Algorithms for Distributed Chirp Assignment in MIMO-FMCW Systems

***Abstract*—We investigate distributed chirp assignment for a network of multiple-input multiple-output frequency-modulated continuous-wave (MIMO–FMCW) radars with limited measurements. A fixed number of chirps is assigned to each radar from a common pool, subject to interference between radars with overlapping fields of view. To accommodate intermittent communication, we formulate a graph-based collision objective incorporating an age-of-information (AoI)-dependent penalty for unavailable assignments. We propose four algorithms: distributed row Gibbs sampling (DRGS), parallel Gibbs sampling using auxiliary Gaussian variables, and block coordinate descent (BCD) with single-radar and two-radar blocks. We establish finite-iteration bounds on the expected objective gap, sufficient conditions for geometric convergence to the Gibbs distribution, and stability bounds under perturbations of the AoI-dependent penalty. For an augmented angle–Doppler dictionary, we derive sufficient conditions for sparse recovery in terms of physical chirp overlap and the coherence of the constituent dictionaries. We further develop a target-mediated cross-echo model and relate the expected interference power to the collision objective. In simulations with two radars, six targets in a common range cell and an input SNR of -15 dB, row BCD achieves a detection rate of 93.0% and a false-discovery rate of 4.1%, compared with 42.4% and 9.3%, respectively, for independent random chirp assignment.**

***Index Terms*—MIMO radar, FMCW radar, interference mitigation, chirp assignment, Gibbs sampling, block coordinate descent, sparse recovery.**

I. Introduction

MULTIPLE-INPUT multiple-output frequency-modulated continuous-wave (MIMO–FMCW) radar exploits spatial and waveform diversity for multitarget localization. High-resolution estimation in the angular and Doppler domains typically requires a large virtual aperture and numerous slow-time measurements, with associated hardware and computational costs [1], [2]. Measurement selection for MIMO radar has consequently received considerable attention. Tohidi et al. [2] consider antenna and pulse placement using the Cramér–Rao lower bound as a design criterion; Rai et al. [1] develop compressive-sensing (CS) techniques for range estimation and joint Doppler–angle estimation from a reduced set of chirps and antenna measurements. These results motivate the use of a fixed measurement budget for each radar.

In a multiradar system, however, measurement selection at an individual radar cannot in general be considered independently of the transmissions of the other radars. Simultaneous transmissions in overlapping sectors may give rise to mutual interference. Depending on the transmitted waveforms, propagation paths, and receiver processing, such interference can raise the noise floor or produce spurious target responses [3]. Orthogonal allocation of the transmission resources avoids simultaneous use but can impose an unnecessarily large resource requirement when interference is confined to a subset of radar pairs. Spatial reuse is consequently of interest, particularly when each radar is required to acquire a prescribed number of measurements within a common chirp pool.

A. Related work

Interference avoidance and interference suppression constitute two complementary approaches to the coexistence problem. RadChat [4] exploits inter-radar communication to coordinate access to the radar spectrum; its synchronization-free extension [5] uses relative timing information. Graph-coloring formulations have also been investigated for dynamic radar resource allocation [6]. Pan et al. [7] formulate joint frequency–time interference avoidance as a repeated anti-coordination game, with mixed strategies updated by regret minimization. In contrast to these formulations, our decision variables are fixed-cardinality subsets of a common set of chirp times. We explicitly allow a physical conflict to exist in the absence of a timely bidirectional communication link.

Interference can also be mitigated through waveform design or processing of the received signal. BlueFMCW [8] employs randomized frequency hopping, whereas [9] exploits sparsity in a chirplet dictionary. Wang et al. [10] formulate interference suppression as the decomposition of a Hankel matrix into sparse and low-rank components, exploiting differences between the dechirped interference and desired beat signals. The present work is complementary to these approaches: it concerns the selection of the measurements themselves. In particular, we extend the limited-measurement sensing formulation of [1] to a network in which the measurement subsets selected at different radars determine the support of the interfering returns.

The optimization problem is closely related to inference in graphical models [11]: the assignment of a radar is a discrete variable with a fixed-cardinality state space, and neighbouring assignments interact through pairwise potentials. Its single-site conditional distribution is a conditional Bernoulli distribution [12], [13]. The auxiliary-variable approach to parallel sampling of pairwise Markov random fields [14], [15] motivates our Gaussian construction. Distributed Glauber dynamics and asynchronous Gibbs sampling have been studied in [16], [17]; these results do not imply invariance of a kernel that substitutes stale rows for its current conditional arguments. We instead incorporate stale information into the objective itself using the AoI framework [18], [19].

B. Contributions

The contributions of this work are summarized as follows.

1) We formulate the fixed-quota chirp assignment problem using separate physical-conflict and timely-communication graphs. Unavailable neighbouring assignments are represented by a quota-preserving prediction model whose weights depend on the age of the most recently received assignment.
2) We develop four distributed scheduling algorithms. The single-radar Gibbs and two-radar BCD subproblems admit exact dynamic-programming implementations. An edge-incidence representation of the interaction term permits Gaussian augmentation and simultaneous sampling of the radar assignments without altering the marginal Gibbs distribution.
3) We establish additive objective-gap bounds for single-radar and paired BCD, a matching criterion for selection of pairs, and finite-time bounds for both Gibbs samplers. We also characterize the sensitivity of the objective gap and the Gibbs distribution to changes in the AoI-dependent costs.
4) We derive an overlap-dependent bound on the coherence between target and interference dictionaries, including the case of leakage outside the shared chirps. In conjunction with the exact recovery condition and results on sparsely corrupted signals [20]–[22], this gives sufficient conditions for recovery from the selected measurements.
5) We develop a range-cell cross-echo model for target-mediated interference, establish a phase-averaged interference-power identity, and evaluate the relationship between collision energy, SINR, detection rate, and false discoveries using Monte Carlo simulations.

The remainder of the paper is organized as follows. Section II describes the system model and scheduling protocol. Sections III and IV present the algorithms and their analysis, respectively. Section V establishes the connection to sparse recovery, Section VI develops the cross-echo model, and Section VII presents the numerical results.

TABLE I
FREQUENTLY USED NOTATION. VECTORS ARE BOLD LOWERCASE; MATRICES ARE BOLD UPPERCASE.

Symbol	Meaning
N, M, P_i	Radars, available chirps, fixed quota of radar i
$\mathcal{S}_i, \boldsymbol{x}_i, \mathbf{X}$	Selected chirps, binary assignment vector, all assignments
$\boldsymbol{r}_i, \theta_i, 2\omega_i$	Radar location, boresight, sector opening angle
$\mathbf{A}^{\mathrm{phys}}, \mathbf{H}, \mathbf{A}$	Physical conflicts, timely communication, available conflicts
$n, t, a_{ij}(n)$	Sensing interval, scheduling iteration, age of a stored assignment
$\mathbf{G}, \boldsymbol{g}_i$	Neighbour chirp counts $\mathbf{AX}$; transpose of its row i
$\pi_{ijp}, \mathbf{F}, \lambda$	Predicted chirp use, resulting cost matrix, and its weight
$E, E_{\mathrm{col}}^{\mathrm{phys}}, \mathcal{G}$	Scheduling cost, physical collision energy, gap $E - E_{\min}$
$\mathcal{B}, D_{\mathcal{B}}$	Radars allowed to change together; maximum available decrease
q_i	$\min(P_i, M - P_i)$; maximum chirps replaced in a row change
$\mathcal{P}, L_{\mathcal{P}}$	Partition into such sets; residual error bound
β, π_β	Inverse temperature and Gibbs distribution
$C_{ij}, \mathbf{C}$	Bounds on Dobrushin influence; matrix $[C_{ij}]$
$\mathbf{B}, \mathbf{L}, \mathbf{Z}$	Edge-incidence matrix, Laplacian, Gaussian edge variables
$\mathbf{D}_i, \mathcal{I}$	Normalized augmented dictionary; column support
$\mathbf{1}, \mathbb{1}\{\cdot\}$	All-ones vector and indicator function

II. SYSTEM MODEL

We consider a multiradar multisector system as depicted in Figure 1. Each radar chooses chirps for one active sector. The sector directions stay fixed while the assignments are computed and used; a change of sectors gives a new assignment problem. We first describe what each radar selects and how shared chirps contribute to collisions. We then account for neighbours whose current assignments are unavailable. We assume perfect calibration, frozen sectors within CPIs and the availability of inter-radar communication functionalities.

A. Chirp Selection and Assignment Matrix

Consider a chirp pool $\mathcal{M} = \{1, \dots, M\}$ shared by N radars. Each radar i selects a subset $\mathcal{S}_i \subseteq \mathcal{M}$ containing exactly P_i chirps, that is, $|\mathcal{S}_i| = P_i$, with $1 \le P_i \le M$. Each index in $\mathcal{M}$ corresponds to a chirp start time.

We can represent the choices of each radar by recording its chirp index selections as a binary entry:

$$x_{ip} = \mathbb{1}\{p \in \mathcal{S}_i\}, \qquad \boldsymbol{x}_i = [x_{i1}, \dots, x_{iM}]^T.$$

The vector $\boldsymbol{x}_i$ has exactly P_i ones. Stacking the vectors as rows gives the assignment matrix

$$\mathbf{X} = \begin{bmatrix} \boldsymbol{x}_1^T \\ \vdots \\ \boldsymbol{x}_N^T \end{bmatrix} \in \{0,1\}^{N \times M}.$$

Thus, this matrix describes every radar's choice: a row identifies a radar and a column identifies a particular chirp. It is a mathematical description of the chirp assignments of the entire multiradar network. The allowed rows and schedules are

$$\mathcal{X}_i = \{\boldsymbol{x} \in \{0,1\}^M : \mathbf{1}^T \boldsymbol{x} = P_i\}, \qquad \mathcal{X} = \prod_i \mathcal{X}_i. \tag{1}$$

Here $\mathbf{1}$ is the all-ones vector. We force all our proposed scheduling algorithms to select rows in $\mathcal{X}_i$, so the fixed chirp quotas P_i always hold.

B. Collision Energy

Which pairs should avoid sharing a chirp? To determine this, we consider radar i to be at position $\boldsymbol{r}_i$, with boresight θ_i and half-opening angle $0 < \omega_i < \pi/2$. We define its active sector to be:

$$\mathcal{W}_i = \{\boldsymbol{r}_i + r\big[\cos\phi, \sin\phi\big]^T : r \ge 0,\ |\phi - \theta_i| \le \omega_i\},$$

where angular differences are interpreted modulo 2π. Two radars have a potential conflict when the interiors of their sectors share a region. For simplicity, we assume infinite radial coverage for each radar, and hence focus on angular overlap.

A conflict graph records these pairs so that scheduling need not repeat the geometric calculation. Each radar is modelled as a node. An undirected edge $\{i, j\} \in \mathcal{E}^{\mathrm{phys}}$ denotes a potential conflict. Its binary adjacency matrix has entries

$$A_{ij}^{\mathrm{phys}} = \mathbb{1}\{\operatorname{int}\mathcal{W}_i \cap \operatorname{int}\mathcal{W}_j \neq \varnothing\}, \quad i \neq j, \qquad A_{ii}^{\mathrm{phys}} = 0.$$

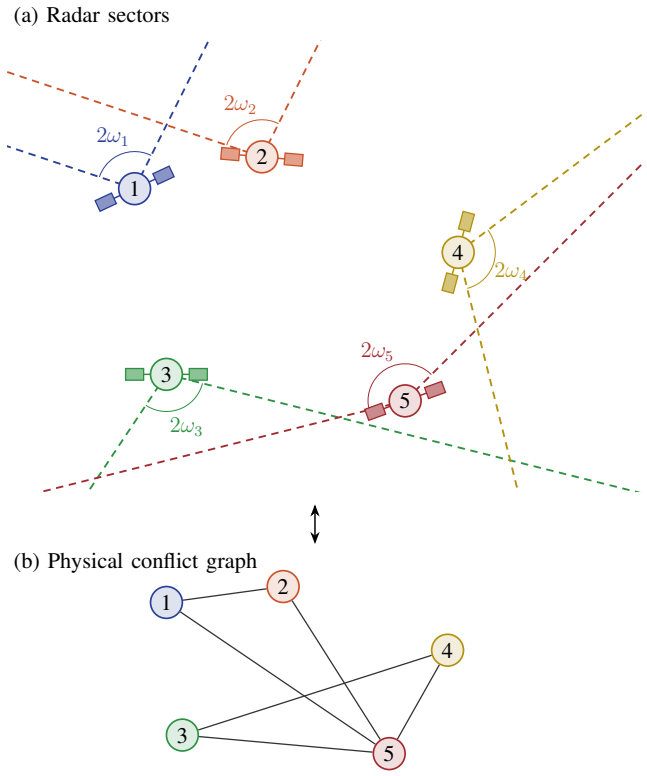

Fig. 1. Five-radar system used throughout. Each pair of dashed rays bounds an unbounded angular sector; the arc specifies which sector is active. An edge denotes a shared region. The physical edges are $12, 15, 25, 34, 35, 45$. The page boundary truncates the drawing only. Crossings of graph edges are not radar nodes.

A conflicting pair shares $|\mathcal{S}_i \cap \mathcal{S}_j|$ chirps, which is the number of intersecting chirps between radars i and j. Adding this count over $\mathcal{E}^{\text{phys}}$ defines the physical collision energy:

$$E_{\text{col}}^{\text{phys}}(\mathbf{X}) = \sum_{\{i,j\}\in\mathcal{E}^{\text{phys}}} |\mathcal{S}_i \cap \mathcal{S}_j|. \tag{2}$$

However, it should be noted that this is a simplified proxy for interference. Received interference power also depends on propagation, timing, and receiver processing [3].

We can represent the collision energy using matrix notation. Observe that chirp p is a member of the set $\mathcal{S}_i \cap \mathcal{S}_j$ iff it is a member of both $\mathcal{S}_i$ and $\mathcal{S}_j$, or equivalently, if both $x_{ip} = 1$ and $x_{jp} = 1$ are true. Also, radars i and j have an edge iff $A_{ij}^{\text{phys}} = 1$. Therefore, to obtain the collision energy, we multiply these three terms and sum over all i, j and p. Summing over all terms counts each undirected pair twice, giving

$$\begin{aligned} E_{\text{col}}^{\text{phys}}(\mathbf{X}) &= \tfrac{1}{2}\sum_{i,p} x_{ip} \sum_{j} A_{ij}^{\text{phys}} x_{jp} \\ &= \tfrac{1}{2}\sum_{i,p} x_{ip} g_{ip}^{\text{phys}} \\ &= \tfrac{1}{2}\langle \mathbf{X}, \mathbf{G}^{\text{phys}} \rangle_F \\ &= \tfrac{1}{2}\,\text{Tr}(\mathbf{X}^T \mathbf{A}^{\text{phys}} \mathbf{X}). \end{aligned}$$

where $\langle \mathbf{U}, \mathbf{V} \rangle_F$ is the Frobenius inner product that sums the entrywise products of $\mathbf{U}$ and $\mathbf{V}$. Each entry $g_{ip}^{\text{phys}} = \sum_j A_{ij}^{\text{phys}} x_{jp}$ of the matrix $\mathbf{G}^{\text{phys}} = \mathbf{A}^{\text{phys}}\mathbf{X}$ counts the physical neighbours of radar i that are using chirp p.

