Finite-Mask Gaussian-Process Reconstruction on a Periodic Sensor Ring: Mask-Geometry Dependence, Fourier-Mode Coupling, and Posterior Trace

Jun Tsuzurugi*

Department of Information Science and Engineering,
Okayama University of Science, Okayama 700-0005, Japan

(Received June 22, 2026; revised July 18, 2026; accepted August 10, 2026; published online September 9, 2026)

We analyze how an arbitrary observation mask on a finite periodic ring of equally spaced sensors affects the posterior covariance, Fourier-mode coupling, and normalized posterior trace in Gaussian-process reconstruction. Under a rotationally stationary prior and homogeneous independent measurement noise, complete observation gives independent scalar posterior formulas for the Fourier modes. For an arbitrary mask, however, the matrix $Q_M = FD_MF^*$ is generally non-diagonal; its off-diagonal entries are finite Fourier components of the realized mask and couple modes in the posterior precision. Consequently, masks with the same unavailable-channel fraction can have different normalized posterior traces because their geometries differ. A dimensionless 64-channel synthetic benchmark illustrates this finite-matrix effect. A circumferential array of equally spaced wall-mounted microphones at a fixed axial station of a circular fan or compressor duct provides one concrete mechanical-engineering interpretation: failed, saturated, corrupted, or dropped-out channels form the observation mask. The analysis is a finite-dimensional reference calculation under the stated rotational-stationarity and common-noise assumptions, not a performance claim for a nonuniform or unequally instrumented operating duct.

Author manuscript of J. Phys. Soc. Jpn. **95**, 104002 (2026).
`https://doi.org/10.7566/JPSJ.95.104002`

1. Introduction

Partially observed finite data are common in Bayesian reconstruction. Sensor records may contain dropouts, masked signal or lattice data may be available only at selected sites, and spatial measurements may be restricted by accessible sampling locations. In such situations, complete observation of every lattice site is an idealization.

Gaussian-process reconstruction is useful because it gives both a posterior mean and a posterior covariance.[1–3] The posterior covariance measures the uncertainty that remains after conditioning on the observation mask, and it connects reconstruction reliability with sampling

*juntuzu@ous.ac.jp

design and missing-data assessment. The practical question is not only how to reconstruct unobserved values, but how reliable that reconstruction is for the actual measurement pattern.

The observed-site count gives only a partial answer: regularly distributed and clustered observations can leave different posterior uncertainties even at the same missing rate. Thus the missing rate is a coarse summary, and the posterior formulation must retain the geometry that enters Σ_M together with the count.

A finite periodic lattice gives a minimal controlled setting in which this geometry dependence can be separated from boundary effects. With a translation-invariant Gaussian prior, complete observation is exactly diagonal in the Fourier basis and gives independent scalar Fourier-mode posterior formulas. Once an arbitrary observation mask is inserted, the likelihood precision is no longer diagonal in that basis; the mask couples Fourier modes and changes the posterior uncertainty even though the prior remains Fourier diagonal.

This paper uses the standard finite-dimensional linear-Gaussian posterior identity as a controlled starting point rather than as a claimed new conditioning theorem. The contribution is to embed that identity in the periodic-lattice Fourier representation with an arbitrary observation mask and to isolate D_M, or equivalently $Q_M = FD_MF^*$, as an explicit finite matrix. Under complete observation, $D_M = I_L$ and the posterior reduces to independent scalar Fourier modes. Under a sparse mask, the prior remains diagonal as $K = F^*\Lambda F$, but the likelihood contribution $b^{-2}Q_M$ generally couples those modes. The matched error is then the normalized posterior trace $L^{-1} \operatorname{Tr} \Sigma_M$, replacing the independent scalar-mode error sum of the complete-observation case.

The finite periodic model serves as a controlled benchmark for partially observed lattice or signal data. It keeps the Gaussian-process prior and the observation geometry separated while remaining exactly finite-dimensional. Connections with spatial-statistical kriging, Wiener filtering, and masked finite-signal analysis are made explicit below.

1.1 Engineering motivation and circumferential microphone-array interpretation

Circumferential arrays of wall-mounted microphones are used to resolve circumferential, or spinning, acoustic modes in circular fan and compressor ducts. Joppa applied a spatial Fourier transform to an equally spaced array to measure spinning modes in a turbofan inlet;[4] Wang et al. used an equidistant circular array for in-duct circumferential-mode measurements of an axial fan/compressor;[5] and Wang et al. reconstructed a full cross-spectral matrix from incomplete circumferential measurements before identifying broadband acoustic modes.[6] Such measurements provide modal information for characterizing duct sound fields and turbomachinery-noise propagation, and hence are important for noise-source characterization and noise assessment in mechanical engineering and aeroacoustics.

Microphone failure, saturation, recording corruption, or communication dropout produces incomplete circumferential samples. The resulting task is therefore not only to reconstruct pressure values at unavailable channels, but also to quantify the uncertainty conditional on the realized arrangement of those channels. This is the engineering connection to the arbitrary finite mask analyzed below.

A concrete engineering interpretation of the finite periodic sensor ring is a circumferential microphone array installed at one fixed axial station of an ideal circular fan or compressor duct. Let L nominally identical wall-mounted microphones be equally spaced at angular positions

$$\theta_j = \frac{2\pi j}{L}, \qquad j = 0, 1, \ldots, L-1. \tag{1}$$

The discrete torus $\mathbb{T}_L$ then represents the angular ordering of the microphones, including the periodic adjacency between channels $L-1$ and zero. For a simultaneous measurement, let $S_j = S(\theta_j)$ denote the mean-removed real-valued acoustic-pressure fluctuation at the j-th microphone. An available channel records

$$t_j = S_j + \xi_j, \tag{2}$$

where the present model assumes independent, identically distributed measurement noise with variance b^2.

The translation-invariant covariance used below follows from explicit second-order assumptions. Suppose that the duct and the nominal sensor layout are rotationally uniform and that, after mean removal, the probability law of the pressure fluctuation is invariant under a common circumferential rotation. Then its covariance depends only on angular separation:

$$\begin{aligned} K_{jj'} &= \mathrm{Cov}\big[S(\theta_j), S(\theta_{j'})\big] \\ &= C\big((\theta_j - \theta_{j'}) \bmod 2\pi\big) \\ &= K_{(j-j') \bmod L}. \end{aligned} \tag{3}$$

Thus K is a circulant matrix. This assumption does not mean that each realized pressure waveform is rotationally symmetric. It means that the mean-removed probability distribution, or at least its second-order covariance, is invariant under a common angular shift.

The same conclusion has a circumferential-mode interpretation. Write the mean-removed pressure field as

$$S(\theta) = \sum_{n\in\mathbb{Z}} a_n e^{in\theta}, \qquad \mathbb{E}[a_n] = 0, \tag{4}$$

with $a_{-n} = a_n^*$ for a real field. If distinct circumferential-mode coefficients are uncorrelated and

$$\mathbb{E}[a_n a_{n'}^*] = P_n \delta_{nn'}, \qquad P_n \geq 0, \tag{5}$$

then

$$K_{jj'} = \sum_{n\in\mathbb{Z}} P_n \exp\big\{in(\theta_j - \theta_{j'})\big\}\,. \tag{6}$$

For the L equally spaced samples, continuous circumferential orders that are congruent modulo L are aliased into the same discrete mode. With the unitary Fourier matrix defined in Sect. 2,

$$FKF^* = \operatorname{diag}(\lambda_0, \ldots, \lambda_{L-1}), \qquad \lambda_q = L \sum_{r\in\mathbb{Z}} P_{q+rL} \geq 0. \tag{7}$$

Fourier diagonalization is therefore not merely a computational assumption in this idealized measurement system; it is the discrete circumferential acoustic mode decomposition of a rotationally stationary covariance.

The binary mask has the direct channel interpretation $M_j = 1$ when microphone j provides a usable record and $M_j = 0$ when that channel is unavailable because of microphone failure, saturation, a corrupted recording, or a communication dropout. The row-selection matrix H retains only the usable channels, and $D_M = H^T H = \operatorname{diag}(M_0, \ldots, M_{L-1})$ records their circumferential acquisition geometry. Unlike K, an arbitrary D_M is not rotationally invariant. Its Fourier representation has entries

$$(Q_M)_{qq'} = (FD_M F^*)_{qq'} = \frac{1}{L}\sum_{j=0}^{L-1} M_j e^{-i(\kappa_q - \kappa_{q'})j}, \tag{8}$$

which are generally nonzero for $q \neq q'$. Hence the off-diagonal elements of Q_M mix circumferential acoustic modes in the posterior precision $\Lambda^{-1} + b^{-2}Q_M$. This is the direct physical interpretation of the mask-induced Fourier-mode coupling studied in this paper.

