Wireless Node Ranking via Fourier-Driven Maximum-Likelihood Estimation

Daniel Schneider
Independent
Germany
mail.schneiderd@gmx.de

Hannes Frey
University of Koblenz
Koblenz, Germany
frey@uni-koblenz.de

***Abstract*—We apply non-abelian Fourier theory to the problem of maximum-likelihood estimation of the order of a line topology, based on signal strength measurements under shadowing only. This yields up to a 60-fold speedup over standard methods for uncorrelated shadowing.**

Moreover, for correlated Gaussian shadowing with a singular covariance (e.g. the joint path-loss model), maximum-likelihood recovery is exact and independent of the shadowing variance (an infinite Le Cam threshold), which our implementation confirms for σ/α as high as 7.0. The key object is the covariance's null space. Its dimension is fixed by the topology; its orientation depends on the correlation function only through the per-link variances.

Although the mechanism formally requires an exact line, it is robust to small perturbations of the formation: an approximate kernel recovers as well as the exact one. We explore the dependence on the perturbation in detail for five nodes.

***Index Terms*—node ranking; RSSI; quadratic assignment; correlation; symmetric group; Fourier analysis; seriation**

I. Introduction

Wireless sensor and ad-hoc networks are frequently deployed in physically elongated, near-linear arrangements — along roads, pipelines, tunnels, building corridors, or the rows of a photovoltaic plant [19] — where the *order* of the nodes along the line is itself the quantity of interest (for routing, in-order data fusion, or fault localization), while GPS or anchor infrastructure is unavailable or too costly [15]. Received signal strength (RSSI), by contrast, is reported by virtually every radio at no extra cost. *Wireless node ranking* is the problem of recovering the order of such nodes from signal-strength measurements alone.

Even in the simple case of a line topology with uniform spacing, the problem is hard once one insists on *best possible* solutions — by which we mean the maximum-likelihood estimator (MLE). While there are heuristics like *order derivation* (AOD) [6], [7] in the wireless literature, the least-squares method (ALS) we proposed in [3] (also cf. [4]) is NP-hard in general. Its benefit is, as we will see in this work, that its solution is the MLE in the case of log-normal shadowing (LNS), i.e. uncorrelated Gaussian shadowing with fixed variance σ^2. As we have pointed out in [4], AOD can dominate ALS when correlation is present. However, AOD is only a heuristic when its underlying assumption does not hold, as is the case, for instance, in the joint path loss models ([13], [4]). Thus it is natural to ask for the MLE also in the case of correlated shadowing.

What comes as a surprise is that the concept of node ranking is known as *seriation* in the mathematical statistics community with — to the best of our knowledge — no connection to the wireless community. In the words of Atkins et al. [9], seriation is the *sequencing of elements in such a way that highly correlated pairs of elements are near each other.* This can clearly be solved by finding the maximizer of the graph correlation, which is the way Kondor states the *quadratic assignment problem* (QAP) in his seminal work [1] — which in turn we tackle with the Fourier Framework he developed there based on the Clausen Fourier Transform [2]. Even more surprising is that the noise-free seriation case is solved by a tool the wireless community knows well — spectral graph theory, with the Fiedler eigenvector at its center [9] — which nonetheless never led it to seriation. We adopt these seriation methods (Fiedler and the minimax-optimal SABRE [22]) as polynomial-time benchmarks. Crucially, these are spectral or convex *relaxations* built for polynomial-time, minimax-rate recovery; none computes the maximum-likelihood estimator, even under i.i.d. Gaussian noise — where our LNS algorithm, solving the exact least-squares QAP, *is* the MLE and already outperforms them. The decisive gap is under *correlated* shadowing: there the MLE is *best possible* — 100% recovery — which the i.i.d.-built seriation methods cannot reach.

Contributions. We give a maximum-likelihood treatment of node ranking. Under uncorrelated LNS the MLE is the QAP, which we solve exactly with a non-abelian Fourier branch-and-bound (*Fourier B&B*) built on the Kondor bound, up to $60\times$ faster than the Gilmore–Lawler bound (GLB). Under correlated shadowing with a *singular* covariance, we prove the MLE is determined entirely by the covariance *kernel* and recovers the order exactly, at every noise level (Theorem 1), which we confirm up to $\sigma/\alpha = 7$. Finally, we investigate how the kernel depends on perturbations both on the topology and the covariance matrix directly.