C. Age of Information

The physical conflict graph by itself does not inform us which radars reliably communicate with each other within the stipulated timeframe. Let $\mathcal{H}$ denote the bidirected graph of timely communication. An edge between i and j means that the required assignments can be exchanged in both directions before an update depends on them.

The available conflict graph retains the physical-conflict pairs that can exchange current assignments in time. Its adjacency matrix is

$$\mathbf{A} = \mathbf{A}^{\text{phys}} \odot \mathbf{H},$$

where $\odot$ denotes the Hadamard Product. Define the set of physical neighbours $\mathcal{N}_i^{\text{phys}} = \{j : A_{ij}^{\text{phys}} = 1\}$ and uncertain neighbours $\mathcal{N}_i^{\text{unc}} = \{j \in \mathcal{N}_i^{\text{phys}} : A_{ij} = 0\}$. The latter are physical neighbours whose current assignments are unavailable.

For the available neighbours, we have

$$\mathbf{G} = \mathbf{A}\mathbf{X}, \qquad g_{ip} = \sum_j A_{ij} x_{jp}.$$

Each entry g_{ip} of $\mathbf{G} = \mathbf{A}\mathbf{X}$ tells radar i how many of its available conflicting neighbours use chirp p. These are the values that it can maintain exactly. Their collision energy is

$$E_{\text{col}}(\mathbf{X}) = \tfrac{1}{2}\sum_{i,p} x_{ip} g_{ip} = \tfrac{1}{2}\langle \mathbf{X}, \mathbf{G} \rangle_F = \tfrac{1}{2}\,\text{Tr}(\mathbf{X}^T \mathbf{A}\mathbf{X}).$$

When $\mathbf{A} = \mathbf{A}^{\text{phys}}$, it follows that $E_{\text{col}}(\mathbf{X}) = E_{\text{col}}^{\text{phys}}(\mathbf{X})$.

D. Stored assignments and Age of Information

How should radar i account for an uncertain neighbour $j \in \mathcal{N}_i^{\text{unc}}$? It may have received an earlier assignment, even though timely exchange is unavailable now ($H_{ij} = 0$). Let $\widetilde{\boldsymbol{x}}_{j|i}$ be the newest stored final row from j. Radar i needs an estimate of whether j uses each chirp when scheduling begins.

We use a simple mixture model: with weight ρ_{ij}, use the stored row; otherwise, treat all size-P_j subsets as equally likely. The latter choice gives each chirp probability P_j/M. If radar i has received at least one assignment from radar j, define

$$\pi_{ijp} = \rho_{ij}\widetilde{x}_{j|i,p} + (1-\rho_{ij})\frac{P_j}{M}. \tag{3}$$

Here π_{ijp} is radar i's modeled probability that neighbour j uses chirp position p. If no assignment from radar j has ever arrived, define

$$\pi_{ijp} = \frac{P_j}{M}.$$

In both cases,

$$\sum_{p=1}^{M} \pi_{ijp} = P_j.$$

Thus the model preserves radar j's fixed chirp quota in expectation. Here π_{ijp} is radar i's modeled probability of neighbour j using chirp p at the start of scheduling. The mixture is an assumption to calibrate, rather than a consequence of the quota alone. Its probabilities sum to P_j. If no row has ever arrived, $\rho_{ij} = 0$.

Older $\widetilde{x}_{j|i,p}$ should receive lesser weight, and conversely, newer $\widetilde{x}_{j|i,p}$ should receive greater weight. To achieve this,

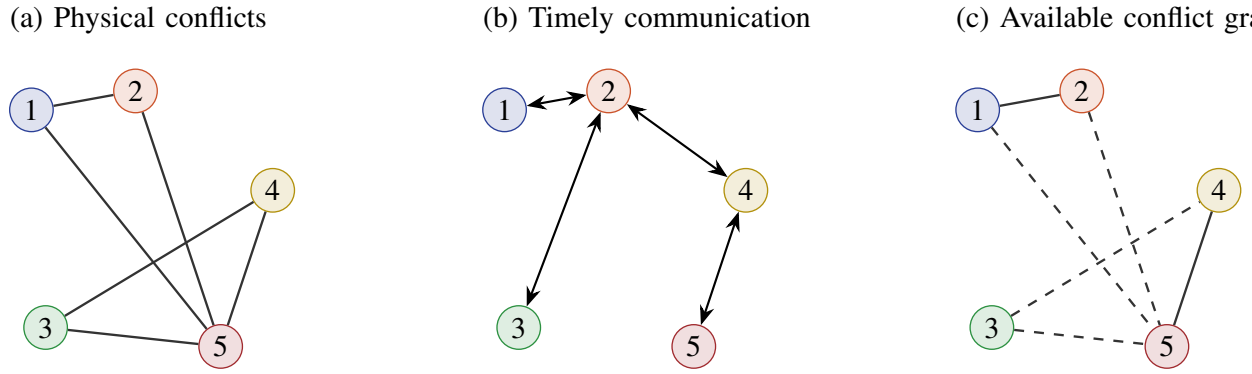

Fig. 2. Physical conflicts, control links, and available conflicts for the same five-radar system. Only direct exchanges meet the update deadline in this example. Thus the available conflict graph $\mathbf{A}$ has edges 12 and 45. In panel (c), dashed lines show the other physical conflicts for comparison; they are not edges of $\mathbf{A}$ and instead contribute stored information to $\mathbf{F}$. A connected control graph alone does not ensure that every current assignment arrives in time.

ρ_{ij} should be a nonincreasing function of the age a_{ij} of the most recently received row. Let us index coherent processing intervals (CPIs) by n and scheduling iterations by t, where $0 \le t \le T$. We consider that each radar starts interval n from its own final assignment for interval $n-1$, i.e., for all i and n,

$$\boldsymbol{x}_{i,n}^{(0)} = \boldsymbol{x}_{i,n-1}^{(T)}$$

Here $\boldsymbol{x}_{i,n}^{(t)}$ represemta radar i's assignment at iteration $t \in \{0,\dots,T\}$ of the scheduling window for CPI n.

For the final assignment $\boldsymbol{x}_{j,k}^{(T)}$ prepared for CPI k, define

$$\tau_{ij}(k) := \min \left\{ \begin{array}{c} n \ge k+1 : \quad \boldsymbol{x}_{j,k}^{(T)} \text{ is available} \\ \text{to radar } i \text{ when scheduling} \\ \text{for CPI } n \text{ begins} \end{array} \right\},$$

For the final row $\mathbf{x}_{j,k}^{(T)}$ produced by radar j during CPI k, let

$$\tau_{ij}(k) := \min \left\{ \begin{array}{l} n \ge k+1 : \mathbf{x}_{j,k}^{(T)} \text{ is available} \\ \text{to radar } i \text{ at the start of CPI } n \end{array} \right\},$$

with $\tau_{ij}(k) = \infty$ if the row is never received. The newest row from j available to i at the start of CPI n is therefore indexed by

$$u_{ij}(n) := \max\{k \le n-1 : \tau_{ij}(k) \le n\}, \tag{4}$$

where $u_{ij}(n) = -\infty$ if the set is empty. We define its age by

$$a_{ij}(n) = \begin{cases} n-1-u_{ij}(n), & u_{ij}(n) > -\infty, \\ \infty, & \text{no row has arrived.} \end{cases} \tag{5}$$

Age zero means that the preceding interval's final row is available. Choose $\rho_{ij} = \rho_{ij}(a_{ij}(n))$ nonincreasing, with $\rho_{ij}(0) = 1$, $\lim_{a\to\infty} \rho_{ij}(a) = 0$, and $\rho_{ij}(\infty) = 0$.

The stored row is replaced only when $u_{ij}(n)$ increases. Thus delayed but newer information is accepted, whereas obsolete deliveries are ignored. Otherwise, finite age increases by one. Note $\mathbf{H} = \mathbf{H}^T$ doesn't imply equal ages in the two directions.

E. The objective function

Summing (3) over the uncertain physical neighbours estimates how many of them use each chirp position. Define

$$f_{ip} = \tfrac{1}{2} \sum_{j \in \mathcal{N}_i^{\text{unc}}} \pi_{ijp}, \qquad \mathbf{F} = \begin{bmatrix} \boldsymbol{f}_1^T \\ \vdots \\ \boldsymbol{f}_N^T \end{bmatrix}.$$

The entry g_{ip} is the known count for neighbours whose current assignments are available, whereas $2f_{ip}$ is the corresponding model-based count for uncertain neighbours. Thus, $\mathbf{G}$ and $2\mathbf{F}$ describe, in essence, the same type of quantity using current and past information, respectively. The factor $1/2$ prevents counting an undirected overlap twice when both endpoint predictions equal the actual rows. Consequently $\langle \boldsymbol{f}_i, \boldsymbol{x}_i \rangle$ is half the predicted overlap at radar i.

We penalize both observed and predicted chirp index reuse:

$$E(\mathbf{X}) = \tfrac{1}{2} \operatorname{Tr}(\mathbf{X}^T \mathbf{A} \mathbf{X}) + \lambda \langle \mathbf{F}, \mathbf{X} \rangle_F, \qquad \lambda \ge 0. \tag{6}$$

The parameter λ sets the importance of predicted overlap relative to observed overlap. Predictions describe initial assignments. They stay fixed during scheduling and do not track later undelivered changes. Thus E is a surrogate for $E_{\text{col}}^{\text{phys}}$.

We denote $\mathbf{X}_n^{(t)}$ to be the system's assignment matrix at iteration t of the prescheduling window of CPI n. During one scheduling iteration, we hold $\mathbf{A}$, $\mathbf{F}$, the quotas, and the algorithm parameters fixed. Why? Since doing so enables us to preserve the detailed balance equations of the Gibbs scheduler. The final assignment is then held fixed for transmission.

Define $E_{\min} = \min_{\mathcal{X}} E$ and the objective gap $\mathcal{G}(\mathbf{X}) = E(\mathbf{X}) - E_{\min}$. BCD seeks small E through local minimization. Gibbs sampling instead assigns probability $\pi_\beta(\mathbf{X}) = Z_\beta^{-1} e^{-\beta E(\mathbf{X})}$, where Z_β is the normalization constant. A finite $\beta \ge 0$ is the inverse temperature: larger values favour lower costs. Note that this distribution over complete schedules is distinct from the neighbour predictions π_{ijp}.

F. Time scales and scheduling protocol

We distinguish three time scales for sensing and distributed coordination, as illustrated in Fig. 3. The chirp slot has duration

T_c; a coherent processing interval (CPI) contains $P_{\max}$ slots. For the parameters of [1], $T_c = 40\,\mu$s and $P_{\max} = 32$, giving a CPI of $1.28\,\mathrm{ms}$. The coherent scene time (CST) is the interval over which the active sectors and the physical conflict graph are assumed constant. We assume that the CST covers many CPIs (of the order of 100 in the illustrative experiments). The communication coherence time (CCT) is the interval over which the timely communication graph $\mathbf{H}$ remains unchanged; it need not coincide with the CST. The CST and CCT are assumed coherence intervals for the respective graphs.

In Fig. 3, the chirp pool coincides with one CPI. A pool comprising multiple consecutive CPIs is also admissible, provided that the required sensing coherence is maintained and the pool lies entirely within a CST. The current assignment is used for successive pools until the activation of a new assignment. Scheduling is initiated at start-up and upon a detected change in the physical or communication topology. The run can overlap sensing with the old assignment, provided that the hardware supports concurrent processing and messaging. The following four phases precede activation:

1) *Handshake:* determine the timely edges, accept the latest stored rows, and freeze $\mathbf{A}, \mathbf{F}, \lambda, \beta$ and the quotas.
2) *Scheduling:* execute one of the algorithms in Section III; required row or Gaussian-vector exchanges must complete before their dependent updates.
3) *Agreement:* confirm a common version and activation boundary for the candidate schedule.
4) *Guard:* reserve time for clock and communication uncertainty before the agreed boundary.

Activation of the new assignment occurs synchronously at the agreed CPI boundary (pool boundary in the multi-CPI case). If preparation fails, the previous assignment is retained. An unacknowledged exchange is treated as a failure of the corresponding required control link; the status of unprobed links remains unknown. A detected change in the communication topology during preparation necessitates either a restart under the updated model or retention of the active assignment. The analysis in Section IV therefore pertains to a fixed model within a completed preparation run. Coordinated activation presupposes sufficient synchronization and successful agreement; it is not an atomic-commit protocol for arbitrary message loss. Preparation latency comprises the computational, communication, agreement and guard durations.

The control loop is summarized in Fig. 4. In the original CPI-indexed notation, $\mathbf{X}_n^{(t)}$ is the tentative assignment prepared for an activation boundary, rather than a requirement to run a scheduler at every CPI. Between activation boundaries the applied schedule is the last committed one; the stored-row age continues to advance with elapsed CPIs. A refreshed row enters the next frozen model. This interpretation retains the within-run hypotheses of the theoretical guarantees.

III. Distributed chirp assignment algorithms

We consider four algorithms for the minimization of (6) over the feasible set (1). The first two use the single-radar conditional cost for stochastic sampling and exact minimization, respectively. Joint minimization over a pair of radars enlarges the class of admissible descent directions, whereas Gaussian augmentation permits parallel sampling. Within a preparation run, $\mathbf{A}$, $\mathbf{F}$, the quotas, λ, and β are held fixed in accordance with Section II-F. The tentative assignments are exchanged along the required control links; the assignment used for sensing remains unchanged until the new version is activated.

A. Random-scan Gibbs sampling and block coordinate descent

Let $\mathbf{G} = \mathbf{A}\mathbf{X}$, and let $\boldsymbol{g}_i^T$ denote its ith row. The terms of $E(\mathbf{X})$ that depend on $\boldsymbol{x}_i$ are linear in that row, since $A_{ii} = 0$. With $\boldsymbol{f}_i^T$ denoting row i of $\mathbf{F}$, the conditional cost vector and the corresponding Gibbs conditional are

$$\boldsymbol{c}_i = \boldsymbol{g}_i + \lambda \boldsymbol{f}_i, \quad p_i(\boldsymbol{x} \mid \mathbf{X}^{(-i)}) \propto e^{-\beta \langle \boldsymbol{c}_i, \boldsymbol{x} \rangle}, \quad \boldsymbol{x} \in \mathcal{X}_i, \quad (7)$$

where $\mathbf{X}^{(-i)}$ denotes the assignments of all radars other than i. In row DRGS, a radar is selected at random with a fixed positive selection probability, and its assignment is drawn from (7). All dependent updates use the new assignment. The resulting random-scan Gibbs kernel is reversible with respect to π_β and, at finite temperature, irreducible and aperiodic. For the analysis we adopt uniform random scan.

