

Over-the-Air Multi-Sensor Collaboration for Resource Efficient Joint Detection

Carlos Feres, *Member, IEEE*, Bernard C. Levy, *Life Fellow, IEEE*, Zhi Ding, *Fellow, IEEE*

Abstract—We develop a resource-efficient framework for collaborative decision-making over distributed sensor networks by proposing a novel over-the-air soft information aggregation. We exploit the natural superposition of wireless transmissions to enable sensors to utilize over-the-air computation to approximate the sufficient statistic for optimum detection over a shared channel. By designing practical transmission and receiver processing in over-the-air computation, the decision-making fusion center can wirelessly obtain a good approximation of the aggregate log-likelihood ratio computed over all observed data with low distortion. Focusing on Neyman-Pearson tests for detection in this new framework, we develop efficient tests and analyze their performance bounds in several common joint detection scenarios. Our results show significant over-the-air collaboration gain even with a few participating sensors. The novel framework exhibits very little performance loss of detection accuracy against traditional multiple access transmission from sensing nodes despite substantial resource savings via over-the-air computation.

Index Terms—Internet of Things, decision-making, collaborative learning, soft information.

I. INTRODUCTION

Internet of Things (IoT) comprises several technologies, protocols, and applications that enable advanced event monitoring, enhance productive processes, and a broad range of diverse solutions [1]. An important class of such applications consists of the detection of distributed phenomena using observations from several sensing nodes (also denoted *sensors*), which share local data with a fusion center node (also denoted *server*). Of course, with an increasing number of sensing nodes, the available resources must be optimized to allow all nodes to collaborate. On the other hand, many IoT applications are only feasible when they are able to operate for long periods of time with no battery replacement, making energy consumption a critical design criterion. However, tackling these challenges should not diminish the importance of guaranteeing the satisfactory performance of the network.

This work aims to develop a practical and simple framework for distributed sensing and integrated decision-making with high resource efficiency and strong performance in terms of decision accuracy. This framework is based on the concept of over-the-air computation for distributed agents/sensors in a network. Given S sensing nodes, the resource efficiency over traditional multiple access protocols for sensor transmission

grows by S -fold. Equally important is the fact that our proposed over-the-air computation allows active sensors to transmit individual local log-likelihood ratio (LLR), which forms a natural and near-optimum approximation of the sufficient statistic (SS) for joint detection and decision-making.

Several works have analyzed detection performance in distributed sensor networks with traditional centralized grant-based multiple access. For example, the authors of [2] studied the detection problem in an IoT network where sensors send decentralized local decisions via a multiple access channel (MAC) to a fusion center, and propose optimal decision-making (in the sense of achievable detection performance) by aggregating as many sensors as the channel rate admits. [2] assumes perfect delivery of local decisions and provides an analytical framework for distributed detection in IoT. Other studies have presented generalizations of this IoT detection framework, including unequal data rates for each sensor [3], Gaussian channel noise effect in sensor decision transmissions [4], coherent channel fading [5], [6] in sensor transmissions, sequential testing [7], correlated sensor observations [8], censoring nodes [9], sparsity of sensor signals [10], [11], and optimality of local decisions [12], among several other extensions and works. In particular, [5] illustrates that under coherent channel fading, equal gain combining (EGC) of local decisions performs well for a wide range of signal-to-noise-ratio (SNR). Furthermore, the authors in [6] proposed a near-optimum aggregation of local sensor decisions based on known Rayleigh fading channel statistics.

However, these works favor traditional sensor transmission of local decision variables over a MAC. This approach, though effective and widely studied, requires high bandwidth use, particularly with increasing number of sensors. The server needs to coordinate network access via scheduling and solicit local signals one sensor at a time. It is thus vital to design efficient protocols that can significantly reduce resource usage in distributed detection without significant performance loss.

As an alternative to traditional polling, grant-free access techniques can improve spectral efficiency. One such technique is blind signal separation [13]–[16], which exploits spatial diversity and high-order statistics to recover simultaneous signals without using pilot sequences. However, these techniques are usually computationally costly, require a large number of samples to identify statistics accurately, and are constrained by the number of antennas of the server node. Thus, these techniques are not appealing for general IoT deployments.

We note that for collaborative decision-making in systems including IoT, data fusion techniques such as over-the-air computation (AirComp) [17], [18] provide a low-

This material is based upon work supported by the National Science Foundation under Grants No. 2002937, No. 2009001 and No. 2029037.

C. Feres was with the University of California, Davis, CA 95616, and is now with Samsung Semiconductor, Inc., San Diego, CA, 92121, USA (email: cjferes@ucdavis.edu). B. C. Levy and Z. Ding are with the Department of Electrical and Computer Engineering, University of California, Davis, CA, 95616, USA (email: bclevy@ucdavis.edu; zding@ucdavis.edu).

cost alternative in saving resources. AirComp has recently been considered for communication-heavy applications such as cooperative spectrum sensing [19] and Federated Learning [20], [21]. Simply put, AirComp exploits the natural aggregation of analog signals simultaneously transmitted by multiple collaborative nodes onto a shared wireless channel. Hence, AirComp improves resource efficiency by utilizing the same channel resource for all participating nodes regardless of their total population size. Based on minimal network coordination and pre- and post-compensation of signals, AirComp provides significant resource savings and relies on very simple access protocol control, which is well suited for collaborative IoT networks. In fact, a simpler aggregation of signals for collaborative estimation in sensor networks was proposed in [22]. As a predecessor of AirComp, the formulation of [22] assumes ideal channels with no fading, which is an important limitation for practical implementation and thus, does not require compensation mechanisms.

In our previous work [23], we developed over-the-air computation to aggregate local measurements from sensors. Thanks to AirComp, sensor data sent in analog signals are naturally superimposed over the MAC. To receive data, the server broadcasts a request, and all sensors respond under synchronization, to simultaneously transmit respective local decision variables. Using different compensation schemes to account for the effect of channel phase and channel noise, we showed that over-the-air sensor collaboration effectively saves resources with mild loss of performance in comparison with scheduled orthogonal sensor transmissions. We further extended such AirComp sensor fusion in [24], where sensors simultaneously transmit locally acquired log-likelihood ratios. At the fusion center, the received signal is a noisy-and-channel weighted sufficient statistic (NCWS) that serves as an approximation of the ideal cumulative LLR over all sensor data. Within the Bayesian paradigm, we showed that under mild conditions and properly designed pre- and post-processing functions, the proposed over-the-air collaboration achieves near-optimum performance and significant resource savings. However, a Bayesian detection is a special case that requires knowledge of some key parameters. For broader real-world applications, we often need to account for possibly unknown decision variables and channel parameters.

In this work, we study over-the-air aggregation for IoT decision-making in more practical scenarios where some parameters might be unknown. To this avail, we adopt a general Neyman-Pearson (NP) framework with the objective of approximating optimum NP performance while achieving substantial resource saving through AirComp, which we denote AirComp-based Federated Decision Making (AirCompFDM). We design several different NP detection tests depending on the knowledge of system and channel parameters. We determine whether or not these NP tests are universally most powerful (UMP), or whether or not we may devise UMP invariant (UMPI) tests without UMP. For performance analysis, we derive analytical approximations of the relative entropy of these tests with respect to the number of participating nodes. This entropy corresponds to the error exponent of both probability of missed detection and false alarm. Our




Fig. 1: Illustration of a sensor network using our proposed over-the-air decision-making paradigm. Sensors observe the event H and transmit log-likelihood ratios via analog signals for over-the-air aggregation at the server.

results show that detection performance has significant gains even with only a modest number of sensors. Additionally, our simple over-the-air protocols and minimal resource usage make our proposed AirCompFDM both feasible and enticing.

This paper is structured as follows. Section II details the system model with over-the-air computation and the use of soft information. Section III presents a classical NP test for AirCompFDM and compares its performance to ideal and traditional schemes. In Section IV we further derive detection tests for means of Gaussian sensor observations depending on the knowledge (or not) of key parameters. Section V studies performance gains thanks to sensor aggregation for all these tests. We corroborate our findings with numerical experiments on Section VI, and Section VII concludes the work.

Notations: In the following, vectors will be denoted with small boldface letters, such as $\mathbf{z}$, with transpose denoted by $\mathbf{z}^T$. Sets are denoted with calligraphic capital letters. $\bar{z}$ represents the complex conjugate of z , and the imaginary unit is denoted as j . $\mathbf{1}$ and $\mathbf{I}$ represent a vector of ones and the identity matrix of appropriate size. Probability, expectation and variance are denoted as $\mathbb{P}(\cdot)$, $\mathbb{E}\{\cdot\}$ and $\text{Var}\{\cdot\}$, respectively.

II. SYSTEM MODEL

Consider a wireless system of single-antenna nodes, where a server node hosts S sensors, as depicted in Fig. 1. The sensors collect data about a particular event H belonging to a discrete set $\{H_0, H_1\}$, which correspond to the two underlying hypotheses. Each sensor i generates N observations $\mathbf{v}_i = [v_{i,1} \ \dots \ v_{i,N}]^T$, for $i \in \mathcal{S} = \{1, \dots, S\}$, under H_0 or H_1 . The observations of the i -th sensor follow conditional probability distribution function (PDF)

$$f_i(\mathbf{v}_i | \theta_j, H_j), \quad j \in \{0, 1\}, \quad i \in \mathcal{S}, \quad (1)$$

conditional on parameter vector θ_j under hypothesis H_j . The task of the server is to decide on a hypothesis, using a test δ over signal z that aggregates information sent from the sensor nodes, subject to system design and channel characteristics in different scenarios (to be specified in Sections II.A through II.C). We then define the parameter vector Θ_j , which contains all relevant system and channel parameters after aggregation.