II. Related Work

Several works address ordering or localization in linear wireless networks without coordinate infrastructure. Lotker et al. [15] introduced range-free ranking via connectivity alone, and Zhu et al. [6] proposed the order-derivation algorithm

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

(AOD), which greedily assigns positions by assuming each node's strongest neighbor lies in the direction of increasing index; this assumption fails at high σ/α under uncorrelated noise, while spatial correlation in fact reinforces it [4]. RSS-based sequencing has further been validated experimentally in buildings [17], [18] and addressed with distributed and self-configuration heuristics for dense linear deployments [16], [19]. None of these works frame the problem as maximum-likelihood estimation or analyze performance under correlated shadowing.

Gudmundson [12] proposed the canonical exponential autocorrelation model $e^{-\Delta/d_0}$ for shadow fading over a spatial separation Δ. Agrawal and Patwari [13] introduced the joint path-loss model (JPLM), deriving inter-link shadowing correlations via a line-integral over a Gaussian random field, with empirical correlations up to 0.33 in multi-hop networks. We use the variance-scaled variant (vsJPLM) introduced in [4] to fix the link-length-dependent variance of the original model.

The ordering problem is equivalent to a QAP [4], [14]: maximizing the graph correlation *is* the QAP [1]. The Gilmore–Lawler bound [11] is the standard branch-and-bound lower bound; the Fourier B&B we use replaces it with the tighter Kondor bound [1].

The seriation literature (see Section III) studies exactly our problem under iid noise. Fogel et al. [20] reformulate seriation as the 2-SUM QAP and derive convex (doubly-stochastic QP) relaxations for noisy recovery. Issartel et al. [22] achieve the minimax-optimal rate $O((\sigma/\gamma)\sqrt{\log n/n})$ under iid sub-Gaussian noise, with γ their signal-strength parameter. Natik and Smith [23] prove that Fiedler seriation is consistent in the noisy graphon model, though at a suboptimal rate compared to SABRE. Both works assume independent noise entries. The regime of our correlation models lies entirely outside the scope of existing seriation theory.

III. Theoretical Background

A. Seriation: Fiedler and SABRE

In 1951, archaeologist W.S. Robinson introduced a matrix condition to seriate pottery types across excavation strata into a temporal sequence [8]: a symmetric $F \in \mathbb{R}^{n\times n}$ is a *Robinson matrix* if its entries decrease monotonically away from the diagonal, i.e. $F_{jk} < F_{ik}$ whenever $k \le i < j$ and $F_{ik} < F_{jk}$ whenever $i < j \le k$. *Seriation* is the problem of recovering a hidden permutation $\pi \in S_n$, where S_n is the *symmetric group*, from a noisy observation $A = F_\pi + N$, where F_π denotes a Robinson F permuted by π and N is a noise matrix.

The noise-free case is solved by the Fiedler vector: for the graph Laplacian $L = D - A$ of the similarity matrix A, with $D = \operatorname{diag}(A\mathbf{1})$, the eigenvector v_2 of the second-smallest eigenvalue $\lambda_2(L)$ is strictly monotone along the true ordering whenever A is exactly Robinson ($N = 0$) [9], so sorting the nodes by its entries recovers π exactly. This is *Fiedler seriation*.

As the noise N grows the Fiedler vector degrades. SABRE (*Seriate by Aggregating Bisections and by Re-Evaluating*, [22]) is the minimax-optimal estimator for the noisy regime and, under correlated shadowing, our strongest polynomial-time baseline. Its central object, playing the role the Fiedler vector plays above, is a *debiased* row distance $\widehat{D}_{ij}$: replacing the noise-inflated self-products of the row inner products by a robust near-neighbor surrogate cancels the noise-variance bias, so $\widehat{D}_{ij}$ consistently estimates the latent inter-node distance. SABRE then recovers π by aggregating local left/right bisections along these distances.

SABRE targets a normalized latent model ($O(1)$ entries, known σ), so our log-distance RSSI matrix $A_{ij} = -10\,\alpha \log_{10}|i-j| + N_{ij}$ needs adapting. Its debiased-distance stage — the source of the noise robustness — carries over unchanged; its bisection-and-refinement stages do not, since their thresholds are tuned to the normalized model and the log-distance signal bends the row geometry into an *arch* (horseshoe) on which any distance-embedding folds. We therefore order the nodes arch-resistantly — from the line-boundary node (least total similarity), a greedy nearest-neighbor path through the debiased distances, refined by 2-opt on the 2-SUM cost — which is exact noise-free.