The distribution in (7) is a conditional Bernoulli distribution on subsets of cardinality P_i [12], [13]. Direct enumeration requires $\binom{M}{P_i}$ weights. Instead, define $R_i(p, r)$ as the total weight of all r-element subsets of $\{p, \dots, M\}$. With boundary value $R_i(M+1, 0) = 1$ and zero weights for infeasible states, the normalizing constant is obtained from the recurrence

$$R_i(p, r) = R_i(p+1, r) + e^{-\beta c_{ip}} R_i(p+1, r-1). \quad (8)$$

A backward evaluation of (8), followed by a forward sampling pass from $(p, r) = (1, P_i)$, gives an exact draw. At state (p, r), chirp p is included with conditional probability $e^{-\beta c_{ip}} R_i(p+1, r-1)/R_i(p, r)$; the remaining quota is decremented on inclusion. Multiplying the successive conditional probabilities gives the mass in (7). The arithmetic complexity is $O(MP_i)$; in the numerical implementation, the recurrence is evaluated in the logarithmic domain.

In block coordinate descent (BCD) [23], minimization with respect to a subset of the variables is performed with all other variables held fixed. In our discrete formulation, let $\mathcal{B} \subseteq \{1, \dots, N\}$ index the radar assignments comprising the block. Write $D_{\mathcal{B}}(\mathbf{X})$ for the resulting objective decrease under exact minimization. We restrict the computational subproblems to singleton and two-radar blocks.

The singleton subproblem reduces to selecting the P_i least-cost chirps in $\boldsymbol{c}_i$. For the two-radar subproblem, independent minimization of the two row costs would retain the obsolete mutual interaction. Removing that interaction gives the external cost vectors $\boldsymbol{a} = \boldsymbol{c}_i - A_{ij}\boldsymbol{x}_j$ and $\boldsymbol{b} = \boldsymbol{c}_j - A_{ij}\boldsymbol{x}_i$. The coupled subproblem is then

$$\sum_{p=1}^{M} (a_p y_p + b_p z_p + A_{ij} y_p z_p), \qquad \boldsymbol{y} \in \mathcal{X}_i, \ \boldsymbol{z} \in \mathcal{X}_j. \quad (9)$$

For each chirp, the four possible binary pairs (y_p, z_p) determine the increments in both quotas and the corresponding stage cost. A DP indexed by the chirp and the two remaining quotas

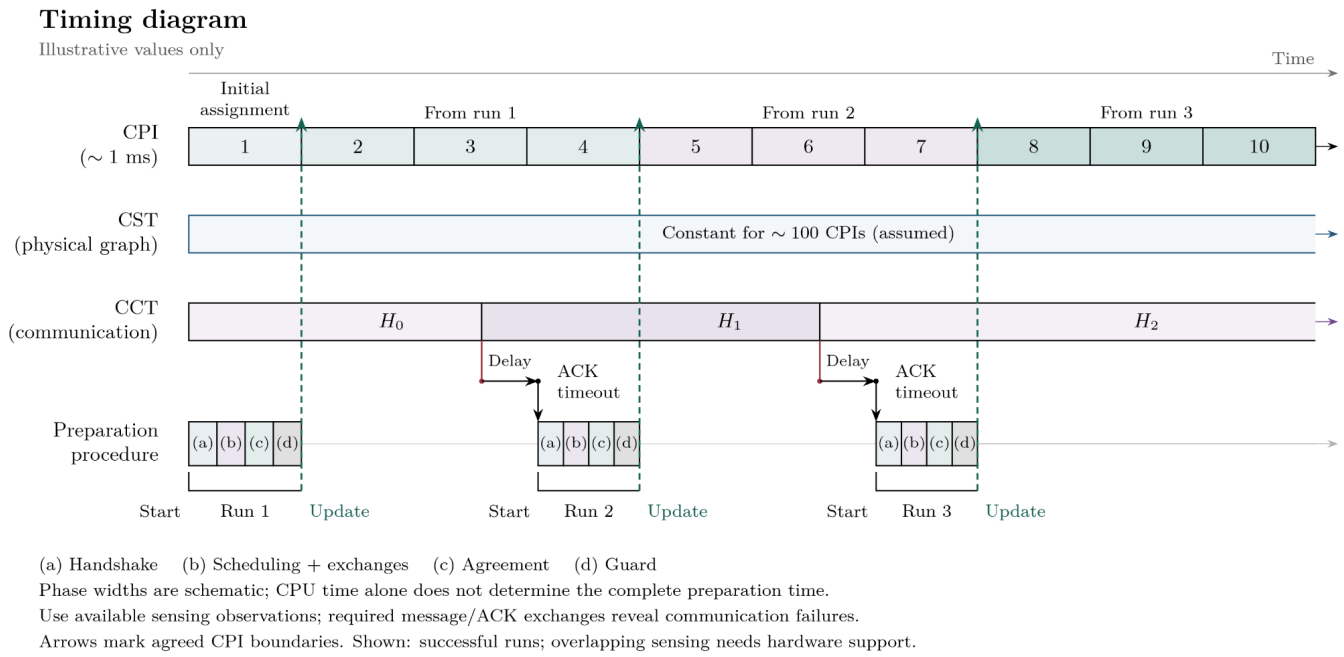

Fig. 3. Event-triggered scheduling on the CPI, CST and CCT time scales (schematic). The physical graph stays fixed over the displayed interval. After a change in the communication graph is detected, a preparation run with phases (a)–(d) computes a candidate; the dashed line marks activation at an agreed boundary. Until then, the previous schedule remains in force. The values and phase widths are illustrative.

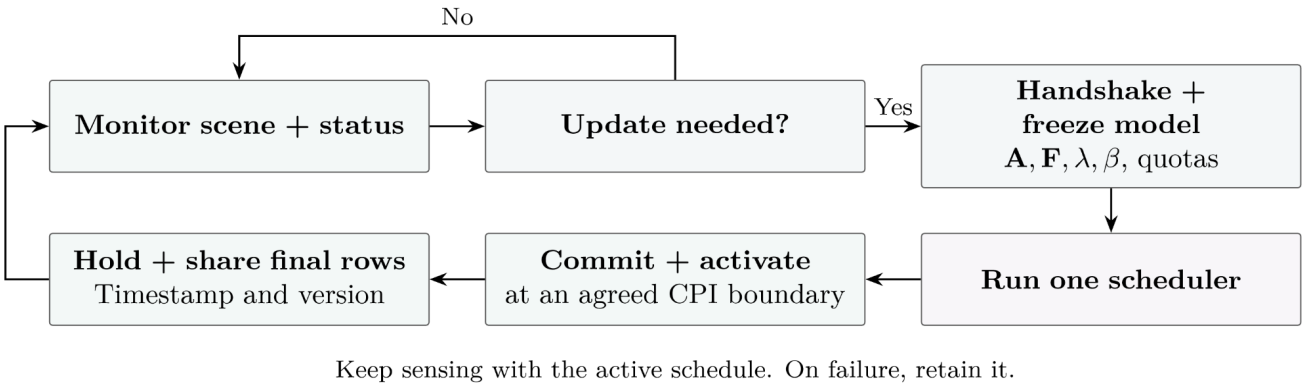

Fig. 4. Event-triggered control flow. The active assignment remains in service during monitoring and preparation. A candidate is activated only after the required agreement; a failed run leaves the previous assignment in use.

solves (9) in $O(M(P_i+1)(P_j+1))$ arithmetic operations. Both updated rows are committed before any dependent update.

We employ a tie-breaking rule that retains the incumbent whenever it is a block minimizer. Strict descent on the finite feasible set implies that only finitely many accepted BCD updates are possible. Blockwise optimality requires that no admissible block yield a strict decrease. At a coordinatewise minimum, a pair with $A_{ij} = 0$ cannot provide a decrease, since its conditional objective is separable.

The neighbour costs admit an incremental implementation. Let $\Delta \boldsymbol{x}_i^{(t)} = \boldsymbol{x}_i^{(t+1)} - \boldsymbol{x}_i^{(t)}$ denote a single-row change and $\boldsymbol{e}_i$ the ith coordinate vector in $\mathbb{R}^N$. The assignment increment is $\boldsymbol{e}_i(\Delta \boldsymbol{x}_i^{(t)})^T$, which gives

$$\mathbf{A}\mathbf{X}^{(t+1)} = \mathbf{A}\mathbf{X}^{(t)} + (\mathbf{A}\boldsymbol{e}_i)(\Delta \boldsymbol{x}_i^{(t)})^T, \tag{10}$$

Equation (10) is a rank-one update of the neighbour-count matrix. A radar replacing r_i chirps need only transmit $2r_i$ signed indices to each available conflicting neighbour. A pair update is the sum of the two corresponding rank-one increments.

B. Gaussian DRGS

Ordinary single-site Gibbs conditionals cannot generally be sampled independently at adjacent vertices without changing

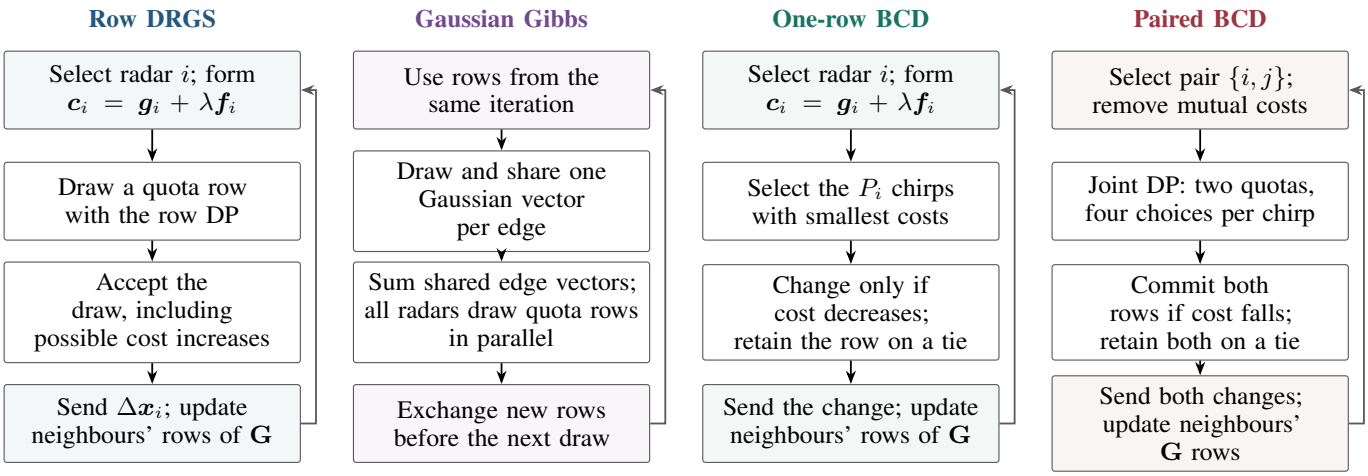

Fig. 5. Schematic representation of the four scheduling algorithms. The row Gibbs and paired BCD subproblems are solved by dynamic programming. Gaussian DRGS alternates shared edge-variable draws with conditionally independent row draws. Single-radar and two-radar updates require delivery of the updated assignments before subsequent dependent updates.

the invariant distribution. Following the auxiliary-Gaussian-variable approach to parallel sampling [14], [15], we construct a joint distribution of assignments and edge variables for which the assignments are conditionally independent across radars.

The construction exploits the fixed-cardinality constraint. Since $\|\boldsymbol{x}_i\|^2 = P_i$, the pairwise overlap satisfies

$$\langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle = \tfrac{1}{2}(P_i + P_j - \|\boldsymbol{x}_i - \boldsymbol{x}_j\|^2).$$

Orient the available conflict edges arbitrarily and let $\mathbf{B} \in \mathbb{R}^{|\mathcal{E}_A| \times N}$ be the incidence matrix, with row $[\boldsymbol{e}_i - \boldsymbol{e}_j]^T$ for edge $i \to j$. Then $\mathbf{L} = \mathbf{B}^T\mathbf{B} = \operatorname{diag}(\boldsymbol{d}) - \mathbf{A}$, where $d_i = \sum_j A_{ij}$ and $\mathcal{E}_A$ is the available edge set. Summation of the preceding identity over the edges gives

$$E(\mathbf{X}) = \tfrac{1}{2}\sum_i d_i P_i - \tfrac{1}{2}\|\mathbf{B}\mathbf{X}\|_F^2 + \lambda \langle \mathbf{F}, \mathbf{X} \rangle_F. \qquad (11)$$

The Hubbard–Stratonovich identity for a positive quadratic form is

$$\exp(\beta \|\boldsymbol{v}\|^2/2) = \mathbb{E}_{\boldsymbol{g} \sim N(\mathbf{0}, \mathbf{I})} \exp(\sqrt{\beta}\langle \boldsymbol{g}, \boldsymbol{v} \rangle). \qquad (12)$$

Applying (12) to the edge differences in (11) introduces a matrix of auxiliary Gaussian variables. Conditional on the assignments, the edge variables are independent normals with means obtained by completing the square. Conditional on the edge variables, the joint mass function of the assignments factorizes over $\mathcal{X}_1, \ldots, \mathcal{X}_N$. Integrating out the auxiliary variables recovers the desired Gibbs law.

Let $\mathbf{Z}$ have rows $\boldsymbol{z}_{ij}^T$ in the edge order of $\mathbf{B}$. A complete iteration of Gaussian DRGS consists of two conditional sampling stages, implemented as follows.

1) For every edge, obtain both endpoint assignments from iteration t. Independently generate
$$\boldsymbol{z}_{ij} = \sqrt{\beta A_{ij}}(\boldsymbol{x}_i - \boldsymbol{x}_j) + \boldsymbol{\varepsilon}_{ij}, \quad \boldsymbol{\varepsilon}_{ij} \sim N(\mathbf{0}, \mathbf{I}_M).$$
2) Initialize each $\boldsymbol{h}_i$ to zero. Both endpoints use the same edge vector. Add $\sqrt{A_{ij}}\boldsymbol{z}_{ij}$ to $\boldsymbol{h}_i$ and subtract it from $\boldsymbol{h}_j$, so $\boldsymbol{h}_i = (\mathbf{B}^T\mathbf{Z})_i^T$.