Now, the technical challenge is to achieve satisfactory detection performance with high efficiency and minimal resource usage. We propose the use of i) over-the-air computation to drastically reduce resource requirements and network coordination; ii) letting nodes transmit local soft scores, maximizing the amount of shared information; and iii) composite testing to account for practical scenarios where the server may not know all the parameters in Θ_j .

Throughout this section, we will describe the key aspects of our framework, namely, over-the-air computation, transmission of soft information, and channel compensation techniques. In Section III we formulate an NP framework for over-the-air decision-making, and further extend it in Section IV where we address composite test design in different practical scenarios.

A. Over-the-Air Aggregation

We follow the model and formulation of [24]. In each transmission slot, sensors simultaneously transmit analog signals to the server over a MAC. We assume that all sensors have acquired network timing and are synchronized at the server, e.g., via round-trip delay information, such that their transmitted signals would aggregate synchronously at the receiving server. Furthermore, we assume that each burst duration is below the coherence time of the wireless channel such that channel gains remain constant within each transmission slot.

AirComp aims to compute an estimation or decision from a nomographic function of distributed data collected locally by participating sensors [17]. At discrete resource (e.g. time slot) k , each sensor i computes a message $u_i(k) = u_i(\mathbf{v}_i; k)$ and transmits a signal $x_i(k)$, $i \in \mathcal{S}$, corresponding to a local pre-processing of message $u_i(k)$, i.e.

$$x_i(k) = \psi_{i,k}(u_i(k)), \quad i \in \mathcal{S}. \quad (2)$$

When all sensors simultaneously transmit individual signals over a shared multiple access channel, AirComp allows the server (fusion) node to receive the cumulative signal

$$y(k) = \sum_{i \in \mathcal{S}} g_i(k)x_i(k) + n(k), \quad (3)$$

where $g_i(k) \in \mathbb{C}/\{0\}$ represents the channel state information (CSI) from the i -th node to the server and we assume channel noise $n(k)$ to be circularly symmetric complex additive white Gaussian noise (AWGN) with power density ω^2 (i.e. variance), independent of channels and signals. From (3), AirComp improves resource efficiency by a factor of S , since a single time slot is required for sensor access instead of S time slots required for sequential sensor polling. Moreover, the AirComp access protocol remains the same regardless of how many sensors may decide to collaborate in a communication round.

The server collects $K \geq 1$ samples of the received signal (3) and makes a decision on a hypothesis using post-processed

samples $z(k) = \Psi_k(y(k))$. Note that if the K samples are obtained within a transmission slot, they all experience the same channel gains. Conversely, if samples are taken in different transmission slots, we assume that channel realizations are independent over slots, yielding independent samples $y(k)$.

B. Sufficient Statistic and Over-the-Air Approximation

In any detection problem, the LLR computed over all observed data corresponds to the optimal sufficient statistic for hypothesis testing [25], [26]. If the server had access to all observations $\mathbf{v}_i$ from S independent sensors, the LLR is

$$\ell(\mathbf{v}_1, \dots, \mathbf{v}_S) = \sum_{i \in \mathcal{S}} \ln \left(\frac{f_i(\mathbf{v}_i | \theta_1)}{f_i(\mathbf{v}_i | \theta_0)} \right) = \sum_{i \in \mathcal{S}} \ell(\mathbf{v}_i), \quad (4)$$

where $\ell(\mathbf{v}_i) = \ln \left[\frac{f_i(\mathbf{v}_i | \theta_1)}{f_i(\mathbf{v}_i | \theta_0)} \right]$ is the LLR computed only with local observations of sensor i . Classical approaches consider sensors that share local decisions using their own $\ell(\mathbf{v}_i)$ [12]. Given (4), if sensors instead transmit their local LLRs by defining $u_i(k) = \ell(\mathbf{v}_i)$, the server only needs to sum them to form a SS λ for optimal decision-making with the test

$$\lambda(\mathbf{v}_1, \dots, \mathbf{v}_S) = \sum_{i \in \mathcal{S}} u_i(k) = \sum_{i \in \mathcal{S}} \ell(\mathbf{v}_i) \stackrel{H_1}{\underset{H_0}{\gtrless}} \eta. \quad (5)$$

Note that using AirComp in ideal noiseless scenarios, the received signal at the server is exactly λ : AirComp yields SS (5) with minimal resource usage. However, in practical scenarios AirComp exhibits amplitude and phase distortions $g_i(k)$ for each sensor signal, plus the effect of channel noise $n(k)$, as depicted in (3). In other words, the received signal $y(k)$ is a sample of a noisy and channel-weighted statistic (NCWS), and the technical challenge is to design pre- and post-processing functions $\psi_{i,k}$ and Ψ_k to obtain an NCWS $\tilde{\lambda}$ that approximates the SS, i.e.

$$\tilde{\lambda}(z) \approx \lambda \quad \text{where} \quad z = [z(1) \quad \dots \quad z(K)]^T$$

$$\text{and} \quad z(k) = \Psi_k \left(\sum_{i \in \mathcal{S}} g_i(k) \psi_{i,k}(u_i(k)) + n(k) \right). \quad (6)$$

C. Channel Compensation in Over-the-Air SS Approximation

Depending on network deployment and practicality of design and implementation, we can devise different compensation schemes for AirCompFDM. In particular, we realize that channel phase compensation is critical for correct decision-making, and although interesting for analysis, the overall impact of channel magnitude compensation is small once the channel phase has been dealt with, as demonstrated by the robustness of the EGC fusion rule [5], [6]. Hence, we focus on channel phase compensation schemes only in this work.

1) *Exact Phase Pre-compensation at Sensors*: If there exists channel reciprocity such as in a time-division duplex (TDD) link between server and sensors, the server broadcast used to acquire timing can also be used to estimate channel phase. If we assume that the estimation is perfect, the sensors can use the pre-processing functions

$$x_i(k) = \frac{\overline{g_i(k)}}{|g_i(k)|} u_i(k), \quad (7)$$

and denoting $a_i(k) = |g_i(k)| > 0$, the received signal is

$$y(k) = \sum_{i \in \mathcal{S}} a_i(k) u_i(k) + n(k), \quad (8)$$

and as the compensated channels are real, the server uses post-processing functions $z_P(k) = \text{Re}\{y(k)\}$. We call this protocol AirCompFDM-P. Note that this setting is similar to the proposed EGC combining of [5], [6], although here the sum of samples is performed over-the-air.

2) *Exact Phase Post-compensation at Server*: In the opposite scenario where there is no channel reciprocity, sensors cannot estimate channels beforehand and are unable to pre-compensate their local LLR signals. Thus, sensors do not perform any pre-processing and we simply let $x_i(k) = u_i(k) \in \mathbb{R}$, resulting in an improper (complex) received signal

$$y(k) = \sum_{i \in \mathcal{S}} g_i(k) u_i(k) + n(k) \in \mathbb{C}. \quad (9)$$

Regardless, we assume that the server is able to estimate the aggregated channel at each sample, $G_k = \sum_{i \in \mathcal{S}} g_i(k) \in \mathbb{C}$, and compensates for its resulting phase. Under this protocol, denoted AirCompFDM-U, we define post-compensated signal samples as

$$z_U(k) = \text{Re}\left\{\frac{\overline{G_k}}{|G_k|} y(k)\right\}. \quad (10)$$

3) *Quantized Phase Pre- and Post-compensation*: Finally, we can consider non-ideal channel phase estimation and we apply discretely quantized phase corrections on both ends. We call AirCompFDM-M when sensors and server perform quantized phase compensation with M uniform angle partitions of $2\pi/M$ radians, with offset of π/M . In particular, $M = 4$ quantizes channel phase into 4 quadrants before compensation. Formally, and assuming TDD channel reciprocity, the i -th sensor estimates channel phase $\angle g_i(k)$ and computes the corresponding t_i -th angle partition, with $t_i \in \{1, \dots, M\}$. In particular, t_i is such that

$$\angle g_i(k) + \frac{\pi}{M} \pmod{2\pi} \in \left[\frac{2\pi(t_i - 1)}{M}, \frac{2\pi t_i}{M}\right), \quad (11)$$

and the sensor then uses the quantized phase precompensation

$$x_i = \exp(-\imath \phi_i) u_i, \quad \phi_i = \frac{2\pi(t_i - 1)}{M}. \quad (12)$$

Denoting $b_i(k) = g_i(k) \exp(-\imath \phi_i)$, the received signal is

$$y(k) = \sum_{i \in \mathcal{S}} b_i(k) u_i(k) + n(k). \quad (13)$$

We apply a similar approach of phase post-compensation with

$$z_M(k) = \text{Re}\{\exp(-\imath \Phi_k) y(k)\}, \quad (14)$$

where Φ_k discretely compensates the aggregated channel gain $\sum_{i \in \mathcal{S}} b_i(k)$ similarly to the quantized precompensation above. We note that this model can be further generalized to consider either only pre-compensation or only post-compensation, according to practical network constraints.

Furthermore, quantized phase precompensation is also possible when there is no TDD channel reciprocity. In this setting, the server would estimate channel phase of each sensor based

on traditional pilots, and then send the quantized angle back to the sensor or user equipment (UE) by using only $\log_2(M)$ bits. Though this step requires additional resource consumption, it only needs to be performed once for the coherence time period of the channel. As wireless sensor networks are typically static, this is still an attractive option in practical deployments.

