_b$, S would have completed its election procedure.

Similar argument also applies to the Energy-Efficient Lightning Protocol. ■

Although according to Theorem 3, multiple winners may be elected, experiment results show that the number of multiple winners is still small (see Section VI-C.3). This is because of the randomness of TOA and TOA recognition time cost (t_{recg}), which equals randomly coloring sensors in some sense. As future work, we plan to further investigate randomized mechanisms to reduce the number of elected sensor within \odot Beep.

VI. EXPERIMENTS AND COMPARISONS

A. Implementation and Lab Environment

We implemented Lightning Protocol on U.C. Berkeley MICA Motes, using their standard sensing boards, and an acoustic sampling rate of 8kHz. The transmission/reception of RF bursts is supported by the RF hardware of Motes [6]. The final total footprint (including TinyOS, Mote's operating system) is 5330 bytes in ROM and 187 bytes in RAM. According to our implementation, $\Delta_{recg} = 0.12\text{ms}$, $T_b = 0.4\text{ms}$ and $T_{burst} = 0.8\text{ms}$.

Several experiments are conducted to evaluate Lightning Protocol and compare it against an ideal DP scheme. Video clips of the demo are available at [8].

The experiments are all conducted in a common office environment filled with daily RF interferences, such as RF interferences from IEEE 802.11a/b/g WLANs (deployed throughout the building), large-scale computer clusters and alternate current power cables (deployed right beneath the floor). The wireless medium also suffers worse large-scale path loss and multipath fading (echoes) effects than open spaces [9].

B. Experiments with Regular Sensor Placement

The experimental setup is as follows:

16 Motes are placed on square grid points, each monitoring a square area of $4\text{ft} \times 4\text{ft}$. Fig. 5 provides a top-view of the layout of sensors, sound source orientation and locations. To evaluate the robustness of the protocol, we use upright speaker oriented along the X-axis (see Fig. 5 and Fig. 6) as directional sound source. Fig. 6(b) demonstrates the irregular intensity field of the speaker playing a hand clap sound.

As a baseline, we also implement a *Data Packet (DP) Protocol*, as shown in Fig. 7. Both Lightning Protocol and DP protocol use the same TOA recognition module.

1) *Localization Accuracy*: Since the purpose of proximity-based localization is to find the closest sensor, we define the following metric to measure localization error e :

$$e \stackrel{\text{def}}{=} \text{dis}(S) - \text{dis}(S^*), \quad (1)$$

where S is the location of elected sensor, S^* is the location of the sensor closest to sound source. $\text{dis}(x)$ measures the Euclidean distance between sensor x and the sound source.

For each location of the sound source in Fig. 5, 10 trials of experiments are carried out using Lightning Protocol and DP protocol respectively. Furthermore, the DP protocol results are ideal in the sense that we only count those trials where no data packet loss occurs. Fig. 8 shows localization error statistics for both schemes. From Fig. 8, we see that Lightning Protocol achieves comparable or even better accuracy than ideal DP

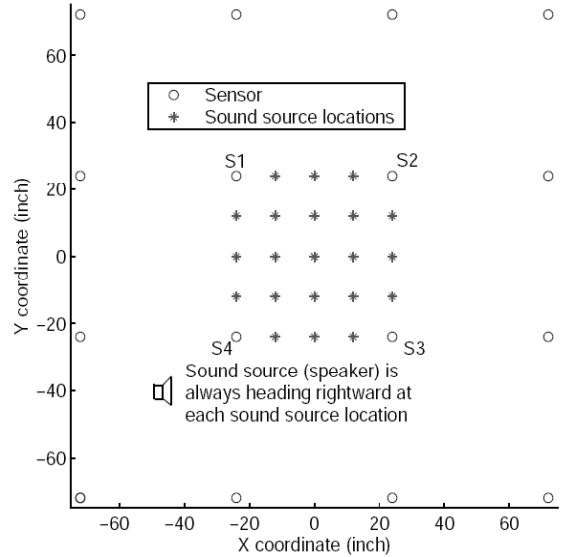


Fig. 5. Top-view of sensor layout, sound source (speaker) orientation and locations