Table I summarizes the correspondence between the mathematical notation and the microphone-array problem.

The posterior mean reconstructs the full pressure vector, including the unavailable channel positions, whereas Σ_M and $E(M)$ quantify the uncertainty that remains for the realized microphone configuration.

The present study does not analyze engine or compressor data. It is an idealized finite-dimensional reference model for this measurement configuration and does not claim reconstruction performance or model fit for an operating machine. Strong support-structure effects, circumferentially nonuniform swirl, nonuniform wall conditions, or other azimuthal inhomogeneities can make K non-circulant, while sensor-specific noise can invalidate the common-noise model $b^2 I_m$. Such systems require a nonstationary covariance or a channel-dependent likelihood and lie outside the present scope. The narrower purpose here is to determine exactly how a realized pattern of unavailable circumferential channels enters Gaussian reconstruction and its posterior uncertainty in the ideal rotationally stationary reference case.

1.2 *Related work and distinction from neighboring methods*

Circumferential microphone measurements and spatial Fourier analysis of duct acoustic modes provide the experimental context for the sensor-ring interpretation. Joppa analyzed equally spaced, wall-mounted circumferential microphone data by a spatial Fourier transform

Table I. Correspondence between the finite periodic model and a circumferential microphone array at one axial station.

Mathematical quantity	Circumferential-array interpretation
L, $\mathbb{T}_L$	Number of nominal channels and their cyclic angular ordering.
j, $\theta_j = 2\pi j/L$	Microphone index and its circumferential angle.
$S_j = S(\theta_j)$	True mean-removed acoustic-pressure fluctuation at channel j.
t_j, $j \in \Omega$	Noisy measurement returned by an available channel.
$b^2 I_m$	Homogeneous independent measurement-noise covariance for the m available channels.
K	Circumferential pressure covariance assumed for the nominal rotationally stationary field.
F, $\widetilde{S}_n$, λ_n	Circumferential discrete-mode transform, modal pressure coefficient, and prior modal variance.
M_j, Ω, m	Channel-availability indicator, set of available channels, and their number.
H, $D_M = H^T H$	Usable-channel selector and full-ring diagonal availability matrix.
$Q_M = F D_M F^*$	Missing-channel-induced mixing matrix between circumferential modes.
Σ_M	Posterior error covariance for the reconstructed full-ring pressure vector.
$G_{\max}$, $2\pi G_{\max}/L$	Largest cyclic gap in channel-index units and the corresponding largest angular gap between available microphones.
$p(M) = 1 - m/L$	Realized fraction of unavailable channels.
$E(M) = L^{-1} \operatorname{Tr} \Sigma_M$	Posterior variance averaged over the entire circumference.

to measure spinning modes in a turbofan inlet.[4] Wang et al. used an equidistant circular microphone array for in-duct circumferential-mode measurements of an axial fan/compressor.[5] For nonsynchronous measurements with a wall-mounted circular array, Wang et al. reconstructed a full cross-spectral matrix from incomplete circumferential measurements and then identified broadband acoustic modes.[6] These works primarily address acoustic-mode measurement or identification, and the last also addresses cross-spectral matrix completion. The present paper instead conditions a Gaussian field on an arbitrary fixed binary pattern of unavailable channels, evaluates the resulting posterior covariance and its normalized trace exactly in finite dimensions, and isolates $Q_M = FD_MF^*$ as the mechanism by which two masks with the same unavailable fraction can yield different uncertainties. It neither proposes a duct-mode identification procedure nor a cross-spectral matrix-completion algorithm.

In the terminology of spatial statistics, the same Gaussian conditioning formula underlies kriging for partially observed spatial fields.[7,8] In the terminology of Gaussian-process regression, the same finite Gaussian conditioning formulas are commonly used.[1–3] The present paper uses these formulas as a finite periodic-lattice benchmark with an arbitrary realized mask, and it exposes the resulting Fourier-mode coupling. The distinction is therefore not a new conditioning identity, but the exact finite representation of the mask term and its posterior-trace consequence in the periodic Fourier basis.

In signal-processing language, the complete-observation periodic case is the finite Fourier-domain analogue of a scalar Wiener filter.[9] When every site is observed and the prior covariance is translation invariant, each Fourier mode is shrunk independently by a scalar factor. With an arbitrary mask, however, the observation operator breaks translation symmetry. The likelihood contribution in Fourier space is then $b^{-2}Q_M$, with

$$Q_M = FD_MF^*, \tag{9}$$

and the posterior precision becomes $\Lambda^{-1} + b^{-2}Q_M$. Since Q_M is generally non-diagonal, the posterior is a coupled finite-mode problem rather than a collection of independent scalar Wiener shrinkages.

This is also the finite-lattice version of the familiar Fourier effect of masking a signal. Multiplication by a real-space mask corresponds in Fourier space to convolution by the Fourier components of that mask, often described as spectral mixing or leakage in finite-signal analysis.[10] In the present Bayesian formulation this effect appears directly inside the posterior precision through $Q_M = FD_MF^*$, whose off-diagonal entries are the finite Fourier components of the realized mask.

The formulation is also distinct from compressed sensing or sparse-recovery theory.[11] The latent signal is modeled by a Gaussian prior, and the evaluated quantities are the posterior covariance and the normalized posterior trace. The aim is to quantify Bayes uncertainty for each finite realized mask.

Scope. The scope of the paper is deliberately finite and mask-conditioned. We do not derive an asymptotic sampling theorem or an off-lattice interpolation theorem. Instead, we focus on the exactly computable posterior covariance for each realized finite observation mask and on the resulting dependence of posterior uncertainty on mask geometry. Random masks and random inter-observation intervals are used as generation and reporting layers for realized finite masks.

Notation bridge to the complete-observation formulation. In the previous complete-observation formulation,[12] the notation $N, A, B, C = A + B$ was used. Table II summarizes how these symbols are carried into the present sparse-observation formulation. The previous formulation is preserved and recovered exactly as the boundary case $H = I_L$, $D_M = I_L$, and $m = L$. Below, after this bridge, that complete-observation work is called the previous formulation.

Table II. Notation bridge between the previous complete-observation formulation and the present sparse-observation formulation. The complete-observation covariance $K + b^2 I_L$ is recovered in the boundary case $H = I_L$, $D_M = H^T H = I_L$, and $m = L$.

Previous complete-observation formulation	Current sparse-observation formulation
Lattice size N	Full latent lattice size L; the number of observed sites is $m = \lvert\Omega\rvert \le L$.
Signal covariance A	Full-lattice prior covariance K.
Noise covariance $B = b^2 I_N$	Observation-space noise covariance $b^2 I_m$.
Covariance sum $C = A + B$	Observation covariance $HKH^T + b^2 I_m$.
Complete observation	Boundary case $H = I_L$, $D_M = H^T H = I_L$, and $m = L$, giving $HKH^T + b^2 I_m = K + b^2 I_L$.

The contribution is threefold. First, the paper fixes a finite-mask posterior notation on the same finite periodic lattice as the complete-observation theory, without claiming a new Gaussian-conditioning identity. Second, it identifies the mask term $Q_M = F D_M F^*$ as the object that is generally non-diagonal in the Fourier basis; its off-diagonal entries are finite Fourier components of the realized mask and convert the scalar Fourier-mode posterior into a coupled finite-mode posterior. Third, it uses the posterior trace to quantify the dependence of reconstruction uncertainty on realized mask geometry, thereby making explicit why the missing rate should be supplemented by mask-geometry information.

2. One-dimensional Periodic Lattice and Admissible Covariance

Definition 1 (One-dimensional finite torus)**.** Let

$$\mathbb{T}_L = \{0, 1, \ldots, L-1\} \tag{10}$$

be the one-dimensional discrete torus. The Fourier index is denoted by $n \in \mathbb{T}_L$, and

$$\kappa_n = \frac{2\pi n}{L}. \tag{11}$$

For a vector $S = (S_0, \dots, S_{L-1})^T$, we use the unitary discrete Fourier transform

$$\widetilde{S}_n = \frac{1}{\sqrt{L}} \sum_{j=0}^{L-1} S_j e^{-i\kappa_n j}, \qquad S_j = \frac{1}{\sqrt{L}} \sum_{n=0}^{L-1} \widetilde{S}_n e^{i\kappa_n j}. \tag{12}$$

Definition 2 (Admissible periodic covariance kernel)**.** A translation-invariant periodic kernel K_r, $r \in \mathbb{T}_L$, is admissible if its finite Fourier eigenvalues

$$\lambda_n = \sum_{r=0}^{L-1} K_r e^{-i\kappa_n r} \tag{13}$$

satisfy

$$\lambda_n \geq 0 \qquad (n \in \mathbb{T}_L). \tag{14}$$

For real fields we also assume $K_{-r} = K_r$.