B. Non-commutative Fourier Transforms

In the following, we will always assume Young's orthogonal representation (YOR) is being used, cf. [21]. Following [1], given a weighted graph $\mathcal{G}$ on n vertices with adjacency matrix $A \in \mathbb{R}^{n\times n}$, define its *graph function* $f_A\colon S_n \to \mathbb{R}$ by $f_A(g) = A_{g(n),\,g(n-1)}$. The function f is constant on each left S_{n-2}-coset gS_{n-2} (the $n(n-1)$ blocks $S_{i,j} = \{g \in S_n \mid g(n) = i,\, g(n-1) = j\}$), which is the key structural property exploited by the Fourier machinery.

Further, define the *graph correlation* of two symmetric $n \times n$ matrices A and B under permutation $\pi \in S_n$ as $f(\pi) = \sum_{i,j} A_{ij}\, B_{\pi(i),\pi(j)} = \operatorname{tr}(A\, P_\pi B P_\pi^\top)$, where P_π is the permutation matrix of π. Finding $\hat\pi = \arg\max_\pi f(\pi)$ is the *quadratic assignment problem* (QAP). The following bound limits the maximum of f over S_n in terms of the order-2 Fourier coefficients ($\lambda = [n], [n-1,1], [n-2,2], [n-2,1,1]$) of the non-commutative *Fourier transform* $\hat f(\lambda) = \sum_{\pi\in S_n} f(\pi)\,\rho_\lambda(\pi)$, where $\rho_\lambda, \lambda \vdash n$, are the irreducible representations, and a linear assignment problem $\mathrm{LAP}(F) := \max_{\pi\in S_n} \operatorname{tr}(P_\pi F)$, solved by the Hungarian algorithm in $O(n^3)$. The *refined Kondor bound* is

$$B_{\mathrm{kond}}(f) = \frac{1}{n!}\Big[(n{-}1)\mathrm{LAP}(F) - (n{-}2)\hat f([n]) + \sum_{\lambda\in\{[n-2,2],\,[n-2,1,1]\}} d_\lambda\, \|\hat f(\lambda)\|_1\Big], \qquad (1)$$

where d_λ is the dimension of ρ_λ, $\|\cdot\|_1$ the trace norm, and $F_{ij} = \sum_{\pi:\,\pi(j)=i} f(\pi)$ the *first-order marginal matrix*. Eq. (1) is our refinement of the bound (5.13) in [1]; the derivation is given in the extended version [5]. We have $\max_{\pi\in S_n} f(\pi) \le B_{\mathrm{kond}}(f)$.

The Fourier transform $\hat f_\lambda$ is computed efficiently via the Clausen decomposition [2], [21]. For each coset representative $g \in S_n/S_{n-2}$, define the coset function $h_g\colon S_{n-2} \to \mathbb{R}$ by $h_g(\tau) = f(g\tau)$. Then $\hat f(\lambda) =$

$\sum_{g\in S_n/S_{n-2}} \rho_\lambda(g) \bigoplus_{\lambda^-} [\hat{h}_g]_{\lambda^-}$ [21], where λ^- ranges over Young diagrams obtained from λ by removing one box and $\bigoplus$ denotes a direct sum of Fourier blocks. This reduces the Fourier transform over S_n to $n(n-1)$ transforms over S_{n-2}, giving the Clausen-type fast Fourier transform on S_n [2].

If f is a graph function, then h_g is constant which leads to the fact that $\hat{f}_\lambda$ is non-zero in YOR only for $\lambda = (n), (n-1,1), (n-2,2)$ and $(n-2,1,1)$, i.e. f is of order 2 and thus *band-limited*. Now f is itself not a graph function, but can be written as sum of graph functions [1], $f(g) = \frac{1}{(n-2)!}\sum_{\pi\in S_n} f_A(g\pi)\, f_B(\pi)$, which by the convolution theorem factors in Fourier space as $\hat{f}(\lambda) = \frac{1}{(n-2)!}\,\hat{f}_A(\lambda)\,(\hat{f}_B(\lambda))^\top$ for $\lambda \vdash n$. Together with the band-limitedness of graph functions, this means $\hat{f}(\lambda) = 0$ unless $\lambda \in \{(n), (n-1,1), (n-2,2), (n-2,1,1)\}$. This structure is what makes the (refined) Kondor bound (1) computable in $O(n^3)$ at each node of the Fourier B&B [1]. Moreover, the bound works particularly well in our node ranking context, pruning many of the inferior branches in the low to mid-noise range.

C. Correlated shadowing models

The shadowing enters the observation model $A = P_{\pi^\star} B P_{\pi^\star}^\top + N$ of Section IV as the additive noise N, with N_{ij} the log-normal shadowing on the link between the nodes at positions i and j. Fast fading is not modeled: it averages out over the RSSI samples collected per link [4]. We recall the spatially correlated models of [4], [13] in this matrix notation.