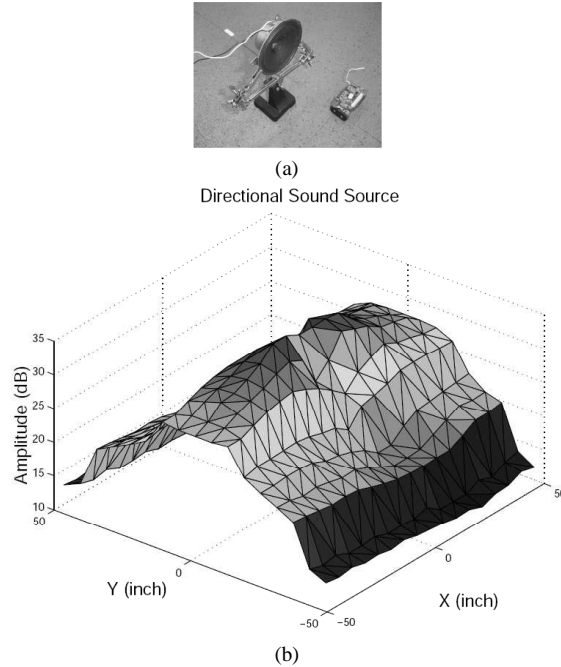


Fig. 6. Sound Source. (a) Speaker (directional sound source) and U.C. Berkeley MICA Mote; (b) Directional sound field of the speaker.

protocol. This is because DP protocol requires accurate clock synchronization to determine which sensor has the earliest TOA³ while Lightning Protocol needs no clock synchronization.⁴

2) *Election Delay*: The election delay T_{elec}^{data} of DP protocol are determined by two factors, i.e., (i) the medium access control protocol used and (ii) the order of transmissions for sensors. In TinyOS 1.0 [5], a simple CSMA/CA mechanism with fixed contention window size cw is implemented. Upon detecting an

³In our implementation, clock synchronization is done by broadcasting a sync-packet from a dedicated synchronization node before each beep. Every sensor, on receiving the sync-packet resets its local clock to 0.

⁴Some errors in our experiment may still look large. This is further explained in Section VII-B.

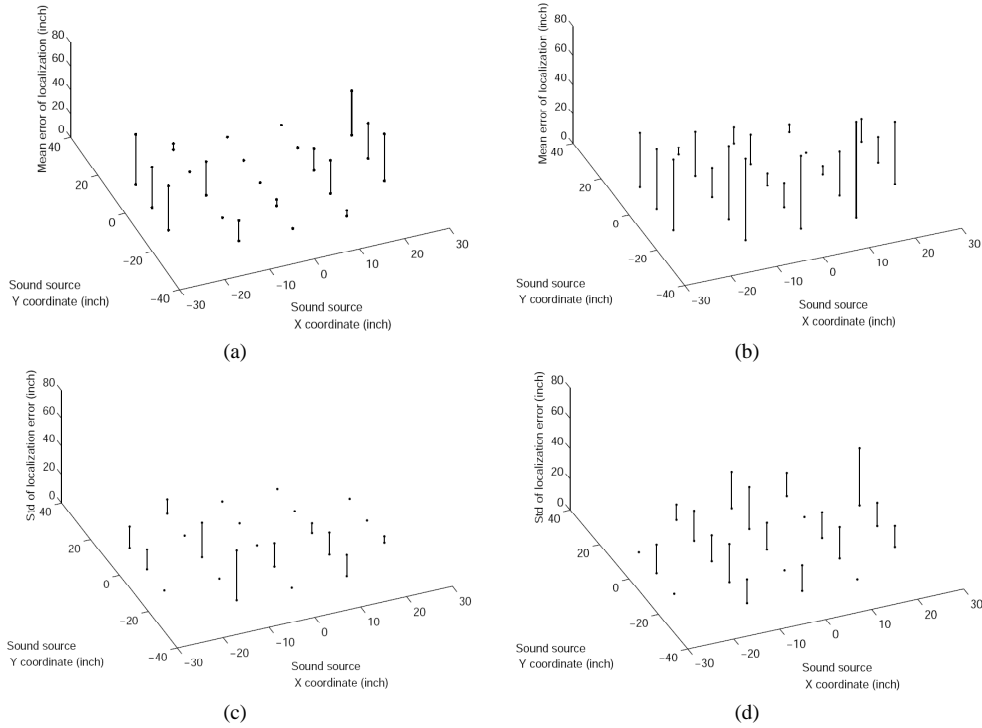


Fig. 8. Localization Error Comparison with Regular Sensor Layout. (a) Mean of Lightning Protocol Localization Error; (b) Mean of Data Packet Protocol Localization Error; (c) Standard Deviation of Lightning Protocol Localization Error; (d) Standard Deviation of Data Packet Protocol Localization Error.