The covariance matrix is

$$K_{ij} = K_{i-j}, \tag{15}$$

where all indices are understood on the finite torus. It is diagonalized as

$$K = F^* \Lambda F, \qquad \Lambda = \mathrm{diag}(\lambda_0, \dots, \lambda_{L-1}), \tag{16}$$

where F is the unitary Fourier matrix.

2.1 Minimum-image Gaussian example

For comparison with the previous periodic formulations, one may use the minimum-image Gaussian

$$K_r^{\min} = a^2 \exp\left[-\frac{\widetilde{\delta}(r)^2}{2\ell^2}\right], \qquad \widetilde{\delta}(r) = \min\{r, L-r\}. \tag{17}$$

Its finite eigenvalues are

$$\lambda_n^{\min} = a^2 \sum_{r=0}^{L-1} \exp\left[-\frac{\widetilde{\delta}(r)^2}{2\ell^2}\right] e^{-i\kappa_n r}. \tag{18}$$

This kernel is used in parameter ranges satisfying

$$\min_n \lambda_n^{\min} \geq 0. \tag{19}$$

3. Sparse-observation Model

Let S_j be the latent zero-mean Gaussian field on the full lattice:

$$S \sim \mathcal{N}(0, K). \tag{20}$$

The sparse-observation formulation separates three layers. The latent field S is always an L-dimensional vector on the full lattice. The observed data t are an m-dimensional vector in observation space. The diagonal matrix D_M is an $L \times L$ full-lattice bookkeeping device for the mask, whereas H is the actual $m \times L$ observation operator that maps the latent lattice field to observation space.

Sparse observations are represented by a mask vector

$$M_j \in \{0,1\}, \qquad j \in \mathbb{T}_L. \tag{21}$$

Here and below, $M = (M_0, \ldots, M_{L-1})$ is the binary mask vector, whereas D_M is the diagonal matrix generated from that vector. We do not use M itself as a matrix. Let

$$\Omega = \{j \in \mathbb{T}_L : M_j = 1\}, \qquad m = |\Omega|. \tag{22}$$

Definition 3 (Arbitrary finite observation mask)**.** A sparse-observation mask is a binary vector

$$M = (M_0, \ldots, M_{L-1}), \qquad M_j \in \{0,1\}. \tag{23}$$

The observed set is

$$\Omega = \{j \in \mathbb{T}_L : M_j = 1\}. \tag{24}$$

In the finite-dimensional formulation, no deterministic condition is imposed on the positions of the observed sites, on the lengths of missing intervals, or on the maximum gap.

Define the observation operator $H \in \mathbb{R}^{m\times L}$ by selecting the rows of the identity matrix corresponding to Ω. The observation model is

$$t = HS + \xi, \qquad \xi \sim \mathcal{N}(0, b^2 I_m). \tag{25}$$

Equivalently, in full-lattice notation,

$$t_j = S_j + \xi_j \qquad (j \in \Omega), \tag{26}$$

where t_j denotes the component of the observation-space vector t associated with the observed site $j \in \Omega$. Thus t is an m-dimensional observation-space vector indexed by the observed set Ω, or by any fixed ordering of that set.

The diagonal mask matrix on the full lattice is

$$D_M = H^T H = \mathrm{diag}(M_0, \ldots, M_{L-1}). \tag{27}$$

Figure 1 illustrates the finite mask geometry as two 64-channel circumferential microphone-array configurations. In the derivation, the realized mask enters the posterior formula through the observation operator H and the full-lattice mask matrix $D_M = H^T H$.

4. Statistical-mechanical Posterior Formulation

The prior density is

$$P(S) = \frac{1}{Z_S} \exp\left[-\frac{1}{2} S^T K^{-1} S\right], \tag{28}$$

assuming that K is strictly positive definite. Appendix A records the zero-eigenvalue convention used when this assumption is relaxed.

The sparse-observation likelihood is

$$P(t|S) = \frac{1}{(2\pi b^2)^{m/2}} \exp\left[-\frac{1}{2b^2}\|t - HS\|^2\right]. \tag{29}$$

Therefore, up to a term independent of S, the conditional Hamiltonian is

$$H(S|t) = \frac{1}{2}S^T K^{-1} S + \frac{1}{2b^2}\|t - HS\|^2. \tag{30}$$

The core object in this section is the exact finite-dimensional Gaussian posterior conditioned on one realized observation mask; this is the finite-mask continuation of the previous complete-observation posterior. For a fixed realized finite mask, all observation geometry enters the quadratic Hamiltonian through the row-selection operator H, or equivalently through the full-lattice mask matrix $D_M = H^T H$.

In this finite Gaussian model, the posterior precision matrix is the Hessian of $H(S|t)$ with respect to the latent field S. Thus the statistical-mechanical calculation reduces to completing the square in a quadratic Hamiltonian. The important finite-mask change is not the conditioning identity itself, but the replacement of the complete-observation likelihood precision $b^{-2}I_L$ by the mask precision $b^{-2}D_M$. In Fourier space this term becomes $b^{-2}FD_MF^*$, and this is the finite-matrix origin of the Fourier-mode coupling appearing below.

The theorem is stated not as a new Gaussian-conditioning theorem, but as the finite-mask notation used throughout the paper. Its role is to display, in the same statement, the latent-space posterior covariance, the observation-space Woodbury form, and the Fourier-space precision containing the mask matrix Q_M. The underlying conditioning step is the standard finite linear-Gaussian identity;[2,3,13] the notation below isolates the mask-dependent Fourier mechanism and the trace quantity $L^{-1}\operatorname{Tr}\Sigma_M$. The theorem is stated for positive-definite K, so that K^{-1} is an ordinary inverse. The positive-semidefinite convention is a technical support convention and is collected in Appendix A.

Theorem 1 (Finite-mask posterior and coupled Fourier representation)**.** *Let $S \sim \mathcal{N}(0, K)$ be a Gaussian field on $\mathbb{T}_L$, and let*

$$t = HS + \xi, \qquad \xi \sim \mathcal{N}(0, b^2 I_m),$$

where H selects the observed sites specified by a fixed realized finite mask M. If K is positive definite and $b^2 > 0$, then the posterior distribution $P(S|t)$ is Gaussian with covariance

$$\Sigma_M = \left(K^{-1} + b^{-2}D_M\right)^{-1} \tag{31}$$

and mean

$$\mu_M = \Sigma_M b^{-2} H^T t. \tag{32}$$

Equivalently, by the Woodbury identity, the same posterior has the observation-space imple-

mentation form with an $m \times m$ inverse,

$$\mu_M = KH^T(HKH^T + b^2 I_m)^{-1} t \tag{33}$$

and

$$\Sigma_M = K - KH^T(HKH^T + b^2 I_m)^{-1} HK. \tag{34}$$

Writing $\widetilde{\Sigma}_M = F\Sigma_M F^*$, the Fourier-space posterior precision is

$$\widetilde{\Sigma}_M^{-1} = \Lambda^{-1} + b^{-2} Q_M, \qquad Q_M = FD_M F^*. \tag{35}$$

Thus the finite-mask posterior is scalar-mode diagonal exactly in cases where the mask term Q_M is diagonal.

Proof. Because K is positive definite in the theorem, the conditional Hamiltonian is an ordinary finite quadratic form. Expanding it, up to an additive term independent of S, gives

$$\begin{aligned} H(S|t) &= \frac{1}{2} S^T K^{-1} S + \frac{1}{2b^2}(t - HS)^T(t - HS) \\ &= \frac{1}{2} S^T \left(K^{-1} + b^{-2} H^T H\right) S - b^{-2} S^T H^T t + \text{const.} \end{aligned} \tag{36}$$

Taking the Hessian of the conditional Hamiltonian with respect to S gives

$$\nabla_S^2 H(S|t) = K^{-1} + b^{-2} H^T H = K^{-1} + b^{-2} D_M. \tag{37}$$

Thus the Hessian is exactly the posterior precision on the full latent lattice. All mask geometry in this precision enters through $D_M = H^T H$. We write it as

$$P_M = K^{-1} + b^{-2} D_M. \tag{38}$$

Completing the square in this quadratic Hamiltonian then identifies the posterior mean and covariance. Specifically,

$$H(S|t) = \frac{1}{2}(S - \mu_M)^T P_M (S - \mu_M) + \text{const.}, \tag{39}$$

where

$$\mu_M = P_M^{-1} b^{-2} H^T t. \tag{40}$$

Therefore the square-completion calculation gives the exact Gaussian posterior covariance

$$\Sigma_M = P_M^{-1} = \left(K^{-1} + b^{-2} D_M\right)^{-1}, \tag{41}$$

and the posterior mean $\mu_M = \Sigma_M b^{-2} H^T t$.