TABLE II
ARITHMETIC COST. COMMUNICATION AND WAITING TIME ARE ADDITIONAL.

Update	Cost	Main operation
Row DRGS	$O(MP_i)$	Subset weights and sampling
Row BCD	$O(M)$	Select lowest costs
Pair BCD	$O(M(P_i+1)(P_j+1))$	Four choices per state
Neighbour update	$O(\deg_{\mathbf{A}}(i) r_i)$	Add changed entries
Gaussian iteration	$O(M\lvert\mathcal{E}_A\rvert + M\sum_i P_i)$	Edge normals and row draws

3) After receiving all incident vectors, each radar independently samples
$$p_i(\boldsymbol{x} \mid \mathbf{Z}) \propto \exp\{\sqrt{\beta}\langle \boldsymbol{h}_i, \boldsymbol{x} \rangle - \beta \langle \lambda \boldsymbol{f}_i, \boldsymbol{x} \rangle\}, \quad \boldsymbol{x} \in \mathcal{X}_i. \qquad (13)$$
4) Deliver the new assignments before calculating the next Gaussian means.

The row draw in the second stage is again a fixed-cardinality conditional Bernoulli draw and can be performed using (8). The two endpoints of an edge must use the same realization of its auxiliary variable; independent copies would change the transition kernel. All edge draws use the same previous iteration.

Table II summarizes arithmetic complexity, with $\mathcal{E}_A = \{\{i, j\} : A_{ij} > 0\}$ and $\deg_{\mathbf{A}}(i)$ denoting the degree of radar i. The Gaussian complexity is aggregate work over the network, rather than parallel elapsed time. Communication, synchronization, and local computation must be accounted for separately.

IV. THEORETICAL GUARANTEES

Let $\mathbf{X}^* \in \arg\min_{\mathbf{X} \in \mathcal{X}} E(\mathbf{X})$ and denote the global minimum by $E_{\min}$. We measure suboptimality by the additive gap

$$\mathcal{G}(\mathbf{X}) = E(\mathbf{X}) - E(\mathbf{X}^*) = E(\mathbf{X}) - E_{\min}.$$

For the deterministic updates, our objective is to bound this gap in terms of the attainable blockwise decreases. For the stochastic updates, we distinguish finite-iteration suboptimality from convergence to the invariant distribution. Unless otherwise specified, all results concern a frozen scheduling instance with finite $\mathbf{F}$ and $\lambda \geq 0$. Guarantees for the physical objective are considered separately in Section V.

A. Objective error of row and pair updates

Coordinatewise optimality need not imply global optimality for the discrete, coupled objective. We derive an additive bound in which the residual is determined by interactions between distinct optimization blocks.

Let $\mathcal{P} = \{\mathcal{B}_1, \ldots, \mathcal{B}_J\}$ be a partition of the radars and let $\mathbf{A}^{\mathrm{cross}}$ be the adjacency matrix obtained by deleting the edges internal to the blocks. In the proposed implementation, each block contains at most two radars. Define $q_i = \min(P_i, M - P_i)$.

For any two feasible assignments of radar i, the difference has at most $2q_i$ nonzero entries of magnitude one. The absolute inner product of two row differences is therefore bounded by $2\min(q_i, q_j)$. Summing this bound over the crossing edges yields the residual

$$L_{\mathcal{P}} = 2 \sum_{i<j} A_{ij}^{\mathrm{cross}} \min(q_i, q_j). \tag{14}$$

For a block $\mathcal{B}$, define the improvement under an exact block minimization as

$$D_{\mathcal{B}}(\mathbf{X}) = E(\mathbf{X}) - \min_{\substack{\mathbf{Y}\in\mathcal{X} \\ \boldsymbol{y}_i = \boldsymbol{x}_i,\ i\notin\mathcal{B}}} E(\mathbf{Y}). \tag{15}$$

By feasibility of the incumbent, $D_{\mathcal{B}}(\mathbf{X}) \geq 0$, with equality precisely at a minimum with respect to block $\mathcal{B}$.

Theorem 1 (BCD error and pairing)**.** *For every feasible schedule,*

$$\mathcal{G}(\mathbf{X}) \leq \sum_{\mathcal{B}\in\mathcal{P}} D_{\mathcal{B}}(\mathbf{X}) + L_{\mathcal{P}}. \tag{16}$$

For this fixed partition and $t \geq 0$, choosing one set uniformly at each step and minimizing its conditional cost gives, with $r = 1 - 1/J$,

$$\mathbb{E}\mathcal{G}(\mathbf{X}^{(t)}) \leq r^t \mathcal{G}(\mathbf{X}^{(0)}) + (1 - r^t) L_{\mathcal{P}}. \tag{17}$$

If no set can improve the output, then $\mathcal{G}(\mathbf{X}) \leq L_{\mathcal{P}}$. A maximum-weight matching with weights $2A_{ij}\min(q_i, q_j)$ minimizes $L_{\mathcal{P}}$ over partitions into pairs and singletons.

The proof compares the incumbent with the blockwise substitutions of a global minimizer; the discrepancy between the sum of the blockwise comparisons and the full comparison consists of crossing-edge bilinear terms (Supplement, Section S3.1). Theorem 1 provides an a posteriori certificate without requiring the value of $E_{\min}$. At a blockwise minimum the certificate reduces to $L_{\mathcal{P}}$. Internalizing an edge in a two-radar block removes its contribution to (14); hence the partition minimizing this bound is obtained by maximum-weight matching. This criterion minimizes the upper bound, not the achieved objective. For adaptive block selection the terminal certificate remains valid if the output is optimal with respect to every block of the comparison partition; the geometric estimate (17) pertains to uniform sampling from a fixed partition.

The random-scan Gibbs update can likewise be analyzed as an inexact coordinate minimization, with an error arising from the conditional entropy. Write $H_{\mathcal{X}} = \log|\mathcal{X}| = \sum_i \log\binom{M}{P_i}$ and let L_{row} denote (14) for the singleton partition. For $\beta > 0$, uniform row DRGS satisfies

$$\begin{aligned}\mathbb{E}\mathcal{G}(\mathbf{X}^{(t)}) \leq\ & r_N^t \mathcal{G}(\mathbf{X}^{(0)}) \\ & + (1 - r_N^t)(L_{\mathrm{row}} + H_{\mathcal{X}}/\beta),\end{aligned} \tag{18}$$

where $r_N = 1 - 1/N$; use $0^0 = 1$ at $t = 0$.

See Supplement, Section S3.6, for the conditional free-energy argument.

In contrast to an estimate based on the stationary Gibbs law, (18) bounds the expected objective at every iteration without a mixing-time assumption.

B. Mixing and objective error of Gibbs sampling

Let $K(\mathbf{X}, \mathbf{Y})$ be the transition kernel of Gaussian DRGS after a complete auxiliary-variable and assignment update. The law after t iterations from initial distribution ν is νK^t.

A sufficient high-temperature condition for geometric convergence can be expressed in terms of the spectral norm of the graph Laplacian. Specifically, we obtain the following spectral-gap and total-variation bounds.

Theorem 2 (Gaussian Gibbs law and spectral gap)**.** *With the stated shared-edge implementation, Gaussian DRGS is reversible for π_β and has positive transition probabilities between every pair of feasible schedules at finite β. Let*

$$s = \tfrac{\beta}{2}\lambda_{\max}(\mathbf{L}) < 1. \tag{19}$$

Its spectral gap satisfies $\mathrm{gap}(K) \geq 1 - s$. For any initial law ν,

$$\|\nu K^t - \pi_\beta\|_{\mathrm{TV}} \leq \min\{1, \tfrac{1}{2}\sqrt{\chi^2(\nu\|\pi_\beta)}\, s^t\}. \tag{20}$$

Here $\chi^2(\nu\|\pi) = \sum_{\mathbf{X}} \nu(\mathbf{X})^2/\pi(\mathbf{X}) - 1$ denotes the chi-squared divergence. Use $s^0 = 1$. At $s = 0$, one Gaussian iteration draws from π_β independently of the previous schedule.

The augmented joint density establishes reversibility. The spectral estimate follows by bounding the fixed-cardinality conditional covariance and applying a variance inequality to the auxiliary-variable distribution [24], [25]; see Supplement, Sections S3.2 and S16.

The invariance statement is independent of the sufficient high-temperature condition in (19). The latter may fail at large β or on graphs with large Laplacian eigenvalues. Consequently neither parallelizability nor invariance alone provides a comparison of finite-time performance with random-scan Gibbs sampling [15].

For the random-scan kernel we use a Dobrushin-type interdependence matrix [26]. The influence of site j on site i is the supremum of the total-variation distance between two single-site conditional laws whose conditioning configurations differ only at j. An entrywise upper bound is

$$C_{ij} = \tanh\!\left(\frac{\beta A_{ij}}{2}\min(q_i, q_j)\right), \quad C_{ii} = 0. \tag{21}$$

with q_i as defined in Section IV-A. Put $\mathbf{C} = [C_{ij}]$. A sufficient contraction condition is the existence of a positive vector $\boldsymbol{w}$ and $\xi < 1$ such that $\mathbf{C}^T\boldsymbol{w} \le \xi\boldsymbol{w}$. Under a coupling in weighted Hamming distance, uniform random scan contracts the expected distance by a factor at most $1 - \alpha/N$, where $\alpha = 1 - \xi$. The resulting TV bound after t single-site updates is $\min\{1, (\sum_i w_i / \min_i w_i)(1 - \alpha/N)^t\}$.

The proof is given in Supplement, Section S3.5.

The mixing bounds yield objective guarantees through the oscillation of the energy function. Let ν_t be the output law, ε_t an upper bound on $\|\nu_t - \pi_\beta\|_{\mathrm{TV}}$, and $R_E = \max_{\mathcal{X}} E - E_{\min}$.

Corollary 3 (Objective error for both Gibbs algorithms)**.** *For either Gibbs algorithm at $\beta > 0$, with ν_t, ε_t, and R_E defined above,*

$$\mathbb{E}_{\nu_t} E - E_{\min} \le \min\{R_E, H_{\mathcal{X}}/\beta + R_E\varepsilon_t\}. \tag{22}$$

Use (20) *for Gaussian DRGS or the preceding row convergence bound for row DRGS, under their respective conditions.*

The proof combines the Gibbs variational principle with the total-variation bound for expectations of bounded functions (Supplement, Section S3.6). The two terms in (22) characterize the equilibrium approximation error and the finite-time sampling error, respectively. A higher inverse temperature improves the entropy bound but can degrade the mixing bound. For row DRGS, the smaller of (18) and (22) can be used whenever both apply.

C. Perturbations of the AoI-dependent field

We next examine the stability of the scheduling problem with respect to a perturbation of its unary costs. This accounts for changes in the AoI-dependent predictions between successive preparation runs.

For a row perturbation $\boldsymbol{v}$, introduce the oscillation seminorm on its feasible fixed-quota assignments,

$$\mathrm{sp}_i(\boldsymbol{v}) = \max_{\boldsymbol{x}\in\mathcal{X}_i}\langle \boldsymbol{v}, \boldsymbol{x}\rangle - \min_{\boldsymbol{x}\in\mathcal{X}_i}\langle \boldsymbol{v}, \boldsymbol{x}\rangle.$$

The extrema are obtained by summing the P_i largest and smallest entries of $\boldsymbol{v}$, respectively. The seminorm is invariant under addition of a constant multiple of $\mathbf{1}$.

Theorem 4 (Information age and objective stability)**.** *Compare two models with the same quotas, fixed $\mathbf{A}$, $\lambda \ge 0$, and finite $\beta \ge 0$. Let $\Delta\mathbf{F} = \mathbf{F}^{\mathrm{new}} - \mathbf{F}$, with row i equal to $\Delta\boldsymbol{f}_i^T$. The largest change in the relative cost of two feasible schedules is*

$$\Omega = \sum_i \mathrm{sp}_i(\lambda\Delta\boldsymbol{f}_i) = \max_{\mathcal{X}}(E^{\mathrm{new}} - E) - \min_{\mathcal{X}}(E^{\mathrm{new}} - E). \tag{23}$$

For every schedule, writing $\mathcal{G}^{\mathrm{new}} = E^{\mathrm{new}} - E^{\mathrm{new}}_{\min}$,

$$|\mathcal{G}^{\mathrm{new}}(\mathbf{X}) - \mathcal{G}(\mathbf{X})| \le \Omega. \tag{24}$$

The two target Gibbs distributions satisfy

$$\|\pi^{\mathrm{new}}_\beta - \pi_\beta\|_{\mathrm{TV}} \le \tanh(\beta\Omega/4). \tag{25}$$

TABLE III
SCHEDULING AND RECOVERY GUARANTEES.

Algorithm	Guarantee	Result and condition
Row BCD	Objective error	Theorem 1; singleton partition
Paired BCD	Objective error and pairing	Theorem 1; fixed partition for rate
Row DRGS	Objective error and TV convergence	(18); contracting influence matrix for TV
Gaussian DRGS	Objective error, spectral gap, and TV convergence	Theorem 2, Corollary 3; $s < 1$ for rate
All four	Objective sensitivity to age and new information	Theorem 4; common quotas
Both Gibbs	Change of target distribution	Theorem 4; common β
All four	Overlap, mutual coherence, and recovery	Proposition 5, Theorem 6; valid overlap and receiver bounds

In particular, suppose only finite ages of directed uncertain relations in $\mathcal{U}$ increase by one, with stored rows unchanged. Define

$$q_{ij} = \min(q_i, q_j), \tag{26}$$
$$W = \frac{\lambda}{2}\sum_{(i,j)\in\mathcal{U}} |\rho_{ij}(a_{ij}+1) - \rho_{ij}(a_{ij})| q_{ij}.$$

Then $\Omega \le W$, and also

$$\sup_{\mathcal{X}} |E^{\mathrm{new}} - E| \le W, \qquad |E^{\mathrm{new}}_{\min} - E_{\min}| \le W. \tag{27}$$

The gap bound follows from the oscillation of the objective perturbation, and the TV estimate follows from the corresponding bound on the likelihood ratio. For aging with fixed stored rows, the range of feasible overlaps bounds the oscillation by W (Supplement, Sections S3.5 and S13). In particular, a schedule with an old objective gap at most b has a new gap at most $b + \Omega$. This provides a quantitative justification for warm-starting the next preparation run. The rowwise evaluation of Ω can account for cancellation between perturbations from different neighbours; W is a sufficient bound expressed solely in terms of the age increments and quota-dependent overlap ranges.