1. /* When a beep takes place */
2. All sensors report their detected TOAs;
3. Pick the sensor with the earliest TOA as the closest sensor;

Fig. 7. Data Packet Protocol

idle channel, sensors that have backlogged packets first back-off randomly (uniformly distributed in $[0, cw]$) and then transmit if the channel remains idle. The current TinyOS implementation does not retransmit data packets in presence of collisions. We measure two delays in the DP protocol. The first one is T_{win1st}^{data} , defined as the time for the winner sensor (i.e. the sensor with the earliest TOA reading. Note the winner sensor is not necessarily the actual closest sensor.) to send out its first data packet; the second metric is T_{any1st}^{data} defined as the time it takes to send out the first data packet from any sensor. Clearly, $T_{win1st}^{data} \geq T_{any1st}^{data}$. And T_{any1st}^{data} is a lower bound of election delay for data packet based schemes.

We measure the election delay of Lightning Protocol and compare it against the delays measured in DP protocol. This comparison errs on the pessimistic side for Lightning Protocol, as DP protocol does not retransmit collided packets.

TABLE II
STATISTICS OF ELECTION DELAY

Metric(ms)	Max	Min	Mean	Standard Deviation
T_{elec}^{light} , winner colored 4	3.6	3.6	3.6	0
T_{win1st}^{data}	101.8	15.4	33.6	14.9
T_{any1st}^{data}	40.8	14.9	20.2	5

Fig. 9(a)(b)(c) show the measurement results of T_{elec}^{light} , T_{win1st}^{data}

and T_{any1st}^{data} respectively.⁵ From Fig. 9(a), we see T_{elec}^{light} is fixed and is only determined by the color of the elected sensor; In Fig. 9(b) and (c), T_{win1st}^{data} and T_{any1st}^{data} are scattered because of the random back-off mechanisms in TinyOS radio stack .

Table II compares statistics of T_{elec}^{light} , T_{win1st}^{data} and T_{any1st}^{data} . Clearly, even the maximum T_{elec}^{light} (when elected sensor is colored 4) is shorter than the minimum T_{win1st}^{data} and T_{any1st}^{data} . This indicates that Lightning Protocol always incurs less election delay than any data packet-based schemes.

3) *Number of Transmissions*: Next, we count the number of transmissions in the Lightning Protocol. From the analysis in Section III-C, we know that in Lightning Protocol, at most 4 sensors would burst and each bursts only once, independent of the total number of sensors hearing the beep. This is proved by the experimental results: 81.4% of the localization trials only involves *one* burst (broadcast); 18.6% of the trials incur two bursts; and none involves more than two bursts.

C. Experiments with Random Sensor Placement

We use the same sound source used in Section VI-B, but deploy sensors in random layouts, as shown in Fig. 10. The random layout in Fig. 10(a) is *uniform*, where 16 sensors are uniformly distributed in a square area of 12ft \times 12ft. On the other hand, the random layout in Fig. 10(b) is *clustered*, where 20 sensors are randomly distributed but deliberately made denser around the

⁵The comparison is made between Basic Lightning Protocol and DP Protocol. Both protocols share the same acoustic propagation delay from sound source to the closest sensor, and share the same TOA recognition module, therefore time zero refers to the time that TOA is recognized at the closest sensor. For Energy-Efficient Lightning Protocol, there should be an additional Δ_{defer} delay. However, if the corresponding DP Protocol is also energy efficient, it should also have the additional Δ_{defer} delay.




Fig. 9. Election Delay Comparison. (a) Lightning Protocol: Election delay distribution; (b) DP Protocol: Distribution of T_{win1st}^{data} (time when the 1st packet is sent out from the winner sensor); (c) DP Protocol: Distribution of T_{any1st}^{data} (time when the 1st packet is sent out from any sensor).

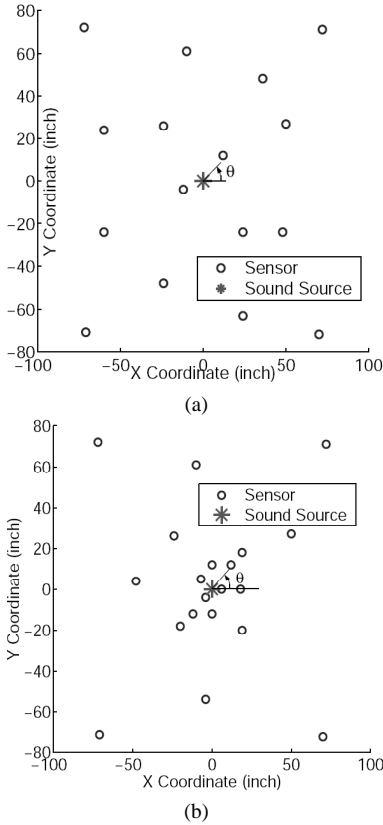


Fig. 10. Random Sensor Layout Testbed Settings. (a) Uniform Random Layout; (b) Clustered Random Layout with Additional Sensors Around the Sound Source.

sound source, so as to intensify election contention within \odot Beep (see Theorem 3 and Corollary 3). For random layout, it is not supposed to elect only *one* winner, therefore sensor coloring is no longer necessary. In practice, all sensors are colored 1, so that each sensor's bursting time is T_{burst} , if there is any.