The observation-space forms follow by applying the Woodbury identity, also called the matrix inversion lemma, to $K^{-1} + b^{-2} H^T H$.[13,14] Appendix B records this algebraic step in the form used for computation. This produces an algebraically identical expression whose matrix inverse is the $m \times m$ observation-space matrix $HKH^T + b^2 I_m$:

$$\Sigma_M = K - KH^T(HKH^T + b^2 I_m)^{-1} HK, \tag{42}$$

and the corresponding posterior mean is

$$\mu_M = KH^T(HKH^T + b^2 I_m)^{-1}t. \tag{43}$$

Finally, multiplying the latent-space precision by the Fourier transform gives $F(K^{-1} + b^{-2}D_M)F^* = \Lambda^{-1} + b^{-2}FD_MF^*$, which is the stated coupled Fourier representation. □

Remark 1. The theorem is conditional on a fixed realized mask. Independent scalar Fourier modes occur in the diagonal-mask boundary cases. Random masks introduced later are outer generation models for realized masks; after a mask is realized, the posterior precision is determined by its corresponding D_M.

5. Fourier Representation and Mode Coupling

In the complete-observation formulation, the prior precision and the likelihood precision are diagonal in the same Fourier basis. This simultaneous diagonalization is what reduces the posterior calculation to independent scalar Fourier modes. In the sparse-observation formulation, the prior keeps this property, but the likelihood precision changes from $b^{-2}I_L$ to $b^{-2}D_M$. This section makes the realized mask geometry visible in Fourier space by writing the mask contribution as $Q_M = FD_MF^*$. Proposition 1 is the explicit finite-mask statement that this matrix generally converts the complete-observation scalar-mode posterior into a coupled finite-mode posterior.

Define

$$\widetilde{\Sigma}_M = F\Sigma_M F^*. \tag{44}$$

Using $K = F^*\Lambda F$, the posterior precision matrix in Fourier space is

$$\widetilde{\Sigma}_M^{-1} = \Lambda^{-1} + b^{-2}FD_MF^*. \tag{45}$$

Define

$$Q_M = FD_MF^*. \tag{46}$$

Its matrix elements are

$$(Q_M)_{nn'} = \frac{1}{L}\sum_{j=0}^{L-1} M_j e^{-i(\kappa_n - \kappa_{n'})j}. \tag{47}$$

Equivalently, define the finite Fourier components of the realized mask by

$$\widehat{M}_q = \sum_{j=0}^{L-1} M_j \exp\left(-\frac{2\pi i q j}{L}\right), \qquad q \in \mathbb{T}_L. \tag{48}$$

Since $\kappa_n - \kappa_{n'} = 2\pi(n - n')/L$, the mask term satisfies

$$(Q_M)_{nn'} = \frac{1}{L}\widehat{M}_{n-n'}, \tag{49}$$

where the index difference is understood modulo L. Thus the off-diagonal elements of Q_M

are not arbitrary matrix elements; they are precisely the finite Fourier coefficients of the realized observation mask. The observation geometry therefore enters the posterior precision through the spectrum of the mask. A mask with a nonzero Fourier component at the difference frequency $n - n'$ couples the prior modes n and n'.

The diagonal entries are

$$(Q_M)_{nn} = \frac{m}{L}. \tag{50}$$

This identity separates the observed-site count from the realized geometry of the mask. The fraction m/L, or equivalently the missing rate $p = 1 - m/L$, fixes only the diagonal entries of Q_M. It does not determine the off-diagonal entries

$$(Q_M)_{nn'} = \frac{1}{L}\widehat{M}_{n-n'} \qquad (n \neq n'), \tag{51}$$

which are the finite Fourier components of the realized mask at nonzero difference frequencies. These components encode how the observed sites are spaced, clustered, or otherwise arranged on the finite torus.

The posterior trace is consequently not a function of the eigenvalue list of Q_M alone. By trace invariance,

$$E(M) = \frac{1}{L}\operatorname{Tr}\Sigma_M = \frac{1}{L}\operatorname{Tr}\left[\left(\Lambda^{-1} + b^{-2}Q_M\right)^{-1}\right]. \tag{52}$$

Thus the finite posterior uncertainty depends on how the mask matrix Q_M, written in the Fourier basis, is positioned relative to the prior spectrum Λ. The diagonal part records the missing rate, but the off-diagonal Fourier components determine the mode couplings that enter the inverse trace. In the complete-observation case, $M_j = 1$ for all j, so $\widehat{M}_q = L\delta_{q0}$ and the off-diagonal mode coupling disappears. For an arbitrary nonconstant mask, nonzero components with $q \neq 0$ generally remain and couple Fourier modes.

Proposition 1 (Fourier mode coupling by a nontrivial mask)**.** *For a binary full-lattice mask, $Q_M = FD_MF^*$ is diagonal in the Fourier basis if and only if $D_M = I_L$ or 0. Hence any true sparse mask with $0 < m < L$ has at least one nonzero off-diagonal Fourier component*

$$\widehat{M}_{n-n'} = \sum_{j=0}^{L-1} M_j e^{-i(\kappa_n - \kappa_{n'})j} \qquad (n \neq n'), \tag{53}$$

and it couples Fourier modes in the posterior precision $\widetilde{\Sigma}_M^{-1} = \Lambda^{-1} + b^{-2}Q_M$.

Proof. If $D_M = I_L$, then $Q_M = I_L$; if $D_M = 0$, then $Q_M = 0$. These are the complete and empty diagonal cases. Conversely, suppose that Q_M is diagonal. Then $D_M = F^*Q_MF$ is a circulant matrix. Since D_M is also diagonal, it must be a scalar multiple of I_L. The binary condition on the diagonal entries then gives either $D_M = I_L$ or 0. Thus any binary mask with $0 < m < L$ necessarily leaves off-diagonal entries in Q_M, equivalently nonzero finite Fourier components at some nonzero difference frequencies. □

The proposition identifies the finite-matrix mechanism behind the deterministic comparisons below. The missing rate fixes $(Q_M)_{nn} = m/L$, whereas the realized mask geometry fixes the off-diagonal Fourier components. These components are precisely the entries that convert the complete-observation scalar-mode posterior into the coupled inverse-trace quantity $E(M)$. Hence two masks with the same p can have different posterior covariances and different posterior traces, as in Table IV.

Remark 2 (Mask spectrum and relative orientation)**.** Although Q_M is unitarily similar to D_M, its eigenvalue list contains only the observed-site count: m eigenvalues are one and $L-m$ eigenvalues are zero. Therefore all masks with the same m have the same eigenvalues of Q_M. The posterior trace, however, is

$$\frac{1}{L}\,\mathrm{Tr}\left[\left(\Lambda^{-1} + b^{-2}Q_M\right)^{-1}\right], \tag{54}$$

and this trace depends on the relative orientation of Q_M and the diagonal prior spectrum Λ, not on the eigenvalues of Q_M alone. Equivalently, Λ^{-1} selects the Fourier basis as a preferred basis, and the entries of Q_M in that basis are the finite mask Fourier components $L^{-1}\widehat{M}_{n-n'}$. For a smooth prior, the larger eigenvalues of Λ often lie in low-frequency modes. The amount of low-frequency mask spectral weight and its coupling to other modes can therefore change the inverse trace even when the missing rate is fixed. This is the Fourier-space reason why regular and clustered masks with the same m can leave different posterior uncertainties.

Figure 2 summarizes the structural distinction between the complete-observation case and the sparse-mask case.

5.1 Complete-observation reduction

When $D_M = I_L$, the posterior covariance becomes diagonal in Fourier space:

$$\widetilde{\Sigma}_{\mathrm{full},n} = \left(\lambda_n^{-1} + b^{-2}\right)^{-1} = \frac{\lambda_n b^2}{\lambda_n + b^2}. \tag{55}$$

The posterior mean has the scalar coefficient

$$\alpha_n = \frac{\lambda_n}{\lambda_n + b^2}. \tag{56}$$

Thus the complete-observation formula of the previous formulation[12)] is retained exactly as the boundary case $D_M = I_L$.

Corollary 1 (Complete-observation matched error)**.** *For the complete-observation mask $D_M = I_L$, the matched reconstruction error is*

$$E_{\mathrm{full}} = \frac{1}{L}\sum_{n=0}^{L-1} \frac{\lambda_n b^2}{\lambda_n + b^2}. \tag{57}$$

6. Matched Mean Squared Error

In the previous complete-observation formulation, the matched error $E_{1\,\mathrm{min}}$ was obtained as a Fourier-mode-wise sum because the posterior covariance was diagonal in the Fourier basis.

Corollary 1 is the corresponding complete-observation reduction in the present notation. For a sparse mask, $\widetilde{\Sigma}_M$ is generally non-diagonal in Fourier space, so the matched error is represented by the full posterior trace rather than by an independent scalar-mode sum. The natural finite-dimensional continuation of the previous $E_{1\,\min}$ is therefore the normalized posterior trace defined below. In the complete-observation limit $D_M = I_L$, this trace reduces to the mode-wise expression in Corollary 1 and hence to the previous complete-observation expression.