III. NEYMAN-PEARSON DETECTION OF SENSOR MEANS

Here we design different tests for collaborative decision making under a NP framework. In the following, we assume that local measurements are independent between sensors and that they all follow the same multivariate Gaussian distribution parametrized by $\theta_j = (m_j, \sigma^2)$ conditioned on hypothesis H_j , i.e., the hypotheses differ only in sensor mean, i.e. $\mathbf{v}_i \sim \mathcal{N}(m_j \mathbf{1}, \sigma^2 \mathbf{I})$, and we further assume that $m_1 > m_0$. Hence, the sensor LLRs correspond to

$$\ell(\mathbf{v}_i) = \frac{1}{2\sigma^2} \left(\|\mathbf{v}_i - m_0 \mathbf{1}\|^2 - \|\mathbf{v}_i - m_1 \mathbf{1}\|^2 \right). \quad (15)$$

Letting $d = (m_1 - m_0)\sqrt{N}/\sigma > 0$ denote sensor hypothesis distance, we have

$$\ell(\mathbf{v}_i) | H_j \sim \mathcal{N}\left((-1)^{1-j} \frac{d^2}{2}, d^2\right), \quad (16)$$

which means that the distribution of local LLRs is simply

$$u_i | \theta_j \sim \mathcal{N}(0.5 q_j d^2, d^2), \quad \theta_j = (q_j, d), \quad q_j = (-1)^{1-j}. \quad (17)$$

Note that if we instead assume $m_0 > m_1$, the signs of the means would be inverted, i.e. $q_0 = 1$ and $q_1 = -1$. Hence, we focus only on the case $m_0 < m_1$ without loss of generality. Note that the test $H_0 : m = m_0$ vs. $H_1 : m = m_1 > m_0$ is equivalent to the test $H_0 : q = q_0 = -1$ vs. $H_1 : q = q_1 = 1$.

We have assumed that sensors have equal sensing quality, related to d , whereas in real-world scenarios this would be unlikely. However, this assumption is not impractical: as d increases with the number of local observations, a bad sensor only needs to consolidate more observations into its local LLR for quality improvement, which applies to *any* underlying distribution. Alternatively, a practical design approach would be to presume a worst-case scenario where all sensors have quality equal to the sensor with the worst quality, which corresponds to our model above. Such design would eventually lead to more robust decision-making by e.g. careful selection of system parameters. Moreover, there is evidence of satisfactory detection performance with simple fusion rules like EGC for sensors with different quality [6, Fig. 6]. These reasons should validate our model with equal sensing quality for all nodes.

Without loss of generality and for ease of exposition, in the rest of this work we shall consider only static channels, as the extension to dynamic channels is straightforward. Hence, for the different compensation schemes of AirCompFDM, the resulting conditional distributions of the samples of the post-processed signal $\mathbf{z}(k)$ are

$$z_P(k) | H_j \sim \mathcal{N}(0.5 q_j A d^2, P d^2 + 0.5 \omega^2), \quad (18a)$$

$$z_M(k) | H_j \sim \mathcal{N}(0.5 q_j B d^2, P d^2 + 0.5 \omega^2), \quad (18b)$$

$$z_U(k) | H_j \sim \mathcal{N}(0.5 q_j |G| d^2, P d^2 + 0.5 \omega^2), \quad (18c)$$



(a) ROC of the tests, with SNR -10 dB.(b) Probability of detection vs. channel SNR, with $P_F = 0.1$.

Fig. 2: Performance of the test δ_1 for different means for known m_0 and m_1 , known channel parameters, and with channel SNR -10 dB, $d = 2$, $S = 2$, and $K = N = 1$. We average 1000 realizations over 1000 independent sensor deployments.

where the sum of channel powers $P = \sum_{i \in \mathcal{S}} |g_i|^2 > 0$ and

$$A = \sum_{i \in \mathcal{S}} a_i = \sum_{i \in \mathcal{S}} |g_i| > 0, \quad (19a)$$

$$B = \text{Re} \left\{ \exp(-i\Phi) \sum_{i \in \mathcal{S}} b_i \right\} \geq 0, \quad (19b)$$

$$G = \sum_{i \in \mathcal{S}} g_i \in \mathbb{C}, \quad (19c)$$

are the sum of channel gains under each compensation scheme. Note that these aggregated gains are not independent of P . We assume that M is large enough to ensure $B \geq 0$. Then, parameter vector Θ_j is

$$\Theta_{j,P} = [A \quad P \quad \omega^2 \quad d^2 \quad q_j]^\top, \quad (20)$$

under perfect channel precompensation, and we obtain $\Theta_{j,M}$ and $\Theta_{j,U}$ by replacing A by B and $|G|$, respectively.

We now derive a clairvoyant NP test where all parameter values are known exactly when implementing the test. Moreover, as the derivation of NP tests is similar for different AirCompFDM compensation schemes, we only show the detailed derivation of tests using AirCompFDM-P, and summarize results for all other schemes. We assume first that all relevant parameters are known. Hence the likelihood ratio test (LRT) at the server is directly computed as

$$L(z) = \exp \left(\frac{\|z + \frac{Ad^2}{2} \mathbf{1}\|^2}{2Pd^2 + \omega^2} - \frac{\|z - \frac{Ad^2}{2} \mathbf{1}\|^2}{2Pd^2 + \omega^2} \right) \underset{H_0}{\overset{H_1}{\geq}} \tau, \quad (21)$$

reducing to the test

$$\delta_1(z) : \tilde{\lambda}_1 = \frac{\mathbf{1}^\top z}{K} \underset{H_0}{\overset{H_1}{\geq}} \sqrt{\frac{2Pd^2 + \omega^2}{2K}} Q^{-1}(\alpha) - \frac{Ad^2}{2}, \quad (22)$$

which has size α and power

$$P_D(\alpha) = 1 - Q(D_P - Q^{-1}(\alpha)). \quad (23)$$

Here we denote the server hypothesis distance under perfect channel phase precompensation as

$$D_P = \frac{Ad^2 \sqrt{K}}{\sqrt{Pd^2 + 0.5\omega^2}}. \quad (24)$$

We obtain similar tests for AirCompFDM- M and - U , with respective thresholds and server hypothesis distances.

As a proof-of-concept demonstration, in Fig. 2, we compare the performance of the NP test $\delta_1(z)$ using all compensation schemes, with $S = 2$ sensors. We assume sensor hypothesis distance $d = 2$ for $K = 1$ and $N = 1$ per sensor, at channel SNR of -10 dB and Rayleigh fading, i.e. $g_i \sim \mathcal{CN}(0, 1)$. Moreover, we compare with the performance bound obtained by assuming a perfect SS in centralized LLR testing, and under a noisy case without channel fading (equivalent to full channel pre-compensation, akin to the method of [22], denoted NSS). To properly capture average performance, we generate random variables in two steps. First, we generate 1000 “sensor deployments”, each consisting of a set of S channels. Then, for each deployment, we generate 1000 realizations of NK observations per sensor, computing K local LLRs per sensor with N observations each. Our results include i) average empirical curves (“empirical”), analytical curves using average D over deployments (“deployment average”), ii) theoretical curves using the approximate expectation of D in (23), i.e. $P_D(\mathbb{E}\{D\})$, for both AirCompFDM-P and -U [24]. We relegate more detailed discussions on testing to Section VI.

The empirical receiver operating characteristic (ROC) is depicted in Fig. 2a. Interestingly, for low channel SNR, low number of sensors, and low d , AirCompFDM-P is close to the performance of fully precompensated channels NSS, and more importantly, achieves comparable performance to a traditional MAC polling scheme of local decisions where channel fading and noise have been perfectly mitigated [2], using a maximum *a posteriori* detector. This detector is equivalent to a majority rule [27] over local decisions $u_i^{\text{poll}}(\tau) \in \{0, 1\}$, computed




Fig. 3: Empirical ROC of test δ_1 with known m_0 and m_1 , known channel parameters, and with $K = N = 1$. We average 1000 realizations over 1000 independent sensor deployments. (a) to (c): Varying sensor hypothesis distance d , with SNR -10dB. (d) to (e): varying SNR with $d = 1$.

using a common threshold τ across sensors. Hence, this detector has theoretical ROC curves

$$P_{F,\text{poll}}(\tau) = \mathbb{P}\left(\sum_{i \in S} u_i^{\text{poll}}(\tau) \geq \frac{S}{2} \middle| H_0\right), \quad (25a)$$

$$P_{D,\text{poll}}(\tau) = \mathbb{P}\left(\sum_{i \in S} u_i^{\text{poll}}(\tau) \geq \frac{S}{2} \middle| H_1\right). \quad (25b)$$

We also note that quantized phase pre-compensation of AirCompFDM- M achieves surprisingly good performance with just $M = 4$ angle regions, being only slightly worse than AirCompFDM-P in all cases. On the other hand, AirCompFDM-U shows worse performance, expected due to the detrimental effect of noncoherent channel aggregation. Furthermore, theoretical approximations are close and slightly overestimate their corresponding empirical performance.

In Fig. 2b we show probability of detection with respect to channel SNR, for equal parameters and a fixed probability of false alarm of 10%. Probability of detection for SS and traditional MAC polling do not depend on channel SNR but are shown across the range for readability. As expected, increasing SNR improves performance regardless of the compensation scheme. For high SNR, AirCompFDM-P only has a small gap with respect to the optimal SS. Moreover, AirCompFDM- M with discrete precompensation of channel phase achieves very similar performance to exact phase precompensation. AirCompFDM-U suffers in performance in comparison, but still achieves performance close to that of traditional MAC polling for higher channel SNR.