It is infeasible to test every random layout, therefore Monte Carlo is carried out. We utilize the feature that the sound source (speaker) has a irregular sound intensity field (see Fig. 6(b)). For each layout of Fig. 10(a) and (b), eight speaker orientations are tried (as shown in Fig. 10(a) and (b), the speaker orientation $\theta = 0^\circ, 45^\circ, 90^\circ, 135^\circ, 180^\circ, 225^\circ, 270^\circ$ and 315°), which equals testing eight different random sensor layouts. For each speaker

orientation, 30 localization trials are carried out.

1) *Localization Accuracy*: As mentioned earlier, for both layouts in Fig. 10(a) and (b), and for each speaker orientation, 30 trials are carried out using Lightning Protocol and another 30 trials are carried out using ideal DP protocol. We use the same metric as defined in Equation (1) (Section VI-B.1) to measure the accuracy of Lightning Protocol, and compare it against the results of DP Protocol. According to Theorem 3, Lightning Protocol may result in more than one winning sensors. In this case, we pessimistically count the sensor farthest away from the sound source as the final winning sensor.

Fig. 11(a.1) and (a.2) shows the localization error statistics of Lightning Protocol and DP protocol respectively with the uniform random layout testbed (see Fig. 10(a)). Fig. 11(b.1) and (b.2) shows the statistics for the clustered random layout testbed (see Fig. 10(b)). From these figures we observe that Lightning Protocol achieves comparable, or even better localization accuracy than DP Protocol. This holds even when there are denser sensors around the sound source, which tends to increase the number of multiple winners of the Lightning Protocol election. The better accuracy of Lightning Protocol is due to its self-synchronization property.

2) *Election Delay*: We also compare the election delays of Lightning Protocol (T_{elec}^{light}) with T_{win1st}^{data} and T_{any1st}^{data} of DP Protocol, where T_{win1st}^{data} is the time cost to send out the first packet from the winner sensor⁶, and T_{any1st}^{data} is the time cost to send out the first packet from any sensor.

Experiment results from both the uniform and clustered random layout testbeds are shown in Fig. 12. Table III compares statistics of T_{elec}^{light} , T_{win1st}^{data} and T_{any1st}^{data} . Clearly, in both random layouts, T_{elec}^{light} is much less than the minimal T_{win1st}^{data} and T_{any1st}^{data} . This indicates that Lightning Protocol, under random layout, still incurs less election delay than any data packet-based schemes.

3) *Number of Transmissions and Winners*: From Section V, we know that owing to the lack of regular layout and coloring, the random layout Lightning Protocol is less efficient in suppressing sensors from contending (bursting) and winning the election. However, experiment results show that the degradation is not significant. Fig. 13(a) and (b) illustrates the statistics of number of bursts per localization trial for the uniform and clustered random layouts respectively. Note that in the clustered random layout, there are more sensors around the sound source. This would

⁶for DP Protocol, there is always only one winner, given that ties are broken randomly.




Fig. 11. Localization Error Comparison with Random Sensor Layout. (a.1) Lightning Protocol with Uniform Random Layout; (a.2) DP Protocol with Uniform Random Layout; (b.1) Lightning Protocol with Clustered Random Layout; (b.2) DP Protocol with Clustered Random Layout.

TABLE III

STATISTICS OF ELECTION DELAY WITH RANDOM SENSOR LAYOUTS

Layout	Metric(ms)	Max	Min	Mean	Standard Deviation
Uniform	T_{elec}^{light}	1.2	1.2	1.2	0
	T_{data}^{win1st}	123.5	15.6	47.3	22.7
	T_{data}^{any1st}	45.7	15.6	18.6	4.1
Clustered	T_{elec}^{light}	1.2	1.2	1.2	0
	T_{data}^{win1st}	99.4	15.6	40.5	18.7
	T_{data}^{any1st}	43.0	15.6	18.7	4.6

further impair the efficiency of burst suppression. However, as observed from the results of Fig. 13(a) and (b), the total number of bursts remains low.