For a fixed realized mask M, define

$$E(M) = \frac{1}{L}\operatorname{Tr}\Sigma_M. \tag{58}$$

This quantity is the Bayes risk of the matched posterior-mean estimator after averaging over both the latent Gaussian field and the observation noise under the same model. Since the posterior covariance in a linear Gaussian model is independent of the realized data values t, the matched error depends on the mask and the model parameters through Σ_M.

Using the Woodbury form,

$$E(M) = \frac{1}{L}\operatorname{Tr} K - \frac{1}{L}\operatorname{Tr}\left[KH^T(HKH^T + b^2 I_m)^{-1}HK\right]. \tag{59}$$

The second term is the variance reduction due to the observed sites. This form also shows why an $m \times m$ inverse is useful when the number of observations is much smaller than L.

Proposition 2 (Monotonicity under adding observations)**.** *Let M and M' be two masks with $D_{M'} - D_M$ positive semidefinite. Then*

$$\Sigma_{M'} \preceq \Sigma_M, \qquad E(M') \leq E(M). \tag{60}$$

Proof. The corresponding posterior precisions satisfy

$$K^{-1} + b^{-2}D_{M'} \succeq K^{-1} + b^{-2}D_M. \tag{61}$$

For positive definite matrices, larger precision implies smaller covariance in the Loewner order. Taking the trace gives the result. This is consistent with the variance-reduction representation above: adding observed sites increases the non-negative reduction from the prior trace and therefore reduces posterior variance. □

Corollary 2 (Bounds between complete and empty observations)**.** *Let*

$$E_{\text{prior}} = \frac{1}{L}\operatorname{Tr} K \tag{62}$$

be the empty-observation error, and let E_{full} be the complete-observation error in Corollary 1. Then, for any observation mask,

$$E_{\text{full}} \leq E(M) \leq E_{\text{prior}}. \tag{63}$$

Proof. The claim follows from Proposition 2 by comparing an arbitrary mask with the complete-observation mask $D_M = I_L$ and the empty-observation mask $D_M = 0$. □

The bounds show that the finite-mask error is continuously ordered between the complete- and empty-observation limits. A practical threshold can be introduced by specifying an external tolerance, for example by declaring a reconstruction acceptable when

$$E(M) < \varepsilon. \tag{64}$$

The tolerance ε is supplied by the application or by an additional limiting problem.

Remark 3 (Continuous degradation)**.** For fixed finite L, fixed covariance K, and finite noise variance $b^2 > 0$, sparse observation changes the posterior covariance through a finite-dimensional matrix formula. The intrinsic quantity is the continuous trace indicator

$$E(M) = \frac{1}{L}\operatorname{Tr}\Sigma_M. \tag{65}$$

A binary decision can be added by specifying an external tolerance or another criterion outside the posterior formula itself.

7. Random Masks and Random Observation Intervals

The posterior formula is conditional on a realized mask. Random masks are used here as a generation and reporting layer: they specify a distribution over finite masks, and the posterior trace is evaluated after each mask has been realized.

As a standard example, consider independent Bernoulli observation,

$$M_j \sim \text{Bernoulli}(\rho), \qquad 0 \le \rho \le 1, \tag{66}$$

independently over j. Here ρ is the Bernoulli observation rate. For each realized mask, the observed-site count is $m = |\Omega|$ and the realized missing rate is $p(M) = 1 - m/L$; the quantity $1-\rho$ is the nominal missing rate of the generation model. For each realized mask,

$$E(M) = \frac{1}{L}\operatorname{Tr}\left[\left(K^{-1} + b^{-2}D_M\right)^{-1}\right] \tag{67}$$

is the deterministic finite-mask posterior trace. The random-mask model induces the outer average

$$\overline{E}(\rho) = \mathbb{E}_M[E(M)], \tag{68}$$

which averages these deterministic posterior traces over realized masks. The mask averaging is performed outside the matrix inverse, after each finite mask has been evaluated.

Proposition 3 (Mean-mask approximation and exact averaging)**.** *Because the matrix inverse is nonlinear,*

$$\mathbb{E}_M\left[\left(K^{-1} + b^{-2}D_M\right)^{-1}\right] \neq \left(K^{-1} + b^{-2}\mathbb{E}_M[D_M]\right)^{-1} \tag{69}$$

in general. For Bernoulli observation,

$$\mathbb{E}_M[D_M] = \rho I, \tag{70}$$

so replacing D_M by ρI gives a mean-mask approximation rather than the exact mask-averaged *posterior covariance.*

For one-dimensional cyclic masks, observation intervals may also be summarized by positive integer gaps $G_1, \ldots, G_m$ between successive observed sites, with

$$\sum_{s=1}^{m} G_s = L. \tag{71}$$

This gap description is a diagnostic of the realized geometry. Once M is fixed, the posterior covariance remains

$$\Sigma_M = \left(K^{-1} + b^{-2} D_M\right)^{-1}. \tag{72}$$

For $m = 0$, $\Sigma_M = K$ and $E(M) = E_{\text{prior}}$. For $m \geq 1$, we report

$$G_{\max} = \max_s G_s, \qquad p(M) = 1 - \frac{m}{L}, \tag{73}$$

together with the normalized trace

$$\frac{E(M)}{E_{\text{prior}}}, \qquad E_{\text{prior}} = \frac{1}{L} \operatorname{Tr} K. \tag{74}$$

In the circumferential-array interpretation, the corresponding largest angular interval between successive available microphones is

$$\Delta\theta_{\max} = \frac{2\pi G_{\max}}{L}. \tag{75}$$

The pair $(p(M), G_{\max})$ is a compact reporting diagnostic that complements the realized-mask posterior covariance.

8. Synthetic Circumferential Microphone-array Benchmark

This section recasts the finite mask-geometry calculations as a dimensionless synthetic benchmark for missing-channel reconstruction in the circumferential microphone array of Sect. 1.1. Consider $L = 64$ nominal microphones equally spaced around one circular-duct cross-section at

$$\theta_j = \frac{2\pi j}{64}, \qquad \Delta\theta = \frac{2\pi}{64} = 5.625^\circ. \tag{76}$$

The benchmark is not based on measured duct data, and the covariance parameters below are not calibrated to a particular acoustic installation. Its purpose is to isolate the effect of channel-availability geometry while retaining the exact finite posterior covariance derived above.

We report the realized-mask trace $E(M) = L^{-1} \operatorname{Tr} \Sigma_M$, interpreted as the residual prediction variance averaged over all 64 nominal microphone positions, the realized unavailable-channel fraction $p(M) = 1 - m/L$, and the largest cyclic gap $G_{\max}$ between successive usable channels. In angular units,

$$\Delta\theta_{\max} = \frac{2\pi G_{\max}}{L} = 5.625^\circ G_{\max}. \tag{77}$$

For deterministic comparisons, we also examine the finite Fourier spectrum of the availability mask. For Bernoulli summaries, ρ denotes the channel-availability rate and $1-\rho$ is used only as the nominal unavailable fraction.

For each prescribed parameter setting, the finite procedure is: construct the admissible covariance, prescribe or generate a channel-availability mask, form H, compute

$$\Sigma_M = K - KH^T(HKH^T + b^2 I_m)^{-1}HK, \tag{78}$$

and report $E(M) = L^{-1}\,\mathrm{Tr}\,\Sigma_M$ together with $p(M) = 1 - m/L$ and, when $m \geq 1$, $G_{\max}$ or its angular equivalent $\Delta\theta_{\max}$.

The covariance K is the minimum-image Gaussian constructed from Eq. (17) with $L = 64$, $a = 1$, and $\ell = 3$, used in the admissible range satisfying Eq. (19). Here $a = 1$ normalizes the pressure-fluctuation variance to $K_{jj} = a^2 = 1$, while the lattice correlation length $\ell = 3$ corresponds to the angular correlation length

$$\ell_\theta = \frac{2\pi\ell}{L} \simeq 0.295 \text{ rad} \simeq 16.9^\circ. \tag{79}$$

The value $b^2 = 1$ is interpreted as the normalized measurement-noise variance. These dimensionless choices preserve the previous covariance and all reported numerical results; they are not claimed to constitute an acoustically calibrated covariance model.