Correlation arises because nearby links cross overlapping terrain. We resolve it onto the $n-1$ unit segments of the line: the segments carry a shadowing vector $\xi = (\xi_1, \dots, \xi_{n-1})^\top$ with covariance C, and a link's shadowing is the sum of the segments it spans, up to a fixed per-link scale, $N = D\,E\,\xi$ with $E \in \{0,1\}^{m\times(n-1)}$ ($m = \binom{n}{2}$) the link–segment incidence matrix and D a positive diagonal. The shadowing is thus a *singular* Gaussian, $N \sim \mathcal{N}(0, K)$ with $K = D\,E\,C\,E^\top D$ and $\operatorname{rank} K = n-1 < m$ for non-singular C. Its kernel collects the linear relations satisfied identically by *every* realization of N: for $D = I$ the additive, or *betweenness*, relations $N_{ik} = N_{ij} + N_{jk}$, in general their D-weighted counterparts. Its dimension is $\binom{n}{2} - (n-1) = \frac{(n-1)(n-2)}{2}$ either way, and this kernel is what the maximum-likelihood estimator of Section IV exploits.

The segment covariance C and the scaling D fix the model. The joint path-loss model (JPLM, [13]) obtains ξ from a line integral of an isotropic Gaussian field, so $\operatorname{Cov}(\xi_k, \xi_l)$ decays with segment separation over a correlation length δ; we report results against $d = d_{12}/\delta$, the neighbor spacing d_{12} in units of δ. Because the original JPLM ($D = I$) gives links of different length different variances, we follow [4] and use the *variance-scaled* JPLM (vsJPLM), which rescales every link to a common variance σ^2 ($\operatorname{diag} K = \sigma^2$) and so keeps the constant-variance property of log-normal shadowing even when $d \gg 1$; this scaling is D, and it trades plain additivity for its weighted form. Uncorrelated log-normal shadowing (LNS), our baseline, is *not* a member of this family: its $K = \sigma^2 I$ has full rank, hence no kernel and no noise-free subspace.

IV. Maximum Likelihood Estimation

We now derive the maximum-likelihood estimator of the ordering $\pi^\star$ from the observation $A = P_{\pi^\star} B P_{\pi^\star}^\top + N$, for both uncorrelated and correlated Gaussian shadowing. Under uncorrelated shadowing the MLE reduces to the least-squares QAP of [3], [4], which we re-establish and solve exactly with the Fourier B&B of Section III-B. Under correlated shadowing this is no longer optimal — that AOD can beat ALS at moderate σ/α already shows the least-squares estimator is misspecified there. The singular covariance of the correlated models, however, makes the correlated MLE not merely better but *exact*: its likelihood is confined to the affine subspace fixed by $\ker K$, and projecting onto that kernel recovers $\pi^\star$ up to reversal at every noise level (Theorem 1), which our experiments confirm at 100% success at $\sigma/\alpha = 7$.

A. Link model

Following our previous work [4], the n nodes lie on a line in an unknown order; write $v_1, \dots, v_n$ for the nodes in their true spatial sequence, with equal neighbor spacing d_{12}, so nodes at positions i and j are at distance $|i-j|\,d_{12}$. Each transmission is attenuated by distance-dependent path loss and by shadowing, so the received power (in dB) on the link between them is

$$-10\,\alpha \log_{10}\big(|i-j|\,d_{12}\big) + \text{const} + N_{ij}, \tag{2}$$

with α the path-loss exponent and N_{ij} the shadowing of Section III-C. Collecting the measurements into a symmetric similarity matrix and absorbing the constant yields the observation model $A = P_{\pi^\star} B P_{\pi^\star}^\top + N$ with $B_{ij} = -10\,\alpha \log_{10}|i-j|$, where B is the noise-free path-loss template of the ordered line, $\pi^\star$ the unknown permutation taking measured indices to line positions, and $N \sim \mathcal{N}(0, K)$ the shadowing. As B is fixed and strictly Robinson (Section III-A), recovering $\pi^\star$ from A is exactly the seriation/QAP of Section III-B; the rest of this section derives its maximum-likelihood solution.