Fig. 14 show the statistics on the number of winners: the majority of the localization trials only elect *one* winner. The percentage of having 2 or 3 winners is low, and there is no trial that elects more than 3 winners.

That Lightning Protocol performs well with random layouts and single coloring can be attributed to the randomness of TOA recognition. Firstly, because of the random layout, the acoustic wave propagation delay to each sensor is randomized. In addition, the TOA is not immediately recognized when the acoustic wave physically reaches a sensor. Instead, there is a random recognition time $t_{recg} \in [0, \Delta_{recg}]$. Both factors contribute to differentiating the time when the bursts start and end. Therefore, majority of redundant RF bursts are still suppressed.

VII. DISCUSSION

A. Multiple Acoustic Events

Lightning Protocol can effectively handle multiple acoustic events (beeps) if the events are separated either temporally by at least $T_{bound}^* + T_{reset}^*$ units of time or spatially by at least 4 times the maximal acoustic audible range (i.e. $4R_{beep}^{max}$, where R_{beep}^{max} is the maximal audible radius from any sound source). The minimal spatial separation requirement is derived as follows (Fig. 15). Suppose beep b_1 and b_2 take place simultaneously, and b_1 and b_2 are $4R_{beep}^{max}$ apart from each other. The most remote sensors that can hear b_1 and b_2 are R_{beep}^{max} away from b_1 and b_2 respectively (denoted as S_1 and S_2 in the Fig. 15). S_2 's RF broadcast only needs to cover a radius of $2R_{beep}^{max}$ to reach all sensors that can hear b_2 . Therefore, the broadcast of S_2 does not interfere with any sensor (say S_1) that can hear b_1 , and vice versa.

Take our experiment testbed settings for example, we assume $R_{beep}^{max} = 20\text{ft}$. this should translate to a separation of beeps by at least 34.3ms in time *or* by at least 80ft in space.

B. Violation of Theoretical Assumptions on Regular Layout

Fig. 8 shows both Lightning and DP protocols may have large error when the sound source is at $(\pm 24, 12)$, $(\pm 24, 0)$ and $(\pm 24, -12)$. This is because our sound source, a directional speaker shown in Fig. 6(a), does *not* fully comply with the bounded directionality assumption (see Section II P2), which says the closest sensor shall always be able to recognize the beep no matter what direction it is oriented toward the sound source. When the speaker is at $(\pm 24, 12)$, $(\pm 24, 0)$ and $(\pm 24, -12)$, the nearest




Fig. 12. Election Delay Comparison with Random Sensor Layouts. (a.1) Lightning Protocol with Uniform Random Layout; (a.2) DP Protocol with Uniform Random Layout: Distribution of T_{win1st}^{data} , i.e. time when the 1st packet is sent out from the winner sensor; (a.3) DP Protocol with Uniform Random Layout: Distribution of T_{any1st}^{data} , i.e. time when the 1st packet is sent out from any sensor; (b.1) Lightning Protocol with Clustered Random Layout; (b.2) DP Protocol with Clustered Random Layout: Distribution of T_{win1st}^{data} , i.e. time when the 1st packet is sent out from the winner sensor; (b.3) DP Protocol with Clustered Random Layout: Distribution of T_{any1st}^{data} , i.e. time when the 1st packet is sent out from any sensor.

sensor is exactly at the speaker’s right/left side (see Fig. 5), where the sound intensity may be too weak, which sometimes causes the sensor’s simple threshold-based algorithm fail to detect the beep. Also, as an early generation of Motes, the MICA Motes’ microphones may occasionally fail to detect the beep, adding another source of error.

Fortunately, Lightning Protocol still works when some sensors fail to detect the beep. Under such cases, the sensor layouts can be regarded as random layouts, which still has good properties. Note when the random layout Lightning Protocol elects multiple winners, we count the winner farthest away from the sound source for localization error. Of all the trials for Fig. 8, we observe only 5.7% trials generate two winners; and no trial generates more than two winners.

C. Multipath Effects

As for wireless multipath, in indoor environments, the common multipath spread is in the order of 0.0001ms. In our Lightning Protocol implementation, $T_{burst} = 0.8\text{ms}$, $T_b = 0.4\text{ms}$, which are way larger than 0.0001ms. This means the wireless channel is flat, and the RF multipath effect can be neglected [9].