For every realized mask, the posterior covariance is evaluated by the Woodbury form above. Figure 3 shows one posterior-mean reconstruction with 32 evenly distributed usable channels. Table III and Fig. 4 are obtained by generating 100 Bernoulli channel-availability masks for each ρ and then computing $E(M)$ separately for each realized mask. Table IV compares two deterministic masks: the regular alternating mask, in which every other channel is usable, and the clustered missing block, in which 32 adjacent channels are unavailable and the other 32 adjacent channels are usable. The latter describes contiguous groupwise missingness without assigning it to a particular failure mechanism. Figure 5 uses a different generation layer from Table III: it fixes $m = 32$ and plots fixed-count random availability masks together with the same two deterministic reference masks. Thus the Bernoulli average in Table III and the fixed-count scatter in Fig. 5 are not the same random experiment; both are reported only to make the finite realized-mask dependence of $E(M)$ transparent. This paragraph fixes only the numerical protocol and adds no new theory.

The single-realization display in Fig. 3 identifies the latent pressure fluctuation, observed noisy microphone values, and posterior mean around the circumference for one regular alternating mask. Its primary horizontal axis is the lattice-site coordinate $j = 0, \dots, 63$, and its lower secondary axis uses $\theta_j = 360^\circ j/64$; hence the sites $j = 0, 16, 32, 48$ correspond to $0, 90, 180, 270^\circ$, respectively. The posterior-trace dependence on channel geometry is summarized in Tables III and IV and in Figs. 4 and 5.

Table III records finite posterior-trace computations for Bernoulli channel-availability

Table III. Results from the synthetic circumferential microphone-array benchmark for Bernoulli channel-availability masks, with $L = 64$, $a = 1$, $\ell = 3$, and $b^2 = 1$. The averages and standard deviations are over 100 mask realizations, except that $\rho = 1$ is deterministic and its standard deviations are therefore zero. The column $1 - \rho$ is the nominal unavailable-channel fraction of the Bernoulli generation model; each realization is evaluated with its own usable-channel count $m = |\Omega|$ and realized unavailable fraction $p(M) = 1 - m/L$. Values of $G_{\max}$ are in channel intervals; one interval is 5.625°. These are synthetic, dimensionless results rather than measurements from a duct experiment.

L	ρ	$1-\rho$	$\langle G_{\max}\rangle$	$\mathrm{sd}(G_{\max})$	$\langle E(M)\rangle$	$\mathrm{sd}(E(M))$
64	1.0	0.0	1.00	0.00	0.1957	0.0000
64	0.8	0.2	3.38	0.80	0.2381	0.0133
64	0.5	0.5	6.35	1.95	0.3488	0.0390
64	0.2	0.8	15.60	5.77	0.6134	0.0774

Table IV. Deterministic comparison in the synthetic circumferential microphone-array benchmark with $L = 64$, $a = 1$, $\ell = 3$, and $b^2 = 1$. Both masks have the same unavailable-channel fraction $p(M) = 1 - m/L = 0.5$. The regular mask uses every other microphone, whereas the clustered mask leaves 32 adjacent channels unavailable without specifying a failure cause. The bracketed angle is $\Delta\theta_{\max} = 2\pi G_{\max}/L$, the largest azimuthal interval between usable channels. The residual prediction variance $E(M)$ is averaged over all nominal microphone positions. The values are synthetic and dimensionless, not measured acoustic data.

L	$p(M)$	Channel geometry	$G_{\max}$	$E(M)$
64	0.5	alternating 32 usable	2	0.3116
64	0.5	contiguous 32 missing	33	0.5643

masks. As the availability rate ρ decreases, the nominal unavailable fraction $1 - \rho$ increases and the average posterior trace increases. Thus reconstruction uncertainty grows when fewer microphone channels are usable, consistently with the monotonicity statement in Proposition 2. Each unit of $G_{\max}$ or its standard deviation corresponds to an azimuthal interval of 5.625°.

Table IV gives the complementary fixed-rate comparison. With $p(M) = 0.5$, the regular alternating mask and the contiguous 32-channel missing block have the same number of usable channels but different posterior traces: 0.3116 and 0.5643, respectively. The latter is about 1.8 times the regular-mask value. Thus, even with the same 32 unavailable channels, their circumferential arrangement changes the full-array average reconstruction uncertainty by a practically substantial factor in this synthetic benchmark. Their largest usable-channel gaps are $\Delta\theta_{\max} = 11.25$ and 185.625°, respectively.

Figure 4 summarizes the Bernoulli results in Table III: the mean residual prediction variance increases with the nominal unavailable-channel fraction $1-\rho$. Figure 5 keeps $p(M) = 0.5$ fixed and plots $E(M)/E_{\text{prior}}$ against $G_{\max}$, with the two deterministic masks from Table IV as references. The random masks in this scatter are fixed-count random masks with exactly $m = 32$ usable channels, generated under the constraint $m/L = 0.5$, rather than Bernoulli masks. It should be read as a finite diagnostic display: at the same usable-channel count, the azimuthal gap geometry supplies information about the posterior trace beyond the unavailable fraction. In this finite fixed-count case study, larger angular gaps tend to accompany larger normalized posterior traces.

The deterministic comparison can also be read directly in Fourier space. Define the finite mask components

$$\widehat{M}_q = \sum_{j=0}^{L-1} M_j \exp\left(-\frac{2\pi i q j}{L}\right), \qquad q \in \mathbb{T}_L. \tag{80}$$

Then

$$(Q_M)_{nn'} = \frac{1}{L}\widehat{M}_{n-n'} \tag{81}$$

with the index difference understood modulo L. The two masks in Table IV have the same value of m, so Q_M has the same eigenvalue list in both cases. What differs is the mask spectrum in the Fourier basis in which Λ is diagonal. For the regular alternating mask at $L = 64$, the nonzero mask components are concentrated at $q = 0$ and the Nyquist frequency $q = 32$. For the clustered block, the mask spectrum is broad and contains large low-frequency components. Since the smooth Gaussian prior has large low-frequency eigenvalues, this difference changes the inverse trace of $\Lambda^{-1} + b^{-2}Q_M$ and gives a direct Fourier-space interpretation of the larger posterior trace in Table IV.

9. Discussion

For the circumferential microphone-array example, the main structural consequence of the mask-conditioned Fourier formulation is the separation between the number of usable channels and the Fourier geometry carried by Q_M. Prior diagonalization survives in periodic Fourier coordinates, whereas posterior mode independence generally does not. The count fixes only the diagonal entries $(Q_M)_{nn} = m/L$, while the off-diagonal entries $(Q_M)_{nn'} = L^{-1}\widehat{M}_{n-n'}$ couple prior modes in $\Lambda^{-1} + b^{-2}Q_M$. Missing-channel reconstruction on the idealized ring is therefore a finite mode-coupling problem rather than a collection of independent scalar shrinkage problems.

For fixed model parameters and a fixed mask, the posterior covariance Σ_M is independent of the observed values t. Hence $E(M) = L^{-1}\operatorname{Tr}\Sigma_M$ can be computed once the mask is known and used as a mask-conditioned uncertainty indicator for comparing missing-channel patterns within this sensor-ring model.

The synthetic 64-channel benchmark shows this dependence without relying on an asymptotic argument. At the same unavailable-channel fraction $p(M) = 0.5$, Table IV gives posterior traces 0.3116 and 0.5643 for the regular alternating mask and the contiguous missing block. The latter uncertainty is about 1.8 times the former even though both configurations contain 32 unavailable channels. Therefore the number of unavailable microphones alone does not determine reconstruction uncertainty; their circumferential placement must also be retained. These values are a dimensionless synthetic comparison and are not a performance claim for an operating duct.

The Fourier representation explains the difference. Although the eigenvalues of $Q_M = FD_MF^*$ are fixed by the count m, the posterior trace is the inverse trace of $\Lambda^{-1} + b^{-2}Q_M$. It therefore depends on the orientation of the mask matrix in the Fourier basis selected by the prior spectrum Λ. This is the finite-matrix mechanism by which mask geometry changes spectral reconstruction uncertainty.

This viewpoint also clarifies the relation to neighboring methods. In Wiener-filter terminology, the complete-observation periodic case is the independent scalar Fourier-shrinkage limit. In the sparse-mask problem, the central object is the non-diagonal mask matrix Q_M inside the posterior precision. Random masks and gap-sensitive quantities are reporting layers outside the conditional posterior formula: a generation model defines an outer average over realized masks, and diagnostics such as $(p(M), G_{\max})$ record geometry beyond the missing rate.

The duct interpretation depends on equally spaced sensors on a circular cross-section, a rotationally stationary pressure covariance, and homogeneous independent sensor noise. For a noncircular duct, an unequally spaced sensor layout, a circumferentially nonuniform mean flow, or azimuthally varying wall conditions, the covariance need not be circulant. Sensor-dependent or correlated noise likewise replaces the common-noise likelihood used here. Under those conditions the scalar Fourier-mode structure and the precision $\Lambda^{-1} + b^{-2}Q_M$ do not remain valid without modification, although the general finite-dimensional Gaussian conditioning identity still applies. Accordingly, the synthetic results should not be generalized to an operating duct before these assumptions are checked.