More generally, for a transition $\mathbf{X} \to \mathbf{Y}$ with old-objective decrease $D = E(\mathbf{X}) - E(\mathbf{Y})$, the perturbed objective satisfies

$$\mathcal{G}^{\mathrm{new}}(\mathbf{Y}) \le \mathcal{G}(\mathbf{X}) - D + \Omega.$$

For BCD, $D \ge 0$, whereas Gibbs transitions need not be descent steps. An upper bound on the old objective gap of the output remains valid for the perturbed model upon adding Ω. By the triangle inequality, an old-target TV error of at most ε gives a new-target error bounded by $\min\{1, \varepsilon + \tanh(\beta\Omega/4)\}$. A simultaneous change in $\mathbf{A}$ contributes at most $\sum_{i<j} |A^{\mathrm{new}}_{ij} - A_{ij}| q_{ij}$ to the oscillation bound. These estimates quantify perturbations of the declared scheduling model; an AoI value does not by itself bound a neighbour's actual change of assignment.

V. PHYSICAL OVERLAP AND SPARSE RECOVERY

We now investigate the implications of chirp assignment for sparse target recovery. The relevant quantity at the receiver is

the physical overlap, rather than the cost of the incomplete-information optimization problem. The analysis is deterministic conditional on the selected measurements and hence applies to the output of any of the four schedulers.

Following range processing, the single-radar joint angle–Doppler estimation problem can be expressed as [1]

$$\boldsymbol{y}_i = \boldsymbol{\Phi}_i(\mathcal{S}_i)\boldsymbol{\theta}_i + \boldsymbol{w}_i, \qquad \boldsymbol{\Phi}_i(\mathcal{S}_i) = \mathbf{T}_i(\mathcal{S}_i) \otimes \mathbf{C}_i. \tag{28}$$

Here $\mathbf{C}_i \in \mathbb{C}^{n_i \times G_\theta}$ is the angular dictionary associated with the array, and $\mathbf{T}_i(\mathcal{S}_i) \in \mathbb{C}^{P_i \times G_D}$ is the Doppler dictionary evaluated at the selected chirp times. The dimensions G_θ and G_D are the angular and Doppler grid sizes, and n_i is the number of spatial measurements. The sparse vector $\boldsymbol{\theta}_i$ encodes the target amplitudes on the joint grid. The Kronecker structure in (28) shows that the quota alone does not characterize the sensing matrix: its slow-time factor depends on the particular choice of chirps.

A. Bounds on overlap from stored assignments

For an unavailable neighbour j, suppose a valid upper bound $r_{j|i}$ is known on the number of chirp substitutions between its stored and final assignments. Define the observed overlap $o_{ij} = \langle \boldsymbol{x}_i, \widetilde{\boldsymbol{x}}_{j|i} \rangle$ and the true overlap $w_{ij}(\mathbf{X}) = |\mathcal{S}_i \cap \mathcal{S}_j|$. The latter satisfies

$$\begin{aligned} \max\{0, P_i + P_j - M, o_{ij} - r_{j|i}\} \le w_{ij}, \\ w_{ij} \le \min\{P_i, P_j, o_{ij} + r_{j|i}\}. \end{aligned} \tag{29}$$

The interval follows from the fixed quotas and the fact that one substitution changes the overlap by at most one. If both directed estimates are available, their intervals may be intersected; denote the resulting upper endpoint by $\bar{w}_{ij}$. For an available conflict edge, the exact overlap replaces this bound. Consequently the sum of the upper endpoints over physical edges bounds $E_{\mathrm{col}}^{\mathrm{phys}}$.

The substitution bound must be valid for the output under consideration. In the absence of a model or deterministic constraint on the evolution of the assignments, the only general quota-based choice is $r_{j|i} = q_j$. A small age does not necessarily imply a small substitution count.

B. Receiver assumptions and coherence

To treat sparse interference explicitly, we adopt an augmented-dictionary formulation as in [21], [22]. Apply the measurement-selection and receiver operators of radar i to its candidate target and interference atoms and collect the resulting columns in $\boldsymbol{\Phi}_i^{\mathrm{aug}}$. Discard columns of zero norm and let $\boldsymbol{\Lambda}_i$ be the diagonal matrix of the positive column norms. With the corresponding unnormalized coefficient vector $\boldsymbol{\vartheta}_i$, the normalized observation model is

$$\begin{aligned} \mathbf{D}_i &:= \widetilde{\boldsymbol{\Phi}}_i^{\mathrm{aug}} = \boldsymbol{\Phi}_{i,\mathrm{nz}}^{\mathrm{aug}} \boldsymbol{\Lambda}_i^{-1}, \\ \boldsymbol{a}_i &:= \widetilde{\boldsymbol{\vartheta}}_i = \boldsymbol{\Lambda}_i \boldsymbol{\vartheta}_i, \qquad \boldsymbol{y}_i = \mathbf{D}_i \boldsymbol{a}_i + \boldsymbol{n}_i. \end{aligned} \tag{30}$$

The normalized dictionary $\mathbf{D}_i$ has unit-norm columns, and rescaling preserves the coefficient support. Coefficients of discarded zero columns are unidentifiable. The disturbance $\boldsymbol{n}_i$ includes measurement noise and any bounded modelling error. Partition the measurement space into mutually orthogonal subspaces indexed by chirp time, padding unobserved slots with zeros. Assume that each unit-norm target atom has squared norm $1/P_i$ in every selected subspace, as in the untapered model (28). For each normalized interference atom associated with neighbour j, assume that its energy fraction ℓ outside the shared chirps satisfies $0 \le \ell \le \delta_{ij} \le 1$. The case $\delta_{ij} = 0$ corresponds to exact confinement to the shared chirps. Timing mismatch and filtering must be represented by a valid leakage bound or included as a bounded residual in $\boldsymbol{n}_i$. For a dictionary with normalized columns $\boldsymbol{d}_a$, define the mutual coherence [20] as

$$\mu(\mathbf{D}_i) = \max_{a \ne b} |\boldsymbol{d}_a^* \boldsymbol{d}_b|. \tag{31}$$

We set empty maxima to zero. Write μ_i^{loc} for the coherence within the target dictionary, μ_i^{int} for the coherence within the interference dictionary, and μ_{ij}^{TI} for the maximum absolute inner product of a target atom with an interference atom associated with j. Then

$$\mu(\mathbf{D}_i) = \max\{\mu_i^{\mathrm{loc}}, \mu_i^{\mathrm{int}}, \max_{j \in \mathcal{N}_i^{\mathrm{phys}}} \mu_{ij}^{\mathrm{TI}}\}. \tag{32}$$

Chirp overlap constrains the cross-dictionary coherence, not the two within-dictionary coherences. The following proposition bounds the cross-coherence in terms of the chirp overlap and leakage, and gives a sufficient condition for exact recovery by basis pursuit (BP) and orthogonal matching pursuit (OMP).

Proposition 5 (Mutual coherence from chirp overlaps)**.** *Assume the receiver model above. For every target column $\boldsymbol{\phi}$ and interference column $\boldsymbol{\psi}$ from j, put $a = w_{ij}/P_i$. Then*

$$\begin{aligned} |\boldsymbol{\phi}^* \boldsymbol{\psi}| &\le g(a, \delta_{ij}), \\ g(a,\delta) &= \begin{cases} \sqrt{a(1-\delta)} + \sqrt{(1-a)\delta}, & a + \delta < 1, \\ 1, & a + \delta \ge 1. \end{cases} \end{aligned} \tag{33}$$

This bound is attained under the stated energy information alone when the required nonzero subspaces exist. A valid upper bound on w_{ij} can replace it in a. With no leakage, the bound is $\sqrt{w_{ij}/P_i}$.

In particular,

$$\mu_{ij}^{\mathrm{TI}} \le g(w_{ij}/P_i, \delta_{ij}). \tag{34}$$

With no leakage, define $\mu_{\max}^{\mathrm{TI}} = \max_{i,j} \mu_{ij}^{\mathrm{TI}}$ and $P_{\min} = \min_i P_i$. Then

$$\mu_{\max}^{\mathrm{TI}}(\mathbf{X}) \le \min\{1, \sqrt{E_{\mathrm{col}}^{\mathrm{phys}}(\mathbf{X})/P_{\min}}\}. \tag{35}$$

Let K_i and Q_i bound the numbers of nonzero target and interference coefficients, respectively, with $K_i + Q_i \ge 1$. Noiseless basis pursuit and normalized OMP recover them if

$$\boxed{\mu(\mathbf{D}_i) < \frac{1}{2(K_i + Q_i) - 1}.} \tag{36}$$

OMP stops at zero residual. The conditions concern retained normalized coefficients and apply to all four schedulers.

The cross-coherence bound follows from Cauchy–Schwarz applied separately to the shared and complementary subspaces,

and the recovery statement follows from the classical mutual-coherence criterion [20] (Supplement, Section S18.1). A reduction in physical overlap tightens the cross-dictionary bound. It need not reduce the coherence of the target dictionary: the selected slow-time samples can produce Doppler ambiguities even when the assignment is collision-free. All three terms in (32) are therefore required for the recovery certificate.

C. Recovery using individual correlations

A scalar mutual-coherence bound can be conservative because it replaces every off-diagonal Gram entry by the largest absolute correlation. To exploit nonuniform correlations, we formulate a support-dependent sufficient condition for the exact recovery condition (ERC) of [20].

Suppress the radar index and let $\mathbf{D}$ denote the normalized augmented dictionary. For a candidate support $\mathcal{I}$ of cardinality $k > 0$, let $\mathbf{D}_{\mathcal{I}}$ be the corresponding subdictionary. Suppose $\mathbf{V}_{\mathcal{I}}$ is symmetric, entrywise nonnegative, has zero diagonal, and bounds the absolute off-diagonal entries of $\mathbf{D}_{\mathcal{I}}^{*}\mathbf{D}_{\mathcal{I}}$. For each $b \notin \mathcal{I}$, suppose $\boldsymbol{v}_b \geq |\mathbf{D}_{\mathcal{I}}^{*}\boldsymbol{d}_b|$ entrywise. The cross-dictionary entries may be bounded using (33).

Define the spectral-radius and inverse-Gram comparison quantities

$$\varrho = \rho(\mathbf{V}_{\mathcal{I}}), \qquad \eta = \max_{b\notin\mathcal{I}} \mathbf{1}^{T}(\mathbf{I} - \mathbf{V}_{\mathcal{I}})^{-1}\boldsymbol{v}_b, \tag{37}$$

where the second quantity is defined if $\varrho < 1$. This condition guarantees invertibility of the restricted Gram matrix. The nonnegative inverse comparison bounds $\max_{b\notin\mathcal{I}} \|\mathbf{D}_{\mathcal{I}}^{\dagger}\boldsymbol{d}_b\|_1$ by η. In particular, $\eta < 1$ implies the ERC.

Theorem 6 (Recovery using individual correlations)**.** *An empty off-support maximum is zero. If $\varrho < 1$ and $\eta < 1$, basis pursuit and normalized OMP recover every noiseless coefficient vector supported on $\mathcal{I}$; OMP stops when its residual is zero.*

For observations $\boldsymbol{y} = \mathbf{D}\boldsymbol{a} + \boldsymbol{n}$ with noise norm $\|\boldsymbol{n}\|_2 \leq \varepsilon$, let $\widehat{\boldsymbol{a}}$ minimize $\|\boldsymbol{z}\|_1$ subject to $\|\mathbf{D}\boldsymbol{z} - \boldsymbol{y}\|_2 \leq \varepsilon$. Under the same conditions $\varrho < 1$ and $\eta < 1$, allowing coefficients outside $\mathcal{I}$,

$$\|\widehat{\boldsymbol{a}} - \boldsymbol{a}\|_1 \leq \frac{4\varepsilon\sqrt{k}}{(1-\eta)\sqrt{1-\varrho}} + \frac{2(1+\eta)}{1-\eta}\|\boldsymbol{a}_{\mathcal{I}^c}\|_1. \tag{38}$$

All statements apply to a fixed output of row DRGS, Gaussian DRGS, row BCD, or paired BCD. A uniform recovery guarantee requires the inequalities for every allowed coefficient support.

The exact-recovery statement follows by comparison of the inverse Gram matrix with $(\mathbf{I} - \mathbf{V}_{\mathcal{I}})^{-1}$ and application of the ERC [20]. The stability estimate follows from residual feasibility and the ℓ_1 optimality inequality; see Supplement, Sections S14.1 and S18.2. The two terms in (38) account for measurement noise and the coefficient mass outside the chosen support, respectively. The bound thus applies to approximately sparse signals and specializes to exact recovery in the noiseless, exactly sparse case. Its constants depend explicitly on the conditioning and separation margins $1 - \varrho$ and $1 - \eta$.

For the uniform comparison $\mathbf{V}_{\mathcal{I}} = \mu(\mathbf{1}\mathbf{1}^{T} - \mathbf{I})$ and $\boldsymbol{v}_b = \mu\mathbf{1}$, one obtains $\eta \leq k\mu/[1-(k-1)\mu]$ provided $(k-1)\mu < 1$. The sufficient condition reduces to $(2k-1)\mu < 1$. The nonuniform comparison can certify supports for which this scalar criterion fails. Exact support identification in the presence of noise additionally requires a lower bound on the nonzero coefficient magnitudes.

Finally, for zero leakage, if the measurement assumptions hold for every admissible output and $\mathbb{E}E_{\mathrm{col}}^{\mathrm{phys}} \leq H$, Jensen's inequality gives

$$\mathbb{E}\mu_{\max}^{\mathrm{TI}} \leq \min\{1, \sqrt{H/P_{\min}}\} \tag{39}$$

where H is an upper bound for the physical collision energy. This expectation bound is not a uniform sparse-recovery guarantee. Probabilistic recovery results, such as [27], rely on specific assumptions on the random supports or coefficients that are not supplied merely by convergence of a scheduler to its Gibbs law.

VI. Cross-echo interference model

Mutual interference between FMCW radars depends on the transmitted chirp slopes, relative timing and propagation paths [3]. For chirps with unequal slopes, the interference at the output of a dechirping receiver may be concentrated in short time intervals; such structure is exploited in [10]. Here we consider a different mechanism, namely the scattering of a neighbouring radar's transmission by a target within the victim radar's field of view. Under aligned, equal-slope chirps, the interfering return can contribute to the same range cell as a desired echo. Its slow-time phase, however, need not coincide with that of the monostatic return. We extend the range–Doppler–angle model of [1] to incorporate these target-mediated returns. Direct coupling, receiver saturation and interference across distinct chirp slots are excluded from this model.