As for acoustic multipath, this paper assumes open space or lightly obstructed environment where acoustic multipath effects is negligible (see Section II A4). In practice, our experiments show that Lightning Protocol can tolerate acoustic multipath effects in a common office environment satisfactorily. Lightning Protocol may even handle worse acoustic multipath effects with extra time cost: First, dense deployment of sensors empirically guarantees

LOS between the sound source and its closest sensor. Therefore the closest sensor still bursts first. Second, Δ_{defer} is increased in case some remote sensor does not have LOS to the sound source but can hear the beep arrived through reflected path. Third, T_{reset}^* is increased to prevent sensors from resetting too early to mistake echoes of the beep as new beeps.

VIII. RELATED WORK

Existing solutions to acoustic localization using wireless sensors mainly fall into two categories.

The first category requires the use of some high-capability nodes to conduct triangulation based on readings gathered from a larger population of low-end sensors [10][11][12][1][2][13][14][15][16]. In Sheng et al. [10] and Aslam et al. [11], recursive particle-filtering algorithms are devised to asymptotically converge to a moving target’s track over time. Though convergence is proven, there is no hard time bound on convergence time. In [12][1][2][13], *Maximum Likelihood*(ML)-based localization methods are proposed based on intensity or TOA readings of an array of sensors. In Kim et al. [12], binary readings are aggregated at a central tracking node where regressions are carried out to best fit the weighted set of readings. In Sheng et al. [1], sound sources are assumed to be omni-directional and the attenuation model is known. Each sensor estimates its distance to sound source based on the detected acoustic intensity. In Wang et al. [2], each sensor reports its local TOA, and the algorithm scans the whole monitored area to find the grid point that best matches the TOA readings




Fig. 13. Number of Bursts per Localization Trial. (a) Lightning Protocol with Uniform Random Layout; (b) Lightning Protocol with Clustered Random Layout.




Fig. 14. Number of Winners per Localization Trial. (a) Lightning Protocol with Uniform Random Layout; (b) Lightning Protocol with Clustered Random Layout.

of the sensor array. Simon et al. [13] implement a sniper gun localization system based on TOA of muzzle blast using Maximum Likelihood estimation on consistent TOA readings. Both VigilNet [14] and ExScal [15][16] systems focus on providing holistic middleware architecture solution for moving target tracking. In VigilNet, the target localization module uses magnetic sensors. Every sensor detecting a target within its sensing range reports its coordinates and timestamp. The group leader uses regression to estimate the most likely track of the moving target. In ExScal, all sensors that detect a target within its sensing range report its coordinates and timestamp. For each time window, the leader calculates the centroid of the convex region that envelops all sensors currently detecting the target as the target's location; and uses correlation over successive time windows to estimate the track of the target.

The second category of acoustic localization schemes is based on *proximity-based localization* among homogeneous wireless sensor nodes [17][18][19][20]. In Liu et al. [17] and Blum et al. [18], all sensors detecting an impulsive sound exchange their sound intensity readings or TOAs using multi-hop data communication. The sensor with the best reading wins the election. Chen et al. [19] devise a back-off based method to accelerate the election process. This scheme works well if the sound sources are omnidirectional and of known intensity. Oh et al. [20] propose using Viterbi algorithm and hidden Markovian model to track targets in a sparse wireless sensor network, where sensors' sensory coverages are non-overlapping. This algorithm is mainly designed for tracking. For locating individual acoustic events, it would




Fig. 15. Sufficient spatial separation of two simultaneous beeps

just give the singular sensor that covers the event's location, as sensors' sensory coverages are non-overlapping. Generally speaking, with densely deployed sensors, proximity-based localization achieves acceptable accuracy. More importantly, since proximity-based localization requires only algorithmic comparisons, it is lightweight and can be carried out using only low-end micro-sensors. A proximity-based localization scheme can also serve the purpose of leader election, which is important for activities such as dynamically collaborating sensors around the sound source to achieve higher localization accuracy.

Compared to all the aforementioned approaches, we claim Lightning Protocol is the first scheme that provides a hard real-time ($O(1)$) guarantee on acoustic event localization, and it only involves a few most basic computations (specifically, a few algorithmic comparison operations). The simplicity of Lightning Protocol allows it being implemented on very cheap low-end wireless sensors. The acoustic localization scheme is based on TOA, but unlike conventional TOA based localization schemes which need clock synchronization (simple schemes may need

coarse-grained clock synchronization as those deployed in S-MAC [21] and Z-MAC [22]; sophisticated schemes [2] may need microsecond level clock synchronization, such as RBS [23]), Lightning Protocol eliminates the need of clock synchronization among sensors. Lightning Protocol is also robust to (bounded) directionality of sound sources and variation of sensor density.