10. Conclusion

Motivated by missing channels in a circumferential microphone array at one circular-duct cross-section, we have formulated the finite linear-Gaussian posterior for an arbitrary realized mask on an equally spaced periodic sensor ring. The standard conditioning identity gives the posterior mean and covariance for each fixed mask; the contribution here is the explicit finite mask-geometry representation, not a new Gaussian-conditioning theorem or a duct-mode identification method.

The complete-observation scalar Fourier formula is recovered at $D_M = I_L$. Sparse masks introduce $Q_M = FD_M F^*$, whose off-diagonal finite Fourier components couple modes in $\Lambda^{-1} + b^{-2} Q_M$. The matched error becomes $E(M) = L^{-1} \operatorname{Tr} \Sigma_M$, a mask-conditioned posterior-trace indicator. The figures display this mechanism through the structural Fourier schematic, the Bernoulli-mask trace summary, and the fixed-rate gap diagnostic.

In the synthetic 64-channel circumferential-array benchmark, the same 50% unavailable-channel fraction gives $E(M) = 0.3116$ for alternating missing channels and $E(M) = 0.5643$ for one contiguous 32-channel missing block. Thus, within the stated ring model, the unavailable fraction is a coarse descriptor and the realized circumferential mask geometry must be retained.

These conclusions are limited to a circular, equally instrumented ring with a circulant prior covariance and common independent sensor noise. Noncircular geometry, unequal sensor spacing, circumferentially nonuniform mean flow, or sensor-dependent noise generally destroys that circulant covariance or the common-noise likelihood, so the scalar Fourier-mode structure used here does not carry over unchanged. The analysis is therefore an exact finite reference calculation under explicit idealizations, not a validation of missing-channel reconstruction performance for fan or compressor ducts in general.

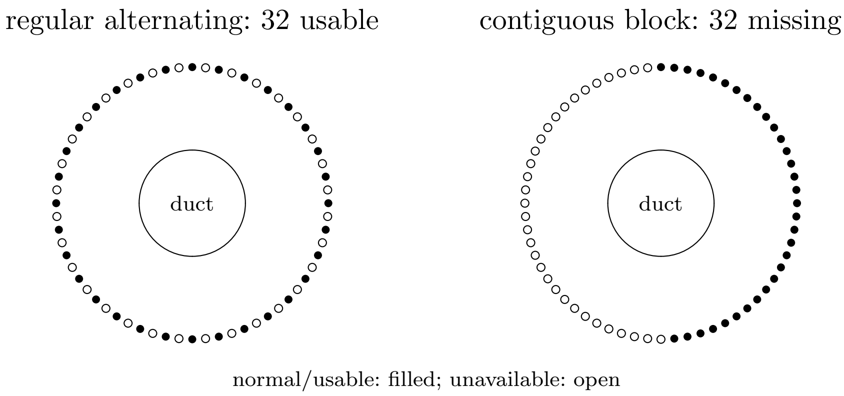

Fig. 1. Mask-geometry schematic for the synthetic circumferential microphone-array benchmark. Each circular-duct cross-section contains 64 equally spaced microphones at angular intervals of 5.625°. The left panel makes every other channel usable; the right panel leaves one contiguous group of 32 channels unavailable. Filled markers denote normal, usable channels and open markers denote unavailable channels. The cyclic ordering joins the top channel to itself after one full turn. The right panel represents groupwise missingness without attributing it to a particular failure mechanism.

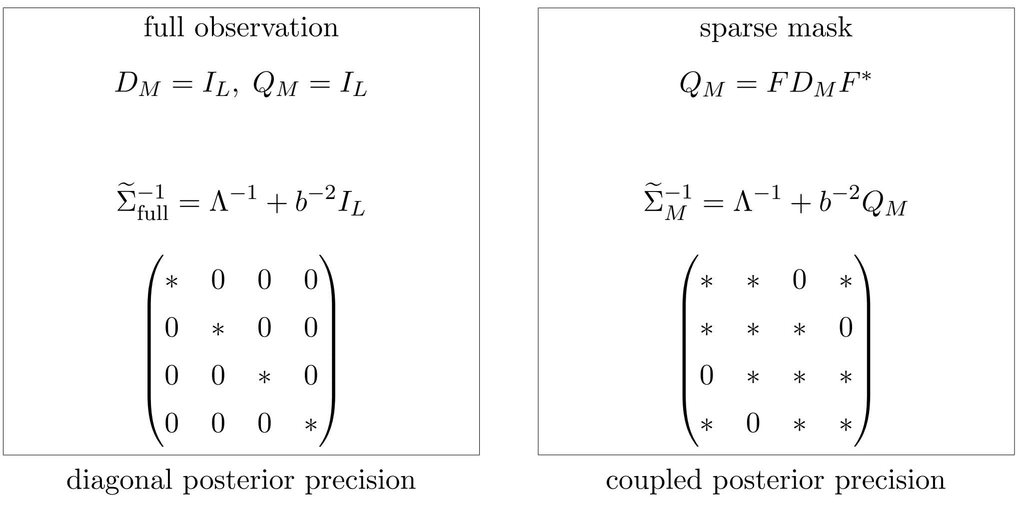

Fig. 2. Fourier-space mode-coupling schematic. Under complete observation, $D_M = I_L$ gives $Q_M = I_L$, so the Fourier-space posterior precision is diagonal. Under a sparse mask, $Q_M = F D_M F^*$ is generally non-diagonal and couples Fourier modes. The schematic summarizes the structural distinction between complete and sparse observation.

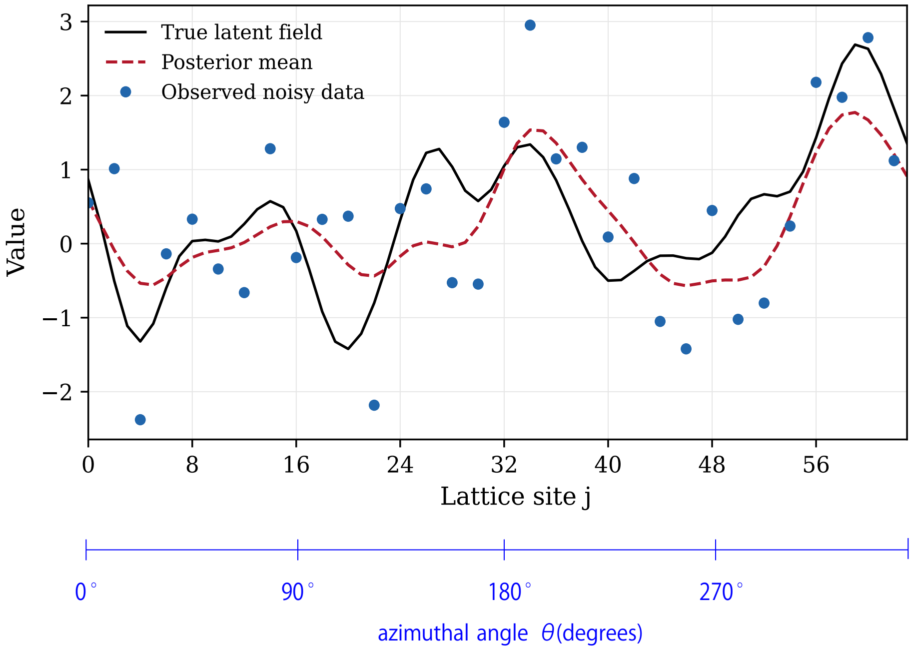

Fig. 3. (Color online) Single-realization reconstruction for the synthetic circumferential microphone-array benchmark with $L = 64$. The solid and dashed curves show the latent pressure fluctuation $S(\theta_j)$ and posterior mean $\mu_{M,j}$, respectively, and the markers show the observed noisy microphone values t_j at the 32 usable sites of a regular alternating mask with $p(M) = 0.5$. The primary horizontal axis gives the lattice-site index $j = 0, \ldots, 63$, and the lower axis gives the azimuthal angle $\theta_j = 360°j/64$; thus 0, 90, 180, and 270° align with $j = 0$, 16, 32, and 48, respectively. The data are synthetic and dimensionless.