A. Single-radar range compression

Consider a radar with Q_i virtual antenna elements and $N_s = f_s T_c$ fast-time samples per chirp. Denote the carrier wavelength by λ_c and the pth selected chirp index in increasing order by $\zeta_{i,p} \in \mathcal{S}_i$. Under the narrowband and stop-and-hop approximations, the response of target k is separable in the spatial, slow-time and fast-time dimensions [1]. In particular, the dechirped response with unit amplitude is

$$\psi_{ik}[q,p,l] = c_{ik}[q]\,b_{ik}[p]\,e^{\mathrm{j}2\pi f_{ik} l/f_s}, \tag{40}$$

where f_{ik} denotes the beat frequency. For monostatic range R_{ik} and radial velocity v_{ik}, its leading terms are $2SR_{ik}/c_0$ and $2v_{ik}/\lambda_c$, subject to the mixer sign convention; S denotes the chirp slope and c_0 the speed of light. The spatial and slow-time steering sequences are given by

$$c_{ik}[q] = e^{\mathrm{j}2\pi(d_{i,q}/\lambda_c)\sin\theta_{ik}}, \tag{41}$$

$$b_{ik}[p] = e^{\mathrm{j}2\pi(2v_{ik}T_c/\lambda_c)\zeta_{i,p}}. \tag{42}$$

Here $d_{i,q}$ denotes the virtual antenna position, i.e., the sum of the corresponding transmit and receive antenna positions. The spatial and slow-time steering sequences give the two factors of the Kronecker-structured sensing matrix in Section V.

Assume additive complex white Gaussian noise with variance σ^2 per fast-time sample. Under a unitary DFT, the range

coefficient of an on-grid target with complex amplitude a_{ik} is $\sqrt{N_s}a_{ik}$. We normalize the output by $\sqrt{N_s}$, so that the target amplitude is unchanged and the resulting noise variance is σ^2/N_s. For off-grid targets the amplitude is additionally weighted by the range-window response; the numerical compression gain considered below assumes on-grid targets.

B. Bistatic paths and collision gating

Consider the target-mediated interference path $j \to k \to i$ in Fig. 6. The propagation delay for this bistatic path is $(R_{jk} + R_{ik})/c_0$, compared with $2R_{ik}/c_0$ for the monostatic path. Consequently, even for identical chirp slopes, the two returns need not have identical beat frequencies. For rectangular fast-time sampling, the normalized response in DFT bin r to a beat frequency f is

$$D_{N_s}(f/f_s - r/N_s) = \frac{1}{N_s}\sum_{l=0}^{N_s-1} e^{\mathrm{j}2\pi(f/f_s - r/N_s)l}. \tag{43}$$

The admitted power is weighted by $|D_{N_s}|^2$. More generally, D_{N_s} is replaced by the normalized Fourier transform of the range window. Thus simultaneous transmission is necessary for interference in the aligned-slot model but does not imply unit admission to the range cell.

Let $\mathcal{K}_i$ denote the set of monostatic targets associated with the range cell under consideration and $\mathcal{I}_{j\to i}$ the set of targets contributing bistatic returns from radar j to that cell. The latter is defined for an ordered pair of radars. In the same-cell approximation used in the numerical experiments, $\mathcal{I}_{j\to i} = \mathcal{K}_i \cap \mathcal{T}_j$, where $\mathcal{T}_j$ denotes the targets illuminated by j. The general model permits a target outside $\mathcal{K}_i$ to contribute to $\mathcal{I}_{j\to i}$ through its bistatic delay.

Let $a^{\times}_{jik}$ denote the complex amplitude of the bistatic return after normalized range compression, including propagation loss, transmit and receive gains, scattering and range admission. For a common target with nonzero monostatic amplitude, define the cross-to-own echo power ratio as

$$\chi_{jik} = \frac{|a^{\times}_{jik}|^2}{|a_{ik}|^2}. \tag{44}$$

For homogeneous coupling we write $\chi_{jik} = g$ and define $\mathrm{ISR}_{\mathrm{echo}} = 10\log_{10} g$. The per-sample interference-to-noise ratio is instead determined by the aggregate interference power: I_i/σ^2 on the raw-input scale for on-bin echoes, and $N_s I_i/\sigma^2$ on the range-compressed scale. Under identical antenna and scattering factors the bistatic-to-monostatic power ratio includes the factors $P^{\mathrm{tx}}_j/P^{\mathrm{tx}}_i$ and R^2_{ik}/R^2_{jk}. For the phase model, let $\vartheta_{ji}[\zeta]$ denote the relative carrier phase. We assume these phases to be independent and uniformly distributed across chirps and ordered radar pairs, with the target-dependent path phases fixed within the CPI. This is an idealization of noncoherent transmitters.

Although the arrival direction at the receiving array is determined by the target position, the transmit-channel response and bistatic Doppler frequency need not coincide with those of the monostatic return. For virtual channel $q \equiv (n, m)$, indexed by transmit channel n and receive channel m, write

$$c^{\times}_{jik}[q] = t_{jik}[n]e^{\mathrm{j}2\pi(d^R_{i,m}/\lambda_c)\sin\theta_{ik}}, \tag{45}$$

$$b^{\times}_{jik}[p] = e^{\mathrm{j}2\pi(v_{jk}+v_{ik})T_c\zeta_{i,p}/\lambda_c}, \tag{46}$$

where $t_{jik}[n]$ denotes the effective transmit-channel signature and $d^R_{i,m}$ the receive antenna position. In the equal-power model below, the steering-sequence entries have unit magnitude; channel-dependent gains may otherwise be incorporated in the power weights.

The received data in the selected range cell, indexed by virtual channel and selected chirp, are therefore given by

$$\begin{aligned} Y_{i,q,p} = &\sum_{k\in\mathcal{K}_i} a_{ik}c_{ik}[q]b_{ik}[p] \\ &+ \sum_{j\in\mathcal{N}^{\mathrm{phys}}_i} x_{j,\zeta_{i,p}}e^{\mathrm{j}\vartheta_{ji}[\zeta_{i,p}]} \\ &\quad\cdot \sum_{k\in\mathcal{I}_{j\to i}} a^{\times}_{jik}c^{\times}_{jik}[q]b^{\times}_{jik}[p] \\ &+ W_{i,q,p}, \quad W_{i,q,p} \sim \mathcal{CN}(0,\sigma^2/N_s). \end{aligned} \tag{47}$$

Collect the spatial and slow-time steering sequences in the columns of $\mathbf{U}_i, \mathbf{V}_i$ for the monostatic returns and $\mathbf{U}^{\times}_{ji}, \mathbf{V}^{\times}_{ji}$ for the bistatic returns. Define $\mathbf{D}_{ji} = \mathrm{diag}(x_{j,\zeta_{i,p}}e^{\mathrm{j}\vartheta_{ji}[\zeta_{i,p}]})_{p=1}^{P_i}$. The matrix form of (47) is

$$\begin{aligned} \mathbf{Y}_i = &\mathbf{U}_i\,\mathrm{diag}(\boldsymbol{a}_i)\mathbf{V}_i^T \\ &+ \sum_{j\in\mathcal{N}^{\mathrm{phys}}_i} \mathbf{U}^{\times}_{ji}\,\mathrm{diag}(\boldsymbol{a}^{\times}_{ji})(\mathbf{V}^{\times}_{ji})^T\mathbf{D}_{ji} + \mathbf{W}_i. \end{aligned} \tag{48}$$

Upon vectorization, the bistatic dictionary is formed from Kronecker products of spatial steering vectors and masked Doppler steering vectors. The latter have support of cardinality at most w_{ij}. Thus the same-slot model satisfies the zero-leakage hypothesis of Proposition 5, and the target–interference coherence is bounded by $\sqrt{w_{ij}/P_i}$ for uniform target energy across the selected slots. No distributional assumption on the phases is required for this bound.

In (47), the observation index is already restricted to $\mathcal{S}_i$; the multiplicative factor $x_{j,\zeta_{i,p}}$ implements the remaining collision constraint (Fig. 7). The interference sum extends over all physical neighbours, irrespective of communication availability. The communication matrix $\mathbf{H}$ affects the optimization problem but not the propagation model.

For the range-cell simulations, we adopt the specialization in which the targets lie in the common illuminated region, the monostatic and bistatic angle–velocity steering sequences coincide, and $|a_{ik}| = 1$. The bistatic coefficient is $a^{\times}_{jik} = \sqrt{g}a_{ik}e^{\mathrm{j}\xi_{jik}}$ with independent uniform target phases ξ_{jik}; the common phase $\vartheta_{ji}[\zeta]$ is sampled independently for each collided chirp. This specialization isolates the effect of chirp-dependent phase perturbations from differences in range and Doppler. It excludes the coherent-copy case, in which the cross return changes the effective target amplitude.

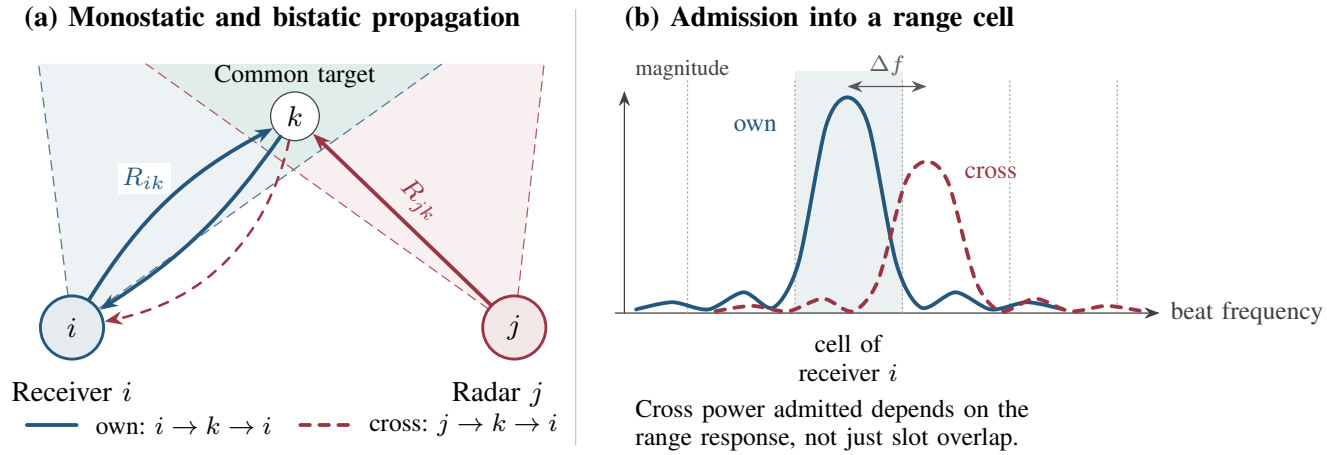

Fig. 6. Own and cross-echo paths for a target in the common illuminated region. The blue path is monostatic; the red path is bistatic. The dashed return is the cross echo. Panel (b) illustrates the range-compression response: distinct path lengths shift the beat frequency, and only the admitted part of the response enters the cell. Curves are schematic, not measured spectra. Direct radar-to-radar coupling is omitted.

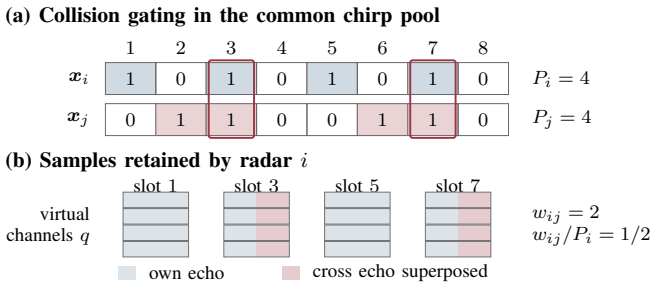

Fig. 7. The collision gate. In this illustrative 4-of-8 assignment, slots 3 and 7 carry cross echoes at radar i; the others contain only its own echoes and noise. Each displayed column represents the virtual-channel samples for that slot.

C. Received power and the collision objective

Let $w_{ij} = \langle \boldsymbol{x}_i, \boldsymbol{x}_j \rangle$ denote the physical chirp overlap. Define the aggregate admitted bistatic power from transmitter j to receiver i as

$$\Xi_{ij} = \sum_{k \in \mathcal{I}_{j \to i}} |a^{\times}_{jik}|^2, \qquad \Xi_{ii} = 0. \tag{49}$$

We average the squared modulus of the interference over virtual channels, selected chirps and phase realizations. Independence of the zero-mean target-specific phase factors eliminates the mixed terms within each bistatic sum; independence of the relative carrier phases eliminates the mixed terms between transmitters. A common slot phase alone would not eliminate the former terms. For unit-modulus steering sequences, $S_{ik} = |a_{ik}|^2$ denotes the monostatic power of target k, and the interference power per range-compressed sample is

$$I_i(\mathbf{X}) = \frac{1}{P_i} \sum_{j \in \mathcal{N}_i^{\text{phys}}} \Xi_{ij} w_{ij}. \tag{50}$$

Accordingly, the per-target SINR at the range-compression output and its input-referred counterpart are

$$\text{SINR}_{ik}^{\text{cell}} = \frac{S_{ik}}{\sigma^2/N_s + I_i(\mathbf{X})}, \tag{51}$$

$$\text{SINR}_{ik}^{\text{eff}} = \frac{S_{ik}}{\sigma^2 + N_s I_i(\mathbf{X})}. \tag{52}$$

The input-referred SINR differs from the range-cell SINR by the compression gain $10 \log_{10} N_s$ in dB and reduces to the input SNR S_{ik}/σ^2 in the absence of cross echoes. Other monostatic targets are treated as signal components to be estimated jointly. For on-bin returns the raw-sample SINR is $S_{ik}/(\sigma^2 + I_i)$; off-bin returns require separate accounting for raw and admitted echo power. Hence the raw-sample SINR is not interchangeable with (52).

The total interference power weighted by the corresponding chirp quotas satisfies

$$\begin{aligned} \sum_i P_i I_i(\mathbf{X}) &= \sum_{i<j} (\Xi_{ij} + \Xi_{ji}) w_{ij} \\ &= \text{Tr}(\mathbf{X}^T \boldsymbol{\Xi} \mathbf{X}), \end{aligned} \tag{53}$$

where $\boldsymbol{\Xi} = [\Xi_{ij}]$ vanishes outside the physical edges. Under homogeneous coupling $\Xi_{ij} = \bar{\Xi} A_{ij}^{\text{phys}}$, the total reduces to $2\bar{\Xi} E_{\text{col}}^{\text{phys}}$. In the heterogeneous case, the symmetrized edge weights $\Xi_{ij} + \Xi_{ji}$ give the corresponding weighted collision objective with the trace convention in (6). If $\xi_- \le (\Xi_{ij} + \Xi_{ji})/2 \le \xi_+$ on every physical edge, the total lies between $2\xi_- E_{\text{col}}^{\text{phys}}$ and $2\xi_+ E_{\text{col}}^{\text{phys}}$. This establishes proportionality under homogeneous coupling and comparison bounds otherwise; it does not identify the physical energy with the available-information surrogate.