In Fig. 3 we show the empirical ROC of test $\delta_1(z)$ under different compensation methods, for different number of

participating sensors S , varying sensor hypothesis distance d (3a-3c) and varying SNR (3d-3f). Overall, the test exhibits performance gains with increasing S in all cases, which are significant for AirCompFDM-P and AirCompFDM- M that use phase precompensation. On the other hand, AirCompFDM-U stagnates and suffers in comparison to either phase precompensation scheme, due to noncoherent channel aggregation. In particular, Figs. 3a to 3c show the ROC for different values of sensor hypothesis distance d for a fixed channel SNR of -10dB. As expected, test performance improves with increasing d across compensation methods. More interestingly, Figs. 3d to 3f depict performance for different SNR values and fixed $d = 1$. In all cases, increasing the number of participant sensors improves test performance beyond the gains obtained by increasing SNR with a lower number of sensors. This fact illustrates the impact of sensor aggregation, overcoming the traditional trade-off between detection performance and energy consumption due to increased transmission power.

Furthermore, we want to highlight that AirCompFDM provides aggregation gains for different distributions of sensor observations. Regrettably, the LLRs of most distributions are usually challenging to derive, except in special cases. Thus, the distribution of the received signal (itself a post-processed version of a weighted sum of LLRs plus Gaussian noise) is usually extremely hard or impossible to obtain, and we are unable to analytically find an appropriate NCWS and its conditional distributions, or even design a test beyond numerical simulations. Thus, in Fig. 4 we offer numerical evidence of the applicability of AirCompFDM in more general scenarios, for all our proposed compensation methods with




Fig. 4: Empirical ROCs for different distributions of sensor observations and channel SNR, with $K = N = 1$. We average 1000 realizations over 1000 independent sensor deployments.

$K = N = 1$, i.e. the least amount of observations and samples, and compare with the perfect SS in centralized LLR testing. We first test $H_0 : \sigma = \sigma_0$ vs. $H_1 : \sigma = \sigma_1 > \sigma_0$ for zero-mean Gaussian observations, i.e. $v_i|H_j \sim \mathcal{N}(\mathbf{0}, \sigma_j^2 \mathbf{I})$ i.i.d. $\forall i \in \mathcal{S}$, with $\sigma_0^2 = 1$ and $\sigma_1^2 = 2$. Figs. 4a and 4b show the ROCs of this test under channel SNR of -10 and 0 dB, respectively. In this test, local LLRs follow biased Gamma distributions, and the distribution of samples $z(k)$ has no closed form. However, the LLR of post-processed signals can be obtained using numerical inversion of the characteristic function under each hypothesis. Comparing Figs. 4a and 4b show the impact of channel noise on detection performance, but AirCompFDM-P and -M still perform adequately for low channel SNR, with performance comparable to ideal SS under 0 dB of channel SNR. Similarly, Fig. 4c depicts the empirical ROCs of test $H_0 : m = m_0$ vs. $H_1 : m = m_1 > m_0$ when sensor observations are contaminated by Laplacian noise, i.e. $v_{i,n}|H_j \sim \text{Laplace}(m_j, \beta)$ i.i.d. $\forall n \in \{1, \dots, N\}, \forall i \in \mathcal{S}$, with antipodal locations $m_0 = -m_1 = -0.3$ and $\beta = 1$, and channel SNR of 0 dB. Here, the distribution of a local LLR is known [28], but the NCWS still is unobtainable in analytical form and its conditional distributions are computed via numerical inversion. In all the cases above, AirCompFDM-P and -M show significant aggregation gains, whereas AirCompFDM-U struggles but still improves performance slightly by mitigating the effect of channel noise. Nevertheless, these results show that AirCompFDM provides collaborative detection regardless of the underlying distribution of sensor measurements, illustrating the generality of our proposed framework.

IV. COMPOSITE NEYMAN-PEARSON TESTING

In this section, we derive further NP tests for more practical scenarios where some parameters in θ_j are unknown, i.e., the server has no access to the exact value of the parameter(s) when implementing the test. Moreover, in these scenarios, it is of great interest to determine whether we can obtain a universally most powerful (UMP) test, and if not, if it is possible to obtain a UMP invariant (UMPI) test instead. When an UMP(I) test is derived, it will exhibit this property regardless of compensation scheme, although these tests are not necessarily optimal performance-wise.

In case of practical deployments, the channel parameters A (corresp. B or $|G|$), P and ω^2 are not known *a priori*, and the server needs to estimate them. In this work, whenever we assume knowledge of a parameter, we mean that a parameter has been estimated perfectly. For completion purposes, we present some simple estimation procedures for each parameter:

- To estimate receiver AWGN noise intensity ω^2 , the server node uses samples of the received signal *before* broadcasting the transmission request to sensors.
- To estimate P , the receiver node can request that sensor transmissions use some known pilot signals and estimate the power of the incoming signal. By knowing the pilot signals and ω^2 , P can be accurately estimated.
- Finally, an estimation of aggregated channel gain A (corresp. B or $|G|$) is also obtained using pilot signals. We can further refine the estimate of sum-gain or sum-power using the other estimate, or a joint estimation procedure, or an iterative process.

To derive analytical insights, we assume again that sensor observations are normally distributed with different means as in the previous section, and that sensor variances are known and equal under both hypotheses. Hence, the distributions of the samples of the post-processed signal correspond to (18) depending on the compensation scheme. We first assume that both sensor means m_0 and m_1 are known. Afterwards, we determine tests for the case when m_1 is unknown and only m_0 is known.

A. Known sensor means

With both m_0 and m_1 known, the following tests are designed for testing $H_0 : m = m_0$ vs. $H_1 : m = m_1 > m_0$. Furthermore, the sensor hypothesis distance d is also known.

1) *Unknown aggregated gain:* We first assume that P and ω^2 are known, but the aggregated gain parameter A (corresp. B or $|G|$) is unknown. Any test design with transmitted signals (17) would not yield UMP tests, as the distributions of the server samples (18) result in a probability of false alarm that depends on the unknown gain parameter. However, as d is known, we can obtain a UMP test if instead the sensors transmit biased LLRs $u_i = \ell(v_i) + d^2/2$. Thus, we have that

$$u_i|H_j \sim \mathcal{N}(jd^2, d^2) \quad (26)$$

and

$$z_P|H_j \sim \mathcal{N}(jAd^2, Pd^2 + 0.5\omega^2), \quad (27a)$$

$$z_M|H_j \sim \mathcal{N}(jBd^2, Pd^2 + 0.5\omega^2), \quad (27b)$$

$$z_U|H_j \sim \mathcal{N}(j|G|d^2, Pd^2 + 0.5\omega^2). \quad (27c)$$

Transmitting biased LLRs, the mean of the server samples under H_0 is zero in all compensation schemes, and hence the unknown gain parameter is not part of H_0 . This allows obtaining a threshold for a given probability of false alarm α . The server LRT yields

$$L(z) = \exp\left(\frac{\|z\|^2}{2Pd^2 + \omega^2} - \frac{\|z - Ad^2\mathbf{1}\|^2}{2Pd^2 + \omega^2}\right) \underset{H_0}{\overset{H_1}{\gtrless}} \tau, \quad (28)$$

reducing to the test

$$\delta_2(z) : \tilde{\lambda}_2 = \frac{\mathbf{1}^\top z}{K} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_2 = \sqrt{\frac{2Pd^2 + \omega^2}{2K}} Q^{-1}(\alpha), \quad (29)$$

which is UMP with power

$$P_D(\alpha) = 1 - Q\left(\frac{Ad^2\sqrt{K}}{\sqrt{Pd^2 + 0.5\omega^2}} - Q^{-1}(\alpha)\right). \quad (30)$$

We have a similar result with other compensation schemes.

2) *Unknown channel parameters*: We assume now that A (corresp. B or $|G|$) unknown, and either P or ω^2 are unknown as well. Hence, there is no UMP test as test threshold and P_F would depend on P and ω^2 . However, we can obtain an UMPI test with a generalized LRT (GLRT). Using biased LLRs as in (26), knowing that the aggregated gain A (corresp. B or $|G|$) are positive, and using the maximum likelihood estimate (MLE) of the mean under H_1 , and the MLE of variance (which requires $K \geq 2$) under both hypotheses, we obtain the GLRT

$$L(z) = \exp\left(\frac{\|z\|^2}{\|z - \frac{\mathbf{1}^\top z}{K}\mathbf{1}\|^2} - \frac{\|z - \frac{\mathbf{1}^\top z}{K}\mathbf{1}\|^2}{\|z - \frac{\mathbf{1}^\top z}{K}\mathbf{1}\|^2}\right) \underset{H_0}{\overset{H_1}{\gtrless}} \tau \quad (31)$$

that reduces to the test

$$\delta_3(z) : \tilde{\lambda}_3 = \frac{\mathbf{1}^\top z}{\sqrt{K}\|z\|} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_3, \quad (32)$$

which is UMPI under positive scale transformations and symmetries w.r.t. the hyperplane orthogonal to $\mathbf{1}$ [26, Chapter 5]. To verify this claim, we use a variable transformation from [26, Equation (5.49)], and obtain the PDF of the NCWS $\tilde{\lambda}_3$ under H_j as

$$f_3(x|H_j) = C_j \sqrt{1-x^2}^{K-3} \exp\left(-\frac{jD^2}{2}\right) \cdot \int_0^\infty r^{K-1} \exp\left(-\frac{r^2}{2\nu^2}\right) \exp(jEr x) dr, \quad (33)$$

for $x \in [-1, 1]$, $j \in \{0, 1\}$,

where $\nu^2 = Pd^2 + 0.5\omega^2$, $D = Ad^2\sqrt{K}/\nu$, $E = D/\nu$, and C_j is a normalization constant under H_j . For other compensation schemes, we only need to replace A by the corresponding

gain parameter in D and E . In particular, under H_0 the PDF reduces to

$$f_3(x|H_0) = \frac{\Gamma(K/2)}{\sqrt{\pi}\Gamma((K-1)/2)} \sqrt{1-x^2}^{K-3} = \tilde{C}_0 \sqrt{1-x^2}^{K-3}, \text{ for } x \in [-1, 1], \quad (34)$$

which is independent of channel parameters and d , and allows us to obtain a threshold η_3 that guarantees a bounded probability of false alarm α , regardless of compensation scheme. Observe that $\tilde{C}_0 \neq C_0$, as it also contains the result of integrating over r with $j = 0$ in (33). Additionally, for $K = 2$ the above PDFs are defined for $x \in (-1, 1)$. In the following, we shall also refer to NCWS $\tilde{\lambda}_3$ as a *cosine* NCWS, as it corresponds to the cosine of the angle between the (normalized) mean vector $d^2\mathbf{1}$ and the sample vector z .