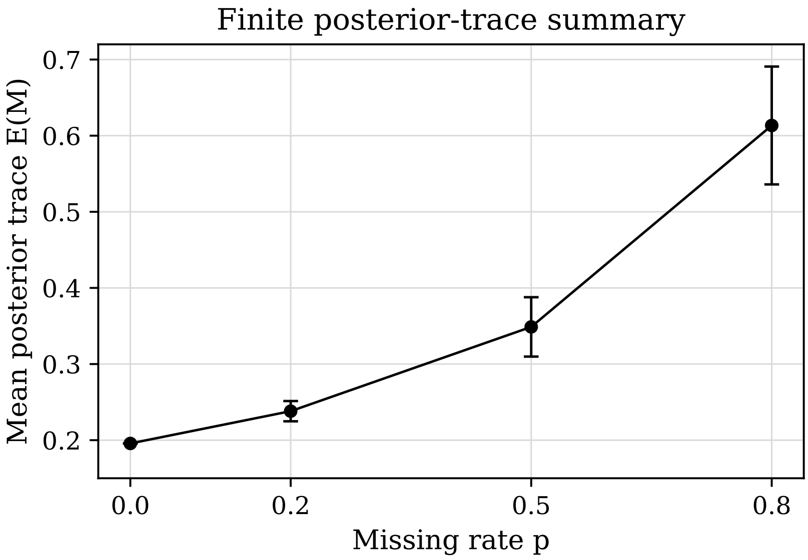

Fig. 4. Posterior-trace summary for the synthetic circumferential microphone-array benchmark based on Table III. The points show the mean full-array residual prediction variance $\langle E(M)\rangle$, and the vertical bars show $\mathrm{sd}(E(M))$, for $L = 64$, $a = 1$, $\ell = 3$, and $b^2 = 1$. The horizontal coordinate is the nominal Bernoulli unavailable-channel fraction $1 - \rho$; each posterior trace is computed after a finite availability mask is realized. The results are synthetic and dimensionless, not measured acoustic data.

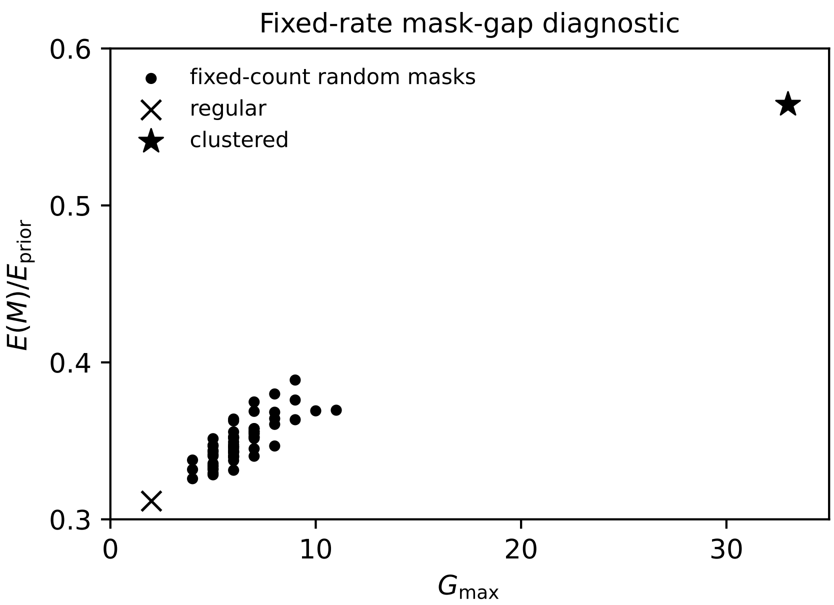

Fig. 5. Gap-diagnostic scatter for the synthetic circumferential microphone-array benchmark with $L = 64$, $m = 32$ ($p(M) = 1 - m/L = 0.5$), $a = 1$, $\ell = 3$, and $b^2 = 1$. Small filled circles are fixed-count random-mask realizations with exactly $m = 32$ usable channels. The regular alternating mask and the contiguous 32-channel missing block from Table IV are included as cross and star reference symbols. One unit on the $G_{\max}$ axis corresponds to 5.625°, so the horizontal coordinate also represents the largest azimuthal interval between usable microphones. This synthetic, dimensionless display indicates that larger angular gaps tend to accompany larger normalized posterior traces in the finite fixed-count case study.

Appendix A: Zero-eigenvalue Convention

This appendix fixes the support convention for positive-semidefinite covariance matrices. The main theorem is stated for the positive-definite case, and the convention below extends the notation to finite covariance supports.

The main text writes inverse-covariance expressions for the strictly positive case. If an admissible finite periodic covariance has zero eigenvalues, then K is positive semidefinite and the Gaussian prior is supported on the range of K. In that case, inverse expressions are interpreted on this support. Equivalently, one may replace K^{-1} by the Moore–Penrose pseudoinverse K^{+} and restrict the posterior calculation to the covariance support.

Let

$$K = F^{*}\Lambda F, \qquad \Lambda = \mathrm{diag}(\lambda_0, \ldots, \lambda_{L-1}), \tag{A·1}$$

with $\lambda_n \geq 0$. The support is spanned by the Fourier modes with $\lambda_n > 0$. On this support, Λ^{-1} is understood as

$$\lambda_n^{-1} \quad \text{for} \quad \lambda_n > 0. \tag{A·2}$$

Zero-variance modes are deterministic under the prior and have zero posterior variance. The posterior covariance formula is applied on the positive-variance subspace, and the covariance in zero-variance directions remains zero.

Thus the positive-semidefinite case is the same finite Gaussian posterior calculation restricted to the covariance support.

Appendix B: Woodbury Identity Form

This appendix records the observation-space implementation form with an $m \times m$ inverse. The identity is the standard Woodbury matrix identity, or matrix inversion lemma, in the form used for finite linear-Gaussian conditioning.[13,14]

For positive-definite K and $b^2 > 0$,

$$\left(K^{-1} + b^{-2}H^{T}H\right)^{-1} = K - KH^{T}(HKH^{T} + b^{2}I_m)^{-1}HK. \tag{B·1}$$

Multiplying the right-hand side by $K^{-1} + b^{-2}H^{T}H$ verifies the identity. The form is useful because the inverse inside the observation-space expression has dimension $m \times m$, whereas the latent-space precision has dimension $L \times L$.

The corresponding posterior mean is

$$\mu_M = KH^{T}(HKH^{T} + b^{2}I_m)^{-1}t. \tag{B·2}$$

Thus the Woodbury form is algebraically the same posterior written in observation-space coordinates.

Appendix C: Coordinate-free Finite Operator Form

This appendix records the coordinate-free finite-dimensional algebra behind the posterior identity used in the main text. It is not a separate dimensional extension of the one-dimensional periodic-lattice formulation.

Let a finite latent Gaussian vector $S \in \mathbb{R}^N$ have covariance K, and let observations be selected by an arbitrary finite operator $H \in \mathbb{R}^{m\times N}$:

$$t = HS + \xi, \qquad \xi \sim \mathcal{N}(0, b^2 I_m). \tag{C·1}$$

Then the same finite linear-Gaussian posterior identity gives

$$\Sigma = \left(K^{-1} + b^{-2}H^T H\right)^{-1} = K - KH^T(HKH^T + b^2 I_m)^{-1}HK. \tag{C·2}$$

If a finite unitary coordinate transform U is chosen, for example the finite unitary Fourier matrix corresponding to a finite periodic model, then

$$U\Sigma^{-1}U^* = UK^{-1}U^* + b^{-2}UH^T HU^*. \tag{C·3}$$

When UKU^* is diagonal, the first term is diagonal and the second term is the transformed mask precision. A sparse observation operator generally makes $UH^T HU^*$ non-diagonal, so the same finite algebra produces coordinate-mode coupling in the transformed representation.

This appendix therefore records only a form-preserving finite operator identity in finite dimensions; it does not assert a continuum limit, a large-system limit, or a sampling-geometry theorem.

References

1) C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, Cambridge, MA, 2006).
2) C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006).
3) D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, U.K., 2003).
4) P. D. Joppa, J. Aircr. **24**, 587 (1987).
5) L.-F. Wang, W.-Y. Qiao, L. Ji, and S.-Y. Yu, J. Aerosp. Power **29**, 917 (2014).
6) R. Wang, M. Yu, Y. Bai, L. Yu, and G. Dong, AIAA J. **61**, 4018 (2023).
7) N. Cressie, *Statistics for Spatial Data* (Wiley, New York, 1993) revised ed.
8) M. L. Stein, *Interpolation of Spatial Data: Some Theory for Kriging* (Springer, New York, 1999).
9) N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (MIT Press, Cambridge, MA, 1949).
10) A. V. Oppenheim and R. W. Schafer, *Discrete-Time Signal Processing* (Prentice-Hall, Upper Saddle River, NJ, 2010) 3rd ed.
11) E. J. Candès and M. B. Wakin, IEEE Signal Process. Mag. **25**, 21 (2008).
12) J. Tsuzurugi, J. Phys. Soc. Jpn. **94**, 104801 (2025).
13) D. A. Harville, *Matrix Algebra from a Statistician's Perspective* (Springer, New York, 1997).
14) G. H. Golub and C. F. Van Loan, *Matrix Computations* (Johns Hopkins University Press, Baltimore, MD, 2013) 4th ed.