Under independent uniform chirp selection, the overlap between two radars is hypergeometrically distributed:

$$\Pr(w_{ij} = h) = \frac{\binom{P_i}{h}\binom{M-P_i}{P_j-h}}{\binom{M}{P_j}}, \tag{54}$$

where impossible binomial coefficients are zero. Thus $\mathbb{E}w_{ij} = P_iP_j/M$ and

$$\mathbb{E}I_i = \sum_j \Xi_{ij}P_j/M. \tag{55}$$

For common quota P, K unit-power targets in the same cell of every radar, common cross-to-own power g and a complete physical graph, let $c_i = \sum_{j\neq i} w_{ij}$. Then

$$I_i = gKc_i/P, \qquad \sum_i I_i = 2gKE_{\mathrm{col}}^{\mathrm{phys}}/P, \tag{56}$$

$$\mathrm{SINR}_{ik}^{\mathrm{eff}} = (\sigma^2 + N_s gKc_i/P)^{-1}, \tag{57}$$

$$\mathbb{E}I_i = gK(N-1)P/M. \tag{58}$$

The ratio of ensemble powers is $[\sigma^2 + N_s gK(N-1)P/M]^{-1}$, whereas the ensemble mean SINR is the expectation of (57); averaging in dB gives a third quantity. These distinctions are relevant in particular when collision-free assignments have nonzero probability. Equation (58) also shows that M proportional to N does not imply vanishing interference on a complete graph, and that the power depends on target congestion within the cell as well as slot occupancy.

Conditioned on $\mathcal{S}_i$, the overlaps with independently sampled neighbouring rows are mutually independent, and their hypergeometric laws do not depend on the realized $\mathcal{S}_i$. Hence the distribution of c_i is given by convolution of the individual overlap laws. In particular, the probability that radar i is free of collisions is

$$\Pr(c_i = 0) = \prod_{j\in\mathcal{N}_i^{\mathrm{phys}}} \frac{\binom{M-P_i}{P_j}}{\binom{M}{P_j}}. \tag{59}$$

The corresponding counts across distinct receiving radars need not be independent. If the conditional detection and false-discovery probabilities given c_i are available, their marginal values follow by averaging over its distribution. These conditional probabilities depend on the receiver and scene statistics. A collision-free assignment removes the bistatic term in (47), but the resulting monostatic dictionary remains dependent on the selected chirp indices.

D. The matched-filter statistic and adaptive detection

For an angle–velocity hypothesis $\omega = (\theta, v)$, form the steering vectors $\boldsymbol{c}_i(\theta)$ and $\boldsymbol{b}_i(v)$ using the actual antenna positions and chirp indices. The unit-norm matched-filter statistic is

$$T_i(\omega) = \frac{|\boldsymbol{c}_i(\theta)^H \mathbf{Y}_i \overline{\boldsymbol{b}_i(v)}|^2}{Q_iP_i}. \tag{60}$$

At the true parameters of target k_0, the monostatic contribution is coherent over all Q_iP_i measurements. In contrast, the assumed independent relative phases eliminate coherent summation of the bistatic contributions across chirps, although spatial coherence may persist. Neglecting the other monostatic returns, or assuming their exact subtraction, the ratio of desired matched-filter power to expected disturbance power becomes

$$\frac{Q_iP_iS_{ik_0}}{\sigma^2/N_s + (Q_i/P_i)\sum_j w_{ij}\sum_{k\in\mathcal{I}_{j\to i}} |a_{jik}^{\times}|^2\kappa_{ik_0,jk}}, \tag{61}$$

where $\kappa_{ik_0,jk} = |\boldsymbol{c}_{ik_0}^H \boldsymbol{c}_{jik}^{\times}|^2/(\|\boldsymbol{c}_{ik_0}\|^2\|\boldsymbol{c}_{jik}^{\times}\|^2)$ is the squared normalized spatial correlation. For isotropically distributed unit-norm interference directions its expectation is $1/Q_i$; for identical steering vectors it is unity. The corresponding interference powers at the matched-filter output differ by a factor of Q_i (approximately 9 dB for eight virtual elements) at fixed average incident interference power. Consequently the sample SINR is not a sufficient statistic for the distribution of (60), and noise-only performance curves need not predict interference-limited detection.

For joint angle–Doppler estimation in the range-cell experiments, we use a grid-based OMP procedure motivated by [1], [20], with cyclic support refinement followed by backward elimination. Each forward iteration selects the atom with the largest absolute correlation with the residual and updates the coefficients by least squares. Unlike the full-waveform OMP+Binary and OMP+Range-OMP receivers, this procedure operates on a prescribed range cell. Let $\mathbf{B}$ denote the matrix of selected unnormalized atoms, $\hat{\boldsymbol{a}}$ their least-squares coefficients and $\boldsymbol{r} = \mathrm{vec}(\mathbf{Y}_i) - \mathbf{B}\hat{\boldsymbol{a}}$ the residual. If $\mathbf{B}$ has full column rank, the increase in residual energy upon deleting atom k and refitting the remaining coefficients is

$$D_k = \frac{|\hat{a}_k|^2}{[(\mathbf{B}^H\mathbf{B})^{-1}]_{kk}}. \tag{62}$$

The numerical procedure uses a pseudoinverse with a numerical floor; the equality in (62) assumes full column rank. Reshape $\boldsymbol{r}$ into the residual matrix $\mathbf{R} \in \mathbb{C}^{Q_i\times P_i}$ and let $\boldsymbol{c}_k$ denote the spatial steering vector of atom k. We estimate the beam-dependent disturbance variance using

$$d_k = P_i - \sum_{\ell\in\mathrm{supp}} \frac{|\boldsymbol{c}_k^H\boldsymbol{c}_\ell|^2}{Q_i^2}, \tag{63}$$

$$\widehat{\sigma}_{k,\mathrm{beam}}^2 = \frac{\|\boldsymbol{c}_k^H\mathbf{R}\|_2^2}{Q_i d_k} \quad (d_k \geq 3). \tag{64}$$

The pruning statistic is $D_k/\widehat{\sigma}_{k,\mathrm{beam}}^2$, set to zero for $d_k < 3$. Backward elimination proceeds by removing the least significant atom until all retained atoms exceed the prescribed threshold. Both the initial search and pruning thresholds are calibrated using noise-only range cells; in the comparison noise-only receiver, the support is retained using only the fixed search threshold. The adaptive rule accounts for residual power in the target beam, but its noise-only false-alarm calibration does not establish constant false-alarm performance in the presence of structured interference.

VII. Numerical results

We evaluate the proposed scheduling algorithms in terms of the collision energy, received interference power and target estimation performance. The numerical studies comprise (i) a verification of the finite-time collision law in an underloaded network, (ii) a comparison between predicted and simulated interference power, and (iii) a comparison of detection and false-discovery rates for the four schedulers and independent random assignment. We use the archived Monte Carlo experiments. Unless specified otherwise, the communication and physical conflict edges coincide and $\mathbf{F} = 0$; empirical performance with stale information is not assessed in these experiments.

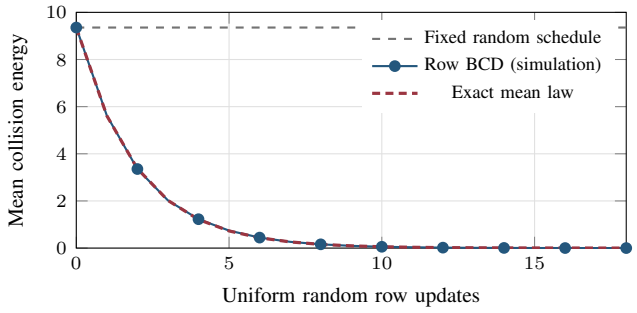

Fig. 8. Five-radar example, six physical edges, full communication, $M = 64$, $P_i = 10$, 20,000 independent initializations. The exact mean law uses the observed initial mean. A globally disjoint allocation has zero energy and is not drawn separately.

A. Scheduling in the absence of overloading

Consider the special case of an underloaded chirp pool with no unary penalty, $\lambda \mathbf{F} = 0$. Let $\mathcal{N}_i = \{j : A_{ij} > 0\}$, and suppose that

$$M \geq P_i + \sum_{j \in \mathcal{N}_i} P_j \quad \text{for every } i. \tag{65}$$

A sufficient condition for (65) is $\sum_i P_i \leq M$. Since the union of the neighbouring assignments has cardinality at most $\sum_{j \in \mathcal{N}_i} P_j$, there exist at least P_i slots of zero conditional cost. Every exact row update therefore removes all collisions incident on the updated vertex without introducing new collisions. A complete sweep yields a collision-free assignment. Under uniform independent random scan, an initial edge contribution remains unchanged until an endpoint is selected and is zero thereafter. Taking expectations edge by edge gives

$$\mathbb{E}[E_{\mathrm{col}}(\mathbf{X}^{(t)}) \mid \mathbf{X}^{(0)}] = (1 - 2/N)^t E_{\mathrm{col}}(\mathbf{X}^{(0)}) \tag{66}$$

for $N \geq 2$, with the initial energy at $t = 0$. The identity pertains to the available collision graph; it also holds for the physical collision graph when all physical edges are available. For nonzero fields, see the sufficient cost-margin condition in the supplementary material.

Figure 8 shows the mean collision energy over 20,000 independent initializations of the five-radar configuration in Fig. 1, with $M = 64$ and $P_i = 10$. The empirical initial mean of 9.3559 is close to the theoretical value $6P_i^2/M = 9.375$. The corresponding values after 4 and 18 random row updates are 1.2252 and 0.0009, respectively. The analytical curve in (66) is evaluated at the empirical initial mean without fitting the decay rate. A deterministic sweep eliminates all collisions in every trial; the same result is obtained in the archived complete-graph configurations with $(N, M) = (2, 32), (4, 64), (6, 64)$.

For the five-radar instance, the median and 95th-percentile execution times of a complete BCD sweep in the archived Python implementation are $164.6\,\mu$s and $248.9\,\mu$s, respectively, on an Intel host. Although both values are below the illustrative CPI duration of 1.28 ms, they exclude the communication, synchronization and guard overheads of the preparation protocol. They do not constitute an end-to-end latency comparison with Gaussian DRGS.

TABLE IV
INPUT-REFERRED SINR (DB): RATIO OF ENSEMBLE POWERS. $P = 10$, $N_s = 200$, INPUT SNR $= -15$ DB, $g = 1$. THE FORMULA IS (67); BCD IS MEASURED.

N	M	K	Random	Formula	BCD
2	32	2	−21.913	−21.949	−14.993
2	32	4	−24.506	−24.497	−15.013
4	64	2	−23.477	−23.407	−15.003
4	64	4	−26.100	−26.092	−14.998

B. Verification of the interference-power law

We consider complete physical and communication graphs with $P_i = 10$, K unit-power targets in each common range cell, and $g = 1$. The phase model is as specified in Section VI. For $N_s = 200$ and an input SNR of -15 dB, the predicted input-referred ratio of ensemble powers for independent random assignment is

$$-10 \log_{10}\left(10^{15/10} + 200K(N-1)P/M\right) \text{ dB}. \tag{67}$$

Equation (67) is obtained directly from (58). All configurations in Table IV satisfy (65). The simulated SINR is calculated from the ensemble means of the desired, cross-echo and noise powers, with the range-compression gain $10 \log_{10} N_s$ subtracted in dB. Over 2,000 paired trials per configuration, the maximum absolute discrepancy between the predicted and simulated random-assignment SINRs is 0.071 dB. In the BCD case, the absence of collisions gives the isolated-radar SNR to within Monte Carlo error. The comparison verifies the power calculation for the simulated model.

C. Range-cell detection: design and metrics

For joint angle–Doppler estimation we consider $K \in \{1, 2, 3, 6\}$ targets per range cell, with grid-aligned angles and velocities and independent unit-modulus complex amplitudes. Distinct target pairs are separated by at least 8 m/s in velocity or 0.25 in sine of angle. The virtual array is a uniform linear array with eight half-wavelength-spaced elements, rather than the sparse random array considered in [1]. The simulation parameters are listed in Table V. Since the range cell is prescribed, the evaluation pertains to two-dimensional estimation rather than end-to-end range–angle–Doppler estimation.

The four schedulers are initialized with the random assignment, which also serves as the fixed baseline. Common random numbers are used for the targets, phases and thermal noise in each paired trial. The computational budgets are $10N$ random row updates for row DRGS and row BCD, ten parallel rounds for Gaussian DRGS, and $5N$ two-row updates for paired BCD. The paired implementation samples from all radar pairs rather than from a fixed partition. These budgets equate nominal row updates, not elapsed time or communication overhead; the fixed-partition bound is not used for this implementation.

Let K, H and V denote the number of true targets, matched estimates and unmatched estimates in a range cell, respectively. The detection rate and false-discovery rate are defined as

$$P_D = \mathbb{E}[H/K], \qquad P_F = \mathbb{E}\left[\frac{V}{\max(1, H + V)}\right]. \tag{68}$$

TABLE V
RANGE-CELL DETECTION EXPERIMENTS.

Parameter	Value
Carrier; chirp duration	24 GHz; 40 μs
Fast-time samples; virtual channels	200; 8
Chirp quota; inverse temperature	10; $\beta = 4$
Interference	$g = 1$, common-cell targets, independent slot and path phases
Search grid	0° to 80° by 2°; -30 to 30 m/s by 1 m/s
Matching tolerances	2°, 1 m/s; one-to-one
Noise-only calibration	Nominal probability 0.02 per empty cell
Two-radar study	$M = 32$, 500 paired trials, both radars scored
Network study	$M = 10N$ or $16N$, 300 paired trials, up to four radars scored
Input SNR	Two radars: -20 to 0 dB in 5 dB steps; network: $-20, -15$ dB

The reported quantities are empirical averages over the evaluated cells. In particular, P_F is the mean false-discovery proportion, not an empty-cell false-alarm probability or the probability of any false estimate. The matching is one-to-one, and duplicate unmatched estimates are false discoveries. The false-discovery proportion is zero for an empty estimate set; consequently P_F must be interpreted jointly with P_D. Since cells within a trial share the scene and assignment, the archived cell-wise standard errors are not treated as independent-trial confidence intervals.

D. Detection performance as a function of target density

Figures 9(a)–(b) illustrate the effect of the number of targets per range cell at -15 dB input SNR. For $K = 1$, independent random assignment and row BCD yield detection rates of 89.7% and 99.9%, respectively. For $K = 6$, the corresponding rates are 42.42% and 93.02%, and the false-discovery rates are 9.28% and 4.13% (Table VI). This corresponds to a 50.60 percentage-point increase in detection and a 5.15 percentage-point reduction in false discoveries. Paired BCD yields 94.08% detection and 3.90% false discoveries. Both BCD schedules are collision-free and achieve the isolated-radar input-referred SINR of -15 dB, compared with -25.66 dB for random assignment. The residual difference in recovery between the two BCD algorithms is attributable to their different measurement matrices and the numerical receiver, rather than to cross echoes.