B. Known sensor mean m_0 and unknown m_1

Now, we consider m_1 unknown such that $m_1 > m_0$, and test $H_0 : m = m_0$ vs. $H_1 : m = m_1 > m_0$. Hence, the local LLRs have conditional distributions

$$\ell(v_i)|H_j \sim \mathcal{N}(0.5q_j d^2, d^2), \quad q_j = (-1)^{1-j}. \quad (35)$$

With m_1 unknown, d is unknown and it is not possible to add a bias $d^2/2$ on signals u_i as in (26). However, note that the distinguishing factor between conditional distributions (35) is only the sign of the mean, conveyed in the parameters q_j .

1) *Known channel parameters*: If A , P and ω^2 are known, the LRT at the server is straightforwardly computed as

$$L(z) = \exp\left(\frac{\|z + \frac{Ad^2}{2}\mathbf{1}\|^2}{2Pd^2 + \omega^2} - \frac{\|z - \frac{Ad^2}{2}\mathbf{1}\|^2}{2Pd^2 + \omega^2}\right) \underset{H_0}{\overset{H_1}{\gtrless}} \tau$$

$$\Rightarrow \delta_4(z) : \tilde{\lambda}_4 = \frac{\mathbf{1}^\top z}{K} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_4 = \frac{2Pd^2 + \omega^2}{2Ad^2K} \ln(\tau), \quad (36)$$

but, different to the test $\delta_1(z)$, here the distribution of $\tilde{\lambda}_4$ under H_0 depends on the unknown d , and we cannot compute a general threshold for bounded probability of false alarm α . Hence, there is no UMP test, and similarly, there are no UMP tests using AirCompFDM- M and - U .

However, we know that the hypothesis information is carried in the sign of the mean q_j . Hence, we now obtain an estimator of d^2 based on the ML estimator of the mean of the received samples, and replace it in the LRT formulation. Thanks to the knowledge of A , we have

$$\hat{d}^2 = \frac{2}{AK} |\mathbf{1}^\top z| \quad (37)$$

and the LRT is

$$L(z) = \exp\left(\frac{\|z + \frac{|\mathbf{1}^\top z|}{K}\mathbf{1}\|^2}{\frac{4P}{AK}|\mathbf{1}^\top z| + \omega^2} - \frac{\|z - \frac{|\mathbf{1}^\top z|}{K}\mathbf{1}\|^2}{\frac{4P}{AK}|\mathbf{1}^\top z| + \omega^2}\right) \underset{H_0}{\overset{H_1}{\gtrless}} \tau$$

$$\Rightarrow \ell(z) = \frac{\text{sgn}(\mathbf{1}^\top z)|\mathbf{1}^\top z|^2}{\frac{P}{A}|\mathbf{1}^\top z| + \frac{K}{4}\omega^2} \underset{H_0}{\overset{H_1}{\gtrless}} \ln(\tau). \quad (38)$$

Note that as $A, P, \omega^2 > 0$, the LLR is monotonically increasing in $\tilde{\lambda} = \mathbf{1}^\top z/K$, because the function

$$h(a) = \frac{a|a|}{b|a| + c} \quad (39)$$

TABLE I: Summary of proposed tests and its principal characteristics, using exact phase precompensation of AirCompFDM-P.

Test	Description	Unknown parameter	Conditional Distributions	UMP(I)?	Power
$\delta_1(\mathbf{z})$	Clairvoyant	All known	$\mathcal{N}(0.5q_j A d^2 K^{-1}, (Pd^2 + 0.5\omega^2)/K^{-2})$	UMP	(23)
$\delta_2(\mathbf{z})$	Average NCWS, biased LLRs	A	$\mathcal{N}(j A d^2 K^{-1}, (Pd^2 + 0.5\omega^2)/K^{-2})$	UMP	(30)
$\delta_3(\mathbf{z})$	Cosine NCWS, biased LLRs	$A \wedge (P \vee \omega^2)$	(33) and (34)	UMPI	Numerical
$\delta_4(\mathbf{z})$	Average NCWS	m_1	$\mathcal{N}(0.5q_j A d^2 K^{-1}, (Pd^2 + 0.5\omega^2)/K^{-2})$	No	-
$\delta_5(\mathbf{z})$	Cosine NCWS	$m_1 \wedge (A \vee P \vee \omega^2)$	(45)	No	-

is monotonically increasing $\forall a \neq 0$ when $b, c > 0$, and hence we can reduce (38) to

$$\tilde{\lambda} = \frac{\mathbf{1}^T \mathbf{z}}{K} \underset{H_0}{\overset{H_1}{\gtrless}} \eta, \quad (40)$$

i.e. we obtain the same test $\delta_4(\mathbf{z})$, and there are no UMP tests for unknown d and known channel parameters.

We note that the estimator of d^2 could be improved. However, this complicates the computation with no additional insight in test design. For example, computing the MLE of d^2 for the case where $\omega^2 = 0$ yields an involved expression, but the test still reduces to $\delta_4(\mathbf{z})$.

2) *Unknown channel parameters:* If A (corresp. B or $|G|$) is unknown, we cannot estimate d^2 using (37). As the means under H_0 and H_1 have equal magnitude, we can use the MLE of the mean magnitude $|\mathbf{1}^T \mathbf{z}|/K$. However, with d^2 unknown, we resort to using MLE of the variance under both hypotheses. Equivalently, we would use the same procedure when either P or ω^2 are unknown. Hence, the GLRT reduces to

$$\ell(\mathbf{z}) = \frac{4\text{sgn}(\mathbf{1}^T \mathbf{z})|\mathbf{1}^T \mathbf{z}|^2}{K\|\mathbf{z}\|^2 - |\mathbf{1}^T \mathbf{z}|^2} \underset{H_0}{\overset{H_1}{\gtrless}} \ln(\tau). \quad (41)$$

We can rewrite the LLR in (41) as a function of the cosine statistic $\tilde{\lambda}_5 = \mathbf{1}^T \mathbf{z}/(\sqrt{K}\|\mathbf{z}\|)$. A little algebra yields

$$\ell(\mathbf{z}) = \frac{4\text{sgn}(\tilde{\lambda}_5)\tilde{\lambda}_5^2}{1 - \tilde{\lambda}_5^2}, \quad (42)$$

which is monotonically increasing in $\tilde{\lambda}_5$ because the function

$$h(a) = \frac{4a|a|}{1 - a^2}, \quad a \in (-1, 1) \quad (43)$$

is monotonically increasing in its domain, where the statistic $\tilde{\lambda}_5$ resides. Hence, we can further reduce (41) to

$$\delta_5(\mathbf{z}) : \tilde{\lambda}_5 = \frac{\mathbf{1}^T \mathbf{z}}{\sqrt{K}\|\mathbf{z}\|} \underset{H_0}{\overset{H_1}{\gtrless}} \eta_5. \quad (44)$$

The PDF of the NCWS $\tilde{\lambda}_5$ under H_j can be obtained as [26, Chapter 5]

$$f_5(x|H_j) = c_j \sqrt{1 - x^2}^{K-3} \exp\left(-\frac{D^2}{8}\right) \cdot \int_0^\infty r^{K-1} \exp\left(-\frac{r^2}{2\nu^2}\right) \exp\left(\frac{q_j E r x}{2}\right) dr, \quad (45)$$

for $x \in [-1, 1]$, $j \in \{0, 1\}$,

where ν^2 , D , E are the same as in (33), c_j are the corresponding normalization constants under each hypothesis, and $q_j = (-1)^{1-j}$. Even when the PDF has no general closed form, it is clear then that under H_0 it depends at least on the unknown d , and hence the test $\delta_5(\mathbf{z})$ cannot be UMP(I).

Note that the cosine NCWS $\tilde{\lambda}_5$ is computed equivalently to the cosine NCWS $\tilde{\lambda}_3$, but it follows different distributions due to the unknown mean m_1 .

C. Summary

Table I organizes our proposed tests and summarizes their properties, referring to the particular equations for each case. We use logical operators “AND” $\wedge$ and “OR” $\vee$ to denote combinations of unknown parameters. All equations apply for exact phase precompensation of AirCompFDM-P. To obtain expressions for AirCompFDM-M and -U, the aggregated gain parameter A needs to be replaced with B and $|G|$, respectively.

V. PERFORMANCE BOUNDS

Under an NP framework, the relative entropy or Kullback-Leibler (KL) divergence governs the decay of probability of error of type I or II. Specifically, the probability of missed detection and false alarm will have exponential decay rate of

$$\text{KL}^I = \int f(x|H_0) \ln\left(\frac{f(x|H_0)}{f(x|H_1)}\right) dx, \quad (46)$$

$$\text{KL}^{II} = \int f(x|H_1) \ln\left(\frac{f(x|H_1)}{f(x|H_0)}\right) dx, \quad (47)$$

respectively. Note that the UMP and UMPI tests derived in the previous section are of type I, setting a bounded probability of false alarm, but can be reformulated into type II straightforwardly. In the following, we study the KL divergence of the tests derived in Section IV, according to the corresponding NCWS and its distributions under each hypothesis.