This paper extends the previous conference version of [24].

IX. CONCLUSION

In this paper, we exploit the fact that electromagnetic waves propagate much faster than acoustic waves to devise Lightning Protocol, which elects the closest sensor to an acoustic event in a network of low-end wireless sensors. This protocol can be used for proximity-based localization or leader election for sensor collaboration. Both theoretical analysis and experimental results are presented. Lightning Protocol is shown to have a very short and bounded delay ($O(1)$). It only incurs $O(1)$ wireless broadcasts. The majority (81.4% with regular sensor layout and $> 75\%$ with random sensor layout) of our experimental trials incur only one wireless broadcast. The protocol does not involve wireless data packet communication, instead, it deploys RF bursts so that concurrent overlapping wireless broadcasts are allowed. This greatly simplifies the design of its wireless communication module, and makes it faster and more reliable. Moreover, the protocol has little computation and storage complexity, and does not require clock synchronization among distributed sensors. Our experiments demonstrate that the accuracy of Lightning Protocol is comparable to and often better than an ideal DP scheme. The protocol is energy-efficient in the sense that sensor nodes only turn on radio in an on-demand fashion in presence of acoustic events for a constant bounded time. Finally, we demonstrate through experiments that Lightning Protocol can handle directional sound sources with variable intensities, and is empirically feasible in a common office environment.

In our future work, we plan to investigate more effective methods to handle multiple sound sources.

ACKNOWLEDGEMENT

This work is supported in part by MURI N00014-01-0576, by NSF ANI 02-21357, NSF CCR 02-09202, NSF CCR 3-25716, NSF CNS 06-49885 SGER, by ONR N00014-02-1-0102, ONR N00014-05-0739, and by a grant from Lockheed Martin and a grant from Rockwell Collins. The first author is also supported by Vodafone Fellowship. Dr. R. Zheng is also supported by NSF CAREER award CNS-0546391. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of sponsors.

The authors thank Ms. Tanya L. Crenshaw and anonymous reviewers for their advice on improving this paper.

REFERENCES

- [1] X. Sheng *et al.*, "Energy based acoustic source localization," in *Proc. of IPSN'03*, ser. LNCS, vol. 2634. Springer, 2003, pp. 285–300.
- [2] Q. Wang *et al.*, "Acoustic target tracking using tiny wireless sensor devices," in *Information Processing in Sensor Networks: 2nd Intl. Workshop (IPSN 2003)*, ser. LNCS, vol. 2634. Springer, 2003, pp. 642–657.
- [3] Q. Wang. Applying lightning protocol to gunshot localization. [Online]. Available: <http://www-rtsl.cs.uiuc.edu/papers/technicalreports.html>