With the thermal-noise-only support-selection rule at the same six-target operating point, the detection and false-discovery rates are 75.28% and 29.29% under random assignment, and 93.62% and 5.81% under row BCD. Relative to the fixed-threshold receiver, the adaptive pruning rule trades a lower false-discovery rate for a lower detection rate in the presence of interference. This comparison demonstrates that the relation between interference power and detection depends on the support-selection rule.

Figures 9(c)–(d) show the scaling with network size for complete graphs. With $M = 10N$, the expected number of overlapping transmissions per selected chirp under independent random assignment is $(N-1)/N$. At $N = 12$ and $K = 6$, the detection rates for random assignment, row DRGS, Gaussian DRGS, row BCD and paired BCD are 4.18%, 58.44%, 5.13%, 86.78% and 86.61%, respectively, with false-discovery rates of 5.62%, 7.10%, 5.94%, 4.78% and 4.77%. The performance of Gaussian DRGS at the specified temperature and iteration budget is substantially inferior to that of row DRGS in this configuration; the invariance of its Gibbs distribution is not a finite-time performance guarantee. The comparison does not optimize the temperature or budget. Variation in the BCD detection rates with M, despite zero collisions, reflects the dependence of the Doppler dictionary and the recovery algorithm on the selected chirp times.

TABLE VI
TWO RADARS, SIX TARGETS PER CELL, $M = 32$, $P = 10$, -15 DB INPUT SNR, ADAPTIVE OMP. MEANS OVER 1,000 SCORED CELLS. SINR IS THE MEAN CONDITIONAL EFFECTIVE SINR IN DB.

Schedule	SINR (dB)	P_D (%)	P_F (%)
Random	-25.66	42.42	9.28
Row DRGS	-15.86	91.65	4.34
Gaussian DRGS	-16.03	91.83	4.67
Row BCD	-15.00	93.02	4.13
Paired BCD	-15.00	94.08	3.90

E. Full-waveform diagnostic and phase sensitivity

We also consider the full-waveform receiver variants OMP+Binary and OMP+Range-OMP in a separate archived two-radar sensitivity experiment. The configuration consists of two 10-of-32 radars and three targets at ranges of 40, 60 and 80 m, with 200 fast-time samples, a carrier of 24 GHz and a bandwidth of 250 MHz. The receiver first estimates candidate ranges before joint angle–Doppler estimation. The principal phase-incoherent experiment uses a cross-to-own echo power ratio of 25 dB and 1,000 paired trials. To examine sensitivity to range admission, the bistatic return is displaced in fast time. These experiments differ from the preceding range-cell studies in interference strength, range processing and the number of admitted echoes.

Figure 10 compares phase-locked, phase-incoherent and range-displaced cross echoes under fixed random scheduling. A phase-locked on-cell copy modifies the coherent target amplitude, whereas independent slot phases perturb the slow-time steering sequence. Range displacement additionally changes the range-search statistic. The difference between the resulting curves supports the use of an explicit phase and admission model in receiver evaluation. It does not establish the physical validity of the assumed oscillator or bistatic propagation statistics.

F. Scope of the numerical evaluation

The interference-power identity is derived under explicit phase and range-admission assumptions and accommodates heterogeneous coupling coefficients. The present simulations, however, predominantly use equal-strength common-cell returns. A more comprehensive waveform model would incorporate geometry-dependent bistatic delay and Doppler, direct coupling, chirp-slope differences and clock offsets. The

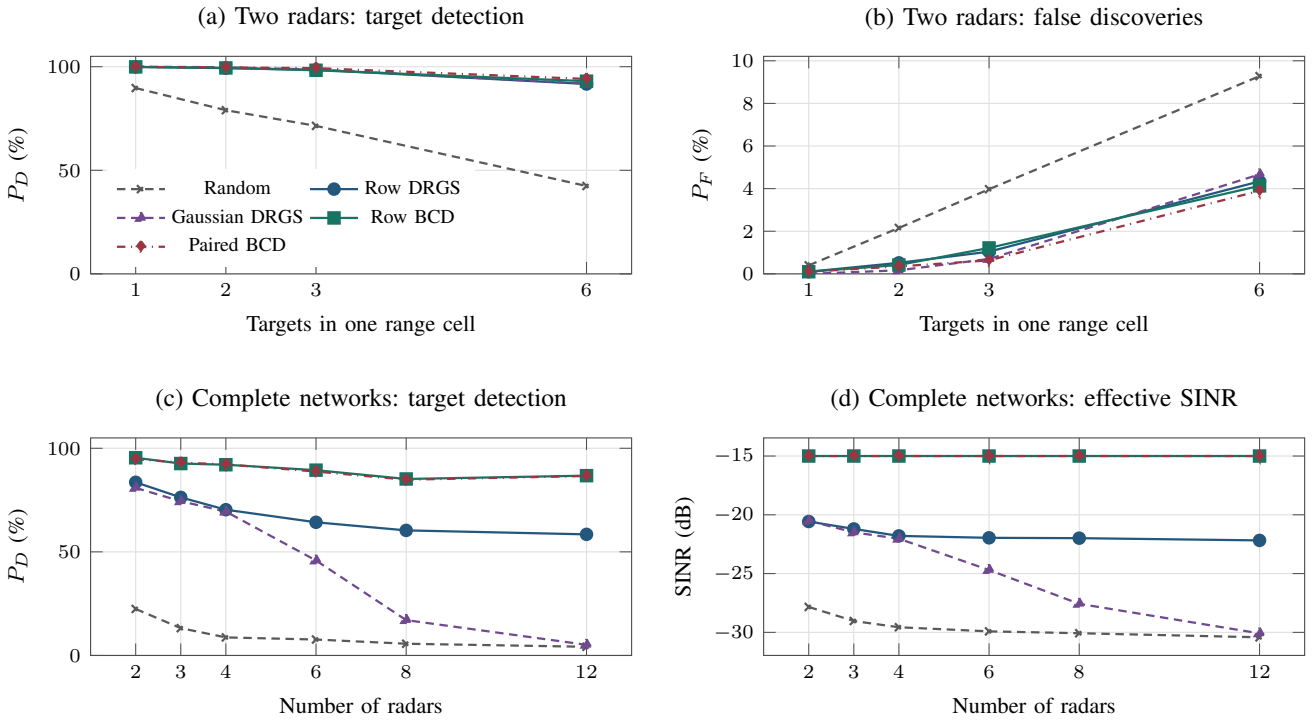

Fig. 9. Range-cell cross-echo simulations at input SNR -15 dB and $g = 1$, using the adaptive OMP receiver. (a)–(b) Two radars with $M = 32$ and 500 paired trials. (c)–(d) Complete graphs with $M = 10N$ and six common targets per cell; 300 paired trials per configuration, up to four victims scored. The plotted effective SINR is the mean of the conditional model values in dB, not the ratio of ensemble powers used in Table IV. Curves reproduce the archived CSV data.

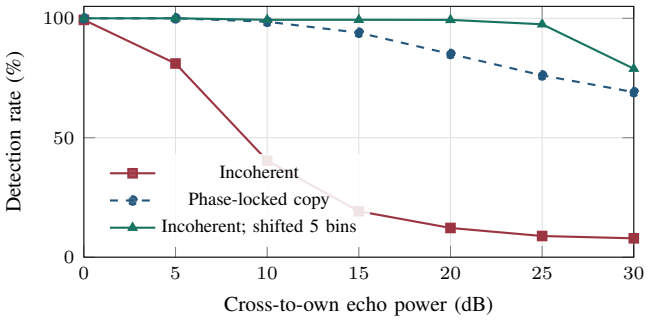

Fig. 10. Full-waveform diagnostic: OMP+Binary, random scheduling, -10 dB input SNR, 500 paired trials per case. The cross-echo phase convention and its range admission change the detection response at the same specified cross-to-own echo power. These are earlier sensitivity experiments, separate from Fig. 9.

adaptive receiver has not been shown to control the false-alarm probability under arbitrary structured interference. Evaluation of the AoI-dependent predictions and end-to-end preparation latency under intermittent communication remains open. When the physical and available graphs differ, elimination of the available-edge collision energy does not imply zero received interference. Finally, the sufficient conditions in Section V are not quantitative predictions of P_D and P_F; the numerical results are conditional on the simulated receiver and scene distributions.

VIII. Conclusion

We have investigated distributed chirp selection for MIMO–FMCW radar networks subject to fixed measurement budgets and incomplete information about neighbouring transmissions. The dual-graph formulation separates physical interference from timely communication and incorporates an AoI-dependent penalty for unavailable assignments. Exact single-radar and two-radar optimization and two Gibbs samplers provide alternative implementations with different computation and communication requirements. The objective-gap, mixing, and perturbation bounds characterize their performance on a fixed scheduling instance and the effect of changes between instances. Under explicit receiver assumptions, the physical chirp overlap controls target–interference coherence; exact and stable recovery additionally depend on the remaining correlations in the augmented dictionary. For the phase-incoherent cross-echo model, the network collision energy is proportional to the expected interference power under homogeneous coupling. The simulations demonstrate substantial improvements in joint angle–Doppler detection when interference and target congestion are appreciable, while also showing that SINR alone does not characterize receiver performance. Evaluation with measured propagation channels, imperfect synchronization, and intermittent control links is an important direction for further work.

References

[1] C. Rai, H. Singh, and A. Chattopadhyay, "Multi-target range, Doppler and angle estimation in MIMO-FMCW radar with limited measurements," *IEEE Transactions on Signal Processing*, vol. 73, pp. 3762–3778, 2025.

[2] E. Tohidi, M. Coutino, S. P. Chepuri, H. Behroozi, M. M. Nayebi, and G. Leus, "Sparse antenna and pulse placement for colocated MIMO radar," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 579–593, 2019.

[3] M. Goppelt, H.-L. Blöcher, and W. Menzel, "Automotive radar—investigation of mutual interference mechanisms," *Advances in Radio Science*, vol. 8, pp. 55–60, 2010.

[4] C. Aydogdu, M. F. Keskin, N. Garcia, H. Wymeersch, and D. W. Bliss, "RadChat: Spectrum sharing for automotive radar interference mitigation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 1, pp. 416–429, 2021.

[5] C. Aydogdu, H. Wymeersch, O. Eriksson, H. Herbertsson, and M. Rydström, "Synchronization-free RadChat for automotive radar interference mitigation," *Sustainability*, vol. 13, no. 12, p. 6891, 2021.

[6] S. Roudiere, V. Martinez, P. Maréchal, and D. Delahaye, "Metaheuristic for optimal dynamic K-coloring application on band sharing for automotive radars," *Sensors*, vol. 23, no. 12, p. 5765, 2023.

[7] Y. Pan, J. Li, L. Xu, S. Sun, and Q. Zhu, "Decentralized no-regret frequency-time scheduling for FMCW radar interference avoidance," arXiv:2512.24619v2, 2026.

[8] T. Moon, J. Park, and S. Kim, "BlueFMCW: Random frequency hopping radar for mitigation of interference and spoofing," arXiv:2008.00624, 2020.

[9] A. Correas-Serrano and M. A. Gonzalez-Huici, "Sparse reconstruction of chirplets for automotive FMCW radar interference mitigation," in *IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, 2019, pp. 1–4.

[10] J. Wang, M. Ding, and A. Yarovoy, "Interference mitigation for FMCW radar with sparse and low-rank hankel matrix decomposition," *IEEE Transactions on Signal Processing*, vol. 70, pp. 822–834, 2022.

[11] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Foundations and Trends in Machine Learning*, vol. 1, no. 1–2, pp. 1–305, 2008.

[12] X.-H. Chen, A. P. Dempster, and J. S. Liu, "Weighted finite population sampling to maximize entropy," *Biometrika*, vol. 81, no. 3, pp. 457–469, 1994.

[13] S. X. Chen and J. S. Liu, "Statistical applications of the poisson-binomial and conditional bernoulli distributions," *Statistica Sinica*, vol. 7, no. 4, pp. 875–892, 1997.

[14] J. Martens and I. Sutskever, "Parallelizable sampling of Markov random fields," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, vol. 9, 2010, pp. 517–524.

[15] C. C. Margossian, C. Zhong, and S. Mukherjee, "Efficient sampling for Ising and Potts models using auxiliary gaussian variables," arXiv:2110.10801v2, 2026.

[16] M. Fischer and M. Ghaffari, "A simple parallel and distributed sampling technique: Local glauber dynamics," in *32nd International Symposium on Distributed Computing (DISC)*, ser. Leibniz International Proceedings in Informatics, vol. 121, 2018, pp. 26:1–26:11.

[17] C. D. Sa, K. Olukotun, and C. Ré, "Ensuring rapid mixing and low bias for asynchronous Gibbs sampling," in *Proceedings of the 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 48, 2016, pp. 1567–1576.

[18] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proceedings IEEE INFOCOM*, 2012, pp. 2731–2735.

[19] A. M. Bedewy, Y. Sun, and N. B. Shroff, "The age of information in multihop networks," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1248–1257, 2019.

[20] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.

[21] C. Studer, P. Kuppinger, G. Pope, and H. Bölcskei, "Recovery of sparsely corrupted signals," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3115–3130, 2012.

[22] C. Studer and R. G. Baraniuk, "Stable restoration and separation of approximately sparse signals," *Applied and Computational Harmonic Analysis*, vol. 37, no. 1, pp. 12–35, 2014.

[23] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of Optimization Theory and Applications*, vol. 109, no. 3, pp. 475–494, 2001.

[24] J. Borcea, P. Brändén, and T. M. Liggett, "Negative dependence and the geometry of polynomials," *Journal of the American Mathematical Society*, vol. 22, no. 2, pp. 521–567, 2009.

[25] E. A. Carlen, D. Cordero-Erausquin, and E. H. Lieb, "Asymmetric covariance estimates of Brascamp–Lieb type and related inequalities for log-concave measures," Author manuscript, 2011.

[26] M. Dyer, L. A. Goldberg, and M. Jerrum, "Matrix norms and rapid mixing for spin systems," *The Annals of Applied Probability*, vol. 19, no. 1, pp. 71–107, 2009.

[27] G. Pope, A. Bracher, and C. Studer, "Probabilistic recovery guarantees for sparsely corrupted signals," *IEEE Transactions on Information Theory*, vol. 59, no. 5, pp. 3104–3116, 2013.