A. Average NCWS

For Gaussian observations with equal variance under both hypotheses, the local LLRs and the post-processed samples of the received signal are also Gaussian under both hypotheses, as shown in (18). Hence, the NCWS $\tilde{\lambda}_1$ also follows Gaussian distributions. In the particular case of exact phase precompensation, the hypothesis distance of the NCWS is

$$D_P = \frac{A d^2 \sqrt{K}}{\sqrt{P d^2 + 0.5 \omega^2}}, \quad (48)$$

and the corresponding KL divergence of test δ_1 is

$$\text{KL}_1^I = \text{KL}_1^{II} = \frac{D_P^2}{2} = \frac{A^2 d^4 K}{2 P d^2 + \omega^2}, \quad (49)$$

and the KL divergence for AirCompFDM-M and -U are obtained similarly. Additionally, note that the NCWS $\tilde{\lambda}_2$ and $\tilde{\lambda}_4$ also follow Gaussian distributions with the same hypothesis distance D , and therefore, the tests $\delta_2(\mathbf{z})$ and $\delta_4(\mathbf{z})$ have the same KL divergence (49).

Of course, D is a random variable that depends on channel realizations, channel noise and compensation scheme. To facilitate the analysis of the behavior of D with respect to the number of participating sensors S , we can approximate its mean value for moderate to large values of S as shown in [24] and obtain a direct relationship between D and S . In the case of phase pre-compensated channels,

$$\mathbb{E}\{D_P^2\} \approx \frac{\text{Var}\{a_i\} + S\mathbb{E}^2\{a_i\}}{\mathbb{E}\{a_i^2\} + 0.5\omega^2/(d^2S)} d^2 K, \quad (50)$$

and with no precompensation,

$$\mathbb{E}\{D_U^2\} \approx \frac{\text{Var}\{g_i\} + S(\mathbb{E}^2\{\text{Re}(g_i)\} + \mathbb{E}^2\{\text{Im}(g_i)\})}{\mathbb{E}\{|g_i|^2\} + 0.5\omega^2/(d^2S)} d^2 K. \quad (51)$$

In [24] we argued that these approximations are informative enough for designing detection strategies under a Bayesian formulation using the corresponding best error exponent, i.e. the Chernoff information $D^2/8$. For large S , the Chernoff information increases linearly with S , and thus the optimal strategy (in terms of detection performance) is to have as many participating sensors as possible. In particular, the Chernoff information of AirCompFDM-P is approximately affine with respect to S as dictated by (50) for *any* channel fading model because $\mathbb{E}^2\{a_i\} > 0$. On the other hand, AirCompFDM-U imposes noncoherent aggregation, and hence its Chernoff information related to (51) suffers with zero-mean channels such as Rayleigh fading, because $\mathbb{E}\{\text{Re}(g_i)\} = \mathbb{E}\{\text{Im}(g_i)\} = 0$ and there is no expected aggregation gains, besides the reduction of effective channel noise [24].

As the KL divergence is only a multiple of the Chernoff information for the average NCWS $\tilde{\lambda}_1$ as shown in (49), we expect the same collaboration gains as in the Bayesian case for each compensation scheme. Moreover, in this work, we also demonstrate the quality of these approximations with respect to S , in particular for low SNR scenarios.

B. Cosine NCWS

When both sensor means are known, the cosine NCWS $\tilde{\lambda}_3$ has conditional PDFs (33) under each hypothesis. Hence, the KL divergences for test $\delta_3(z)$ cannot be obtained in closed form, and furthermore, it is not straightforward to study the effect of S in the KL divergence. As an alternative analysis, we can approximate the KL divergences. We first replace D^2 by its approximate mean above, and furthermore, derive corresponding approximations of the mean of E and ν^2 . Then, the KL divergence will be a complex expression that explicitly depends on S . In the following, we derive the approximations using exact phase precompensation, and approximates for other compensation schemes can be obtained similarly. It is straightforward to see that

$$E_P = \frac{Ad^2\sqrt{K}}{Pd^2 + 0.5\omega^2} = \frac{A/S}{P/S + 0.5\omega^2/(Sd^2)}\sqrt{K} \quad (52)$$

Invoking the Strong Law of Large Numbers (SLLN) [25], $P/S \xrightarrow{a.s.} \mathbb{E}\{a^2\}$ as $S \rightarrow \infty$, and since ω^2 is constant,

$$\frac{P}{S} + 0.5\frac{\omega^2}{Sd^2} \xrightarrow{a.s.} \mathbb{E}\{a^2\} \text{ when } S \rightarrow \infty. \quad (53)$$

On the other hand, the SLLN also states that $A/S \xrightarrow{a.s.} \mathbb{E}\{a\}$ as $S \rightarrow \infty$. Hence, since the ratio of almost surely converging sequences converges almost surely to the ratio of their limits

$$E_P \xrightarrow{a.s.} \frac{\mathbb{E}\{a\}}{\mathbb{E}\{a^2\}}\sqrt{K} \text{ when } S \rightarrow \infty. \quad (54)$$

If we assume that the SLLN holds for moderate values of S , we obtain the approximation¹

$$E_P \approx \mathbb{E}\{E_P\} = \frac{\mathbb{E}\{a\}}{\mathbb{E}\{a^2\} + 0.5\omega^2/(Sd^2)}\sqrt{K}. \quad (55)$$

For the reciprocal of ν^2 , we observe that

$$\frac{r^2}{2\nu^2} = \frac{r^2/(Sd^2)}{2P/S + \omega^2/(Sd^2)} \quad (56)$$

and using the same procedure as above, where r^2 is a constant for the purposes of expectation,

$$\frac{r^2}{2\nu^2} \xrightarrow{a.s.} 0 \text{ when } S \rightarrow \infty. \quad (57)$$

Assuming again that the SLLN holds for moderate values of S , we obtain the approximation

$$\frac{r^2}{2\nu^2} \approx \mathbb{E}\left\{\frac{r^2}{2\nu^2}\right\} = \frac{r^2}{2\mathbb{E}\{a^2\}d^2S + \omega^2}. \quad (58)$$

Replacing in the PDFs and after some algebra, the approximate KL divergence of type I and type II for NCWS $\tilde{\lambda}_3$ become

$$\text{KL}_3^I \approx \frac{\mathbb{E}\{D^2\}}{2} - \ln\left(\frac{C_1}{C_0}\right) - \tilde{C}_0 \int_{-1}^1 \sqrt{1-x^2}^{K-3} \ln(R_3(x)) dx, \quad (59)$$

$$\text{KL}_3^{II} \approx \ln\left(\frac{C_1}{C_0}\right) - \frac{\mathbb{E}\{D^2\}}{2} + C_1 \int_{-1}^1 \sqrt{1-x^2}^{K-3} R_3(x) \ln(R_3(x)) dx, \quad (60)$$

respectively, where

$$R_3(x) = \int_0^\infty r^{K-1} \exp\left(-\mathbb{E}\left\{\frac{r^2}{2\nu^2}\right\} + \frac{\mathbb{E}\{E\}rx}{2}\right) dr. \quad (61)$$

Even with the above approximations, it is not clear how the KL divergence depends on S . Hence, we obtain numerical computations of the approximate KL divergence and compare them with the average KL divergence of empirical channel realizations. Nevertheless, similarly to the case of average NCWS, we expect that KL divergence increases with S for AirCompFDM-P, and that it stagnates for AirCompFDM-U.

Finally, when m_1 is unknown, the cosine NCWS $\tilde{\lambda}_5$ has conditional PDFs $f_5(x|H_j)$ as shown in (45). Again, the KL divergences cannot be obtained in closed form, even for given d^2 , A , P and ω^2 , and moreover, the PDF under the null hypothesis depends on all these parameters as well. However,

¹We note that for AirCompFDM-U, $|G|$ is the magnitude of a complex number. Thus, instead of invoking SLLN, we would need to invoke the Central Limit Theorem and the Continuous Mapping Theorem [25] to obtain an approximation, akin to our derivation in [24].




Fig. 5: Empirical ROC of test $\delta_2(\mathbf{z})$ for different number of sensors S and SNR values, with $d = K = N = 1$. We average 1000 realizations over 1000 independent sensor deployments.





Fig. 6: KL divergence of type I for test $\delta_2(\mathbf{z})$ with respect to the number of sensors S , for different values of channel SNR, and $d = N = K = 1$. We average 1000 independent sensor deployments for empirical results, and use approximations (50) and (51) for AirCompFDM-P and -U, respectively.

a simple variable transformation shows that $f_5(x|H_0) = f_5(-x|H_1)$ for $x \in (-1, 1)$, and hence $c_0 = c_1 = c$. Moreover, KL divergences of type I and II are equal as well. Therefore, using the approximate means of D^2 , E and r^2/ν^2 , the KL divergences of type I and II for NCWS $\tilde{\lambda}_5$ are approximately

$$\text{KL}_5^I = \text{KL}_5^{II} \approx c \exp\left(-\frac{\mathbb{E}\{D^2\}}{8}\right) \cdot \int_{-1}^1 \sqrt{1-x^2}^{K-3} R_5(-x) \ln\left(\frac{R_5(-x)}{R_5(x)}\right) dx, \quad (62)$$

respectively, where

$$R_5(x) = \int_0^\infty r^{K-1} \exp\left(-\mathbb{E}\left\{\frac{r^2}{2\nu^2}\right\} + \frac{\mathbb{E}\{E\}rx}{2}\right) dr. \quad (63)$$

VI. NUMERICAL EXPERIMENTS

To illustrate the performance gains of AirCompFDM, we set up several different network settings. Unless otherwise stated, we simulate S sensors and one server, all equipped

with a single antenna. We assume each sensor obtains $N = 1$ observation, independent among sensors, and contaminated with i.i.d. Gaussian measurement noise, parameterized by the sensor hypothesis distance $d = |m_1 - m_0|/\sigma$. Without loss of generality, we assume antipodal means, i.e. $m_0 = -m_1 < 0$, and hence for each value of d , we fix $\sigma = 1$ and obtain $m_1 = -m_0 = 0.5\sigma/d$. We simulate i.i.d. Rayleigh channels, i.e. $g_i \sim \mathcal{N}(0, 1/2) + j\mathcal{N}(0, 1/2)$. Channel noise is AWGN with intensity ω^2 corresponding to a given average SNR for a single sensor. The server makes a decision using K samples of the received analog signal, and we focus on the lowest amount of samples possible for any given test, i.e. $K = 1$ or $K = 2$. As stated in Section III, we generate 1000 independent sensor deployments, i.e., static channel realizations for K samples and S sensors, and perform 1000 independent Monte Carlo simulations with different sensor observations, allowing us to faithfully reproduce average behavior of our tests.