- [4] (2003, Aug.) Mica2 wireless measurement system datasheet. [Online]. Available: http://www.xbow.com/Products/Product\pdf_files/Wireless\pdf/6020-0042-04_A_MICA2.pdf
- [5] (2004) TinyOS a component-based OS for the networked sensor regime. [Online]. Available: <http://webs.cs.berkeley.edu/tos/>
- [6] (2004, Apr.) RFM TR1000 916.50MHz Hybrid Transceiver. [Online]. Available: http://today.cs.berkeley.edu/tos/hardware/design/data_sheets/RFM.pdf
- [7] F. Cali *et al.*, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit," *IEEE/ACM Trans. on Networking*, vol. 8, no. 6, pp. 785–799, Dec. 2000.
- [8] (2004, Apr.) Lightning Protocol Demo. [Online]. Available: <http://www-rtsl.cs.uiuc.edu/papers/LightningDemo.html>
- [9] T.S.Rappaport, *Wireless communications: principles and practice (2nd Ed.)*. Prentice Hall, 2004.
- [10] X. Sheng *et al.*, "Distributed particle filter with gmm approximation for multiple targets localization and tracking in wireless sensor network," in *Proc. of IPSN'05*, Apr. 2005, pp. 181–188.
- [11] J. Aslam *et al.*, "Tracking a moving object with a binary sensor network," in *Proc. of SenSys'03*, 2003, pp. 150–161.
- [12] W. Kim *et al.*, "On target tracking with binary proximity sensors," in *Proc. of IPSN'05*, Apr. 2005, pp. 301–308.
- [13] G. Simon *et al.*, "Sensor network-based countersniper system," in *Proc. of SenSys'04*, 2004, pp. 1–12.
- [14] T. He *et al.*, "VigilNet: An integrated sensor network system for energy-efficient surveillance," *ACM Trans. on Sensor Networks*, vol. 2, no. 1, pp. 1–38, 2006.
- [15] A. Arora *et al.*, "ExScal: Elements of an extreme scale wireless sensor network," in *Proc. of IEEE RTCSA'05*, 2005, pp. 102–108.
- [16] S. Bapat *et al.*, "Analyzing the yield of ExScal, a large-scale wireless sensor network experiment," in *Proc. of ICNP'05*, 2005, pp. 53–62.
- [17] J. Liu *et al.*, "Distributed group management for track initiation and maintenance in target localization applications," in *Proc. of IPSN 2003*, ser. LNCS, vol. 2634. Springer, 2003, pp. 113–128.
- [18] B. Blum *et al.*, "An entity maintenance and connection service for sensor networks," in *Proc. of MobiSys'03*, 2003.
- [19] W.-P. Chen *et al.*, "Dynamic clustering for acoustic target tracking in wireless sensor networks," *IEEE Trans. on Mobile Computing*, vol. 3, no. 3, pp. 258–271, 2004.
- [20] S. Oh *et al.*, "Tracking on a graph," in *Proc. of IPSN'05*, Apr. 2005, pp. 195–202.
- [21] W. Ye, J. Heidemann, and D. Estrin, "Medium access control with coordinated adaptive sleeping for wireless sensor networks," *IEEE/ACM Trans. on Networking*, vol. 12, no. 3, pp. 493–506, 2004.
- [22] I. Rhee *et al.*, "Z-MAC: a hybrid MAC for wireless sensor networks," in *Proc. of SenSys'05*, Nov. 2005, pp. 90–101.
- [23] J. Elson *et al.*, "Fine-grained network time synchronization using reference broadcasts," in *Proc. of OSDI'02*, Dec. 2002, pp. 147–163.
- [24] Q. Wang *et al.*, "Lightning: A fast and lightweight acoustic localization protocol using low-end wireless micro-sensors," in *Proc. of the 25th IEEE International Real-Time Systems Symposium (RTSS 2004)*, Dec. 2004, pp. 371–381.



Qixin Wang was born in 1977. He received the B.E. and M.E. degrees from the Department of Computer Science and Technology at Tsinghua University, Beijing, China, in 1999 and 2001, respectively. He is currently pursuing the Ph.D. degree in the Department of Computer Science at University of Illinois at Urbana-Champaign. His research interests include real-time/embedded systems and networking, wireless communications, and their applications in industrial control, medicine, and assisted living. Qixin Wang is a student member of IEEE and ACM.



Rong Zheng received her Ph.D. degree from Department of Computer Science, University of Illinois at Urbana-Champaign in May 2004 and earned her M.E. and B.E. in Electrical Engineering in May 1998 and June 1996 from Tsinghua University, P.R. China. She is now an assistant professor at the Department of Computer Science, University of Houston. Zheng is the recipient of the National Science Foundation CAREER Award in 2006. Her current research interests include resource management of large-scale distributed systems, capacity limits and coverage

property of wireless sensor networks, instrumentation and prototyping of wireless systems.



Ajay Tirumala obtained a Ph.D in computer science from the University of Illinois at Urbana Champaign in 2006. He was a member of the Real-Time Systems Laboratory. He is currently a Senior Architect at Nvidia Corporation.



Xue Liu received the BS degree in Applied Mathematics and MEng degree in Control Theory and Applications from Tsinghua University. He received the PhD degree in Computer Science from University of Illinois at Urbana-Champaign in 2006. He is currently an assistant professor in the School of Computer Science at McGill University. He had also worked briefly at HP Labs and IBM T. J. Watson Research Center. His research interests include real-time and embedded computing, performance and power management of server systems, sensor networks, fault tolerance, and control. He has authored/coauthored more than 20 refereed publications in leading conferences and journals in these fields. He is a member of the IEEE.

works, fault tolerance, and control. He has authored/coauthored more than 20 refereed publications in leading conferences and journals in these fields. He is a member of the IEEE.



Lui Sha graduated with Ph.D. from Carnegie Mellon University in 1985 and worked in SEI from 85 to 98. He has been professor and then Donald B. Gillies Chair professor of Computer Science in University of Illinois at Urbana-Champaign since 98. He is a fellow of the ACM and the IEEE. His work on real-time computing was cited as a major accomplishment in National Academy of Science's report: Computing the Future, supported by open standards, and cited as a key element to the success of many national high technology projects including

GPS upgrade, the Mars Pathfinder, and the International Space Station.