A. Average NCWS

Fig. 5 shows the ROC for test $\delta_2(\mathbf{z})$ for sensors sending biased LLRs as in (26), for different compensation schemes, and $d = N = K = 1$. Note that this performance is equal to




Fig. 7: Empirical ROC of test $\delta_3(\mathbf{z})$ for different number of sensors S and SNR values, with $K = 2$ and $d = N = 1$. We average 1000 realizations over 1000 independent sensor deployments.





Fig. 8: KL divergence of type I for test $\delta_3(\mathbf{z})$ with respect to the number of sensors S , for different values of channel SNR, $d = N = 1$ and $K = 2$. We average 1000 independent sensor deployments for empirical results, and use (59) with corresponding approximate parameters for AirCompFDM-P or -U.

the one of the clairvoyant test $\delta_1(\mathbf{z})$, as the hypothesis distance D is equivalent for equal parameter realizations.

Fig. 6 depicts the KL divergence of test $\delta_2(\mathbf{z})$, obtained by averaging empirical simulations and compares with our proposed approximation, for Rayleigh channels and $d = N = K = 1$. For AirCompFDM-P, collaboration gains are significant with increasing S . Moreover, our approximation is sharp and only slightly underestimates the average empirical KL divergence in all cases. On the other hand, AirCompFDM-U does not exhibit such collaboration gains. Regardless, our approximation is close to the empirical average, presenting a small overestimation for larger values of S , even accounting for the larger variability of the compensation of noncoherent channel aggregation. Nevertheless, our approximation helps to devise detection strategies. Moreover, our results indicate that detection performance is maximized when *all* active sensors collaborate, regardless of channel SNR.

Of course, the system designer can always improve performance by increasing transmission power, but as we can see, the KL divergence only increases by about 6% in AirCompFDM-P and 7.7% in the worst case of AirCompFDM-U. The alternative is to increase d by increasing the number of local

observations N , which has no impact on resource usage and requires marginal additional energy consumption. Resource usage does not increase because all observations are consolidated in a single LLR u_i , i.e., using a single transmission slot over the air. Naturally, collecting more data consumes some energy. However, energy used in sensing is in practice significantly lower than the energy use in communication signaling by a typical IoT node.

B. Cosine NCWS

Fig. 7 depicts the empirical ROCs of test $\delta_3(\mathbf{z})$, using cosine NCWS $\tilde{\lambda}_3$ with known sensor means, under different compensation schemes, $K = 2$ and $d = N = 1$. In all compensation schemes, we can see that sensor aggregation improves performance significantly even for a modest number of sensors. In particular, phase precompensation outperforms AirCompFDM-U in terms of probability of detection for lower values of probabilities of false alarm.

Fig. 8 shows the average empirical and approximate KL divergence of type I of test $\delta_3(\mathbf{z})$ with respect to the number of users S , under equal simulation parameters. For AirCompFDM-P, our approximation barely underestimates the




Fig. 9: Empirical ROC of test $\delta_5(\mathbf{z})$ for different number of sensors S and SNR values, with $K = 2$ and $d = N = 1$. We average 1000 realizations over 1000 independent sensor deployments.





Fig. 10: KL divergence of type I for test $\delta_5(\mathbf{z})$ with respect to the number of sensors S , for different values of channel SNR, $d = N = 1$ and $K = 2$. We average 1000 independent sensor deployments for empirical results, and use (62) with corresponding approximate parameters for AirCompFDM-P or -U.

average empirical KL divergence for all SNR values, providing a sharp lower bound for performance gains. Thus, using these approximations as design criteria would lead, on average, to slightly higher performance than expected. Regardless, it is clear that the KL divergence increases with S , showing the gains of over-the-air aggregation. Moreover, increasing SNR improves the KL divergence by about 6.2%, which further demonstrates the gains obtained by sensor collaboration. In the case of AirCompFDM-U, there is a larger variance in empirical results. However, our approximation is again close, with a small overestimation over the empirical KL divergence with increasing S . In particular, the worst overestimation is about 7.5% over the empirical KL divergence. Nevertheless, our approximation behaves similarly to the empirical KL divergence for increasing S , and hence it is helpful when designing a decision-making strategy or choosing deployment parameters.

Fig. 9 shows the empirical ROCs of test $\delta_5(\mathbf{z})$, where m_1 is unknown and we use the cosine NCWS λ_5 , for all compensation schemes, with $K = 2$ and $d = N = 1$. Although this test shares similarities with $\delta_3(\mathbf{z})$, they shall not be confused as the conditional PDFs under each hypothesis

are different. Hence, given that in this case the samples have antipodal means under each hypothesis instead of zero vs. non-zero mean as in $\delta_3(\mathbf{z})$, performance is better compared to the ROCs of test $\delta_3(\mathbf{z})$ across compensation schemes and channel realizations. Conversely, $\delta_3(\mathbf{z})$ is an UMPI test, whereas $\delta_5(\mathbf{z})$ is not. Again, phase precompensation provides considerable performance gains, and are much greater than those of post-compensation.

Fig. 10 depicts the KL divergence of type I for test $\delta_5(\mathbf{z})$ with respect to the number of sensors S , for different SNR values, $K = 2$ and $d = N = 1$. Our approximate KL divergence with AirCompFDM-P is, again, very close to the empirical average for all S , slightly underestimating the empirical KL divergence. KL divergences for high SNR values are only 4.9% larger than the KL divergence with SNR of -10 dB, on average. Using AirCompFDM-U, our approximation is close in all cases. It is particularly sharp for $S \leq 18$ and low SNR, and it overestimates the empirical KL divergence by less than 10% for $S > 20$ for all SNR values. Again, channel noise is mitigated with increasing S , and higher SNR leads to a lower number of participating sensors required to achieve maximum detection performance.




Fig. 11: ROC of legitimate server and eavesdropper for different number of sensors S , with channel SNR -10dB , and $d = N = K = 1$. We average 1000 realizations over 1000 independent sensor deployments.

C. Privacy against eavesdroppers

Our results above show that phase pre-compensation (either perfect or quantized) yields significant performance gains across system parameters, compared with the aggregated phase post-compensation of AirCompFDM-U. We want to emphasize that, from a practical perspective, this performance gap also corresponds to a privacy gap when a third-party node attempts to perform the same decision.

Consider a network using phase pre-compensation, and an eavesdropper node that aims to emulate the decision-making of the legitimate server. The sensors have precompensated the phase of their channels g_i to the legitimate server, which are different than those from sensors to eavesdropper g_i^e . Hence, the best that the eavesdropper can hope to achieve (without the help from sensor precoding) is exact phase post-compensation of its received signal

$$y^e = \sum_{i \in \mathcal{S}} g_i^e \frac{\bar{g}_i}{|g_i|} u_i(k) + n = \sum_{i \in \mathcal{S}} \tilde{g}_i^e u_i(k) + n, \quad (64)$$

with performance equal to that of AirCompFDM-U over channels $\tilde{g}_i^e$.

Fig. 11 compares the performance of a legitimate server using AirCompFDM-P and AirCompFDM-M to the performance of an eavesdropper using AirCompFDM-U, for varying number of sensors, channel SNR of -10dB , and $d = N = K = 1$. For $S = 1$, the performance is the same in all compensation methods, as pre- and post-compensation is equivalent. Nevertheless, as S increases, detection performance of the legitimate server using phase pre-compensation improves at a much faster rate than the performance of the eavesdropper using its best-case strategy. This is particularly interesting for lower values of P_F , which correspond to practical test designs, where the gap is quite significant.

VII. CONCLUSIONS

In this work, we proposed a resource-efficient framework for collaborative decision-making in sensor networks using over-

the-air aggregation of soft information. Thanks to over-the-air computation, we exploit the natural mixing of signals in wireless systems to reduce network coordination and resource usage. We propose different pre- and post-compensation to account for the detrimental effect of channel phases, and design hypothesis testing under these conditions

By exploiting the natural mixture of signals in the multiple-access channel, sensors share locally computed log-likelihood ratios and the server observes the resulting aggregated signal. With careful design of pre- and post-processing, the received signal corresponds to a good approximation of the log-likelihood ratio computed over all observed data under ideal communication conditions, as shown in Section III.

Furthermore, in Section IV we designed composite tests for over-the-air collaborative detection, depending on the knowledge of relevant system and channel parameters. Our results show significant performance increase due to over-the-air aggregation with simple protocol and minimal resource requirements and coordination.

Our proposed framework has important applicability in low-cost low-power wireless sensor networks. We will investigate, in future work, over-the-air decision-making schemes that consider self-censoring devices to further reduce energy consumption, while still providing satisfactory detection performance.

REFERENCES

- [1] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, "Internet of Things: A survey on enabling technologies, protocols, and applications," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 4, pp. 2347–2376, Fourthquarter 2015.
- [2] J.-F. Chamberland and V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 407–416, 2003.
- [3] A. Tarighati and J. Jaldén, "Rate allocation for decentralized detection in wireless sensor networks," in *IEEE 16th Int. Workshop Signal Process. Advances Wireless Commun. (SPAWC)*, 2015, pp. 341–345.
- [4] S. Kar, R. Tandon, H. V. Poor, and S. Cui, "Distributed detection in noisy sensor networks," in *2011 IEEE Int. Symp. Inf. Theory Proc.*, 2011, pp. 2856–2860.
- [5] B. Chen, R. Jiang, T. Kasetkasem, and P. Varshney, "Channel aware decision fusion in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 52, no. 12, pp. 3454–3458, 2004.
- [6] R. Niu, B. Chen, and P. Varshney, "Fusion of decisions transmitted over Rayleigh fading channels in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 54, no. 3, pp. 1018–1027, 2006.
- [7] S. Li, X. Li, X. Wang, and J. Liu, "Decentralized sequential composite hypothesis test based on one-bit communication," *IEEE Trans. Inf. Theory*, vol. 63, no. 6, pp. 3405–3424, 2017.
- [8] J.-F. Chamberland and V. Veeravalli, "How dense should a sensor network be for detection with correlated observations?" *IEEE Trans. Inf. Theory*, vol. 52, no. 11, pp. 5099–5106, 2006.
- [9] S. Appadwedula, V. V. Veeravalli, and D. L. Jones, "Decentralized detection with censoring sensors," *IEEE Trans. Signal Process.*, vol. 56, no. 4, pp. 1362–1373, 2008.
- [10] X. Wang, G. Li, and P. K. Varshney, "Detection of sparse signals in sensor networks via locally most powerful tests," *IEEE Signal Process. Lett.*, vol. 25, no. 9, pp. 1418–1422, 2018.
- [11] X. Wang, G. Li, C. Quan, and P. K. Varshney, "Distributed detection of sparse stochastic signals with quantized measurements: The generalized Gaussian case," *IEEE Trans. Signal Process.*, vol. 67, no. 18, pp. 4886–4898, 2019.
- [12] B. Chen and P. Willett, "On the optimality of the likelihood-ratio test for local sensor decision rules in the presence of nonideal channels," *IEEE Trans. Inf. Theory*, vol. 51, no. 2, pp. 693–699, 2005.
- [13] A. Azari, P. Popovski, G. Miao, and C. Stefanovic, "Grant-free radio access for short-packet communications over 5G networks," in *2017 IEEE Global Commun. Conf. (GLOBECOM)*, Dec 2017, pp. 1–7.

- [14] J. Dong and Y. Shi, "Nonconvex demixing from bilinear measurements," *IEEE Trans. Signal Process.*, vol. 66, no. 19, pp. 5152–5166, Oct 2018.
- [15] C. Feres and Z. Ding, "Wirtinger Flow meets Constant Modulus Algorithm: Revisiting signal recovery for grant-free access," *IEEE Trans. Signal Process.*, vol. 69, pp. 6515–6529, 2021.
- [16] C. Feres and Z. Ding, "A Riemannian geometric approach to blind signal recovery for grant-free radio network access," *IEEE Trans. Signal Process.*, vol. 70, pp. 1734–1748, 2022.
- [17] M. Goldenbaum, H. Boche, and S. Stańczak, "Harnessing interference for analog function computation in wireless sensor networks," *IEEE Trans. Signal Process.*, vol. 61, no. 20, pp. 4893–4906, 2013.
- [18] O. Abari, H. Rahul, and D. Katabi, "Over-the-air function computation in sensor networks," 2016, [Online]. Available: arXiv:1612.02307 [cs.NI].
- [19] L. Chen, N. Zhao, Y. Chen, F. R. Yu, and G. Wei, "Over-the-air computation for cooperative wideband spectrum sensing and performance analysis," *IEEE Trans. Veh. Technol.*, vol. 67, no. 11, pp. 10 603–10 614, 2018.
- [20] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning based on over-the-air computation," in *2019 IEEE Int. Conf. Commun. (ICC)*, 2019, pp. 1–6.
- [21] H. Ma, X. Yuan, D. Fan, Z. Ding, X. Wang, and J. Fang, "Over-the-air federated multi-task learning," in *2022 IEEE Int. Conf. Commun. (ICC)*, 2022, pp. 5184–5189.
- [22] S. Marano, V. Matta, L. Tong, and P. Willett, "A likelihood-based multiple access for estimation in sensor networks," *IEEE Trans. Signal Process.*, vol. 55, no. 11, pp. 5155–5166, 2007.
- [23] C. Feres, B. C. Levy, and Z. Ding, "Over-the-air collaborative learning in joint decision making," in *2022 IEEE Global Commun. Conf. (GLOBECOM)*, 2022, pp. 3581–3586.
- [24] C. Feres, B. C. Levy, and Z. Ding, "Bayesian decision making via over-the-air soft information aggregation," in *2023 IEEE Int. Conf. Commun. (ICC)*, 2023, to appear in.
- [25] M. Loève, *Probability Theory I*, 4th ed. Springer, New York, NY, 1977.
- [26] B. C. Levy, *Principles of Signal Detection and Parameter Estimation*, 1st ed. Springer, Boston, MA, 2007.
- [27] J.-F. Chamberland and V. V. Veeravalli, "Wireless sensors in distributed detection applications," *IEEE Signal Process. Mag.*, vol. 24, no. 3, pp. 16–25, 2007.
- [28] R. J. Marks, G. L. Wise, D. G. Haldeman, and J. L. Whited, "Detection in Laplace noise," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-14, no. 6, pp. 866–872, 1978.



Bernard C. Levy (Life Fellow, IEEE) received the diploma of Ingénieur Civil des Mines from the Ecole Nationale Supérieure des Mines Paris, Paris, France, in 1974, and the Ph.D. in Electrical Engineering from Stanford University, Stanford, CA, USA, in 1979. From July 1979 to June 1987 he was Assistant and then Associate Professor in the Department of Electrical Engineering and Computer Science at M.I.T. Since July 1987, he has been with the University of California at Davis, Davis, CA, USA, where he was a Professor of Electrical Engineering and a member of the Graduate Group in Applied Mathematics until July 2020 and is currently an Emeritus Professor. He was the Chair of the Department of Electrical and Computer Engineering with UC Davis from 1996 to 2000. He was a Visiting Scientist at the Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA) in Rennes, France from January to July 1993, and at the Institut National de Recherche en Informatique et Automatique (INRIA), in Rocquencourt, France, from September to December 2001. He is the author of the books *Principles of Signal Detection and Parameter Estimation* (New York, NY: Springer, 2008) and *Random Processes with Applications to Circuits and Communications* (Cham, Switzerland: Springer Nature, 2019). His research interests include statistical signal processing, estimation, detection, and circuits applications of signal processing. Prof. Levy served as Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS I, of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS II, of the *EURASIP J. on Advances in Signal Processing*, and of *Signal Processing*. He is a member of SIAM.



Zhi Ding (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, USA, in 1990. He is currently with the Department of Electrical and Computer Engineering, University of California at Davis (UC Davis), Davis, CA, USA, where he holds the position of Distinguished Professor. From 1990 to 2000, he was a Faculty Member of Auburn University and later with The University of Iowa, Iowa City, IA, USA. In 2000, he joined the College of Engineering with UC Davis. He is a coauthor of the textbook: *Modern Digital and Analog Communication Systems*, 5th edition, Oxford University Press, 2019. His research interests and expertise cover the areas of wireless networking, communications, signal processing, multimedia, and learning.

Prof. Ding supervised more than 30 Ph.D. dissertations since joining UC Davis. His research team of enthusiastic researchers works very closely with industry to solve practical problems and contributes to technological advances. His team has collaborated with researchers around the world and welcomes self-motivated young talents as new members. Prof. Ding is currently the Chief Information Officer and Chief Marketing Officer of the IEEE Communications Society. He was an Associate Editor for IEEE TRANSACTIONS ON SIGNAL PROCESSING from 1994 to 1997, 2001 to 2004, and an Associate Editor of IEEE SIGNAL PROCESSING LETTERS 2002–2005. He was a Member of the Technical Committee on Statistical Signal and Array Processing and a Member of the Technical Committee on Signal Processing for Communications (1994–2003). Prof. Ding was the General Chair of the 2016 IEEE International Conference on Acoustics, Speech, and Signal Processing and the Technical Program Chair of the 2006 IEEE Globecom. He was also an IEEE *Distinguished Lecturer* (Circuits and Systems Society, 2004–2006, Communications Society, 2008–2009). He served as IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS Steering Committee Member (2007–2009) and its Chair (2009–2010). Prof. Ding was the recipient of the IEEE Communication Society's WTC Award in 2012 and the IEEE Communication Society's Education Award in 2020.



Carlos Feres (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering from Pontificia Universidad Católica de Chile, Santiago, Chile, in 2011 and 2013, respectively, and the Ph.D. degree in electrical engineering from the University of California at Davis, Davis, CA, USA, in 2021. He is currently a Senior Engineer in Modem R&D with Samsung Semiconductor Inc., San Diego, CA, USA. His research interests include blind estimation, cross-layer design and optimization, MIMO wireless networks, wireless sensor networks, non-convex op-

timization, and collaborative and federated learning.