B. Least squares

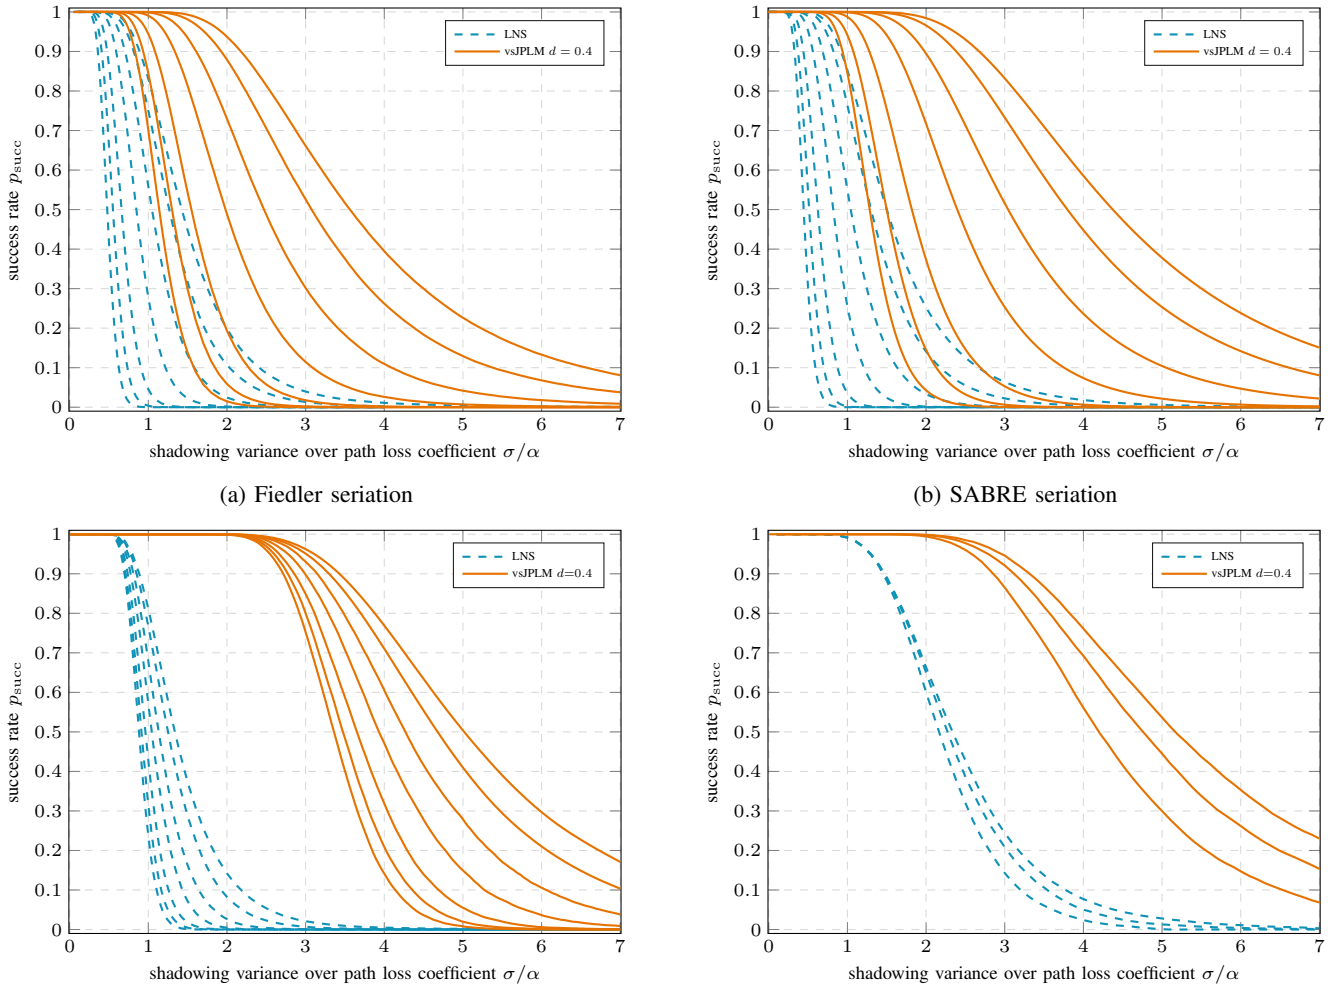
Under uncorrelated log-normal shadowing, link errors are i.i.d. Gaussian in the dB domain: $N_{ij} = N_{ji} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$ for $i < j$, with zero diagonal. Independence on the upper triangle factorizes the joint density, giving

$$\ell(\pi) \;=\; \text{const} \;-\; \frac{1}{2\sigma^2}\sum_{i<j}\big(A_{ij} - (P_\pi B P_\pi^\top)_{ij}\big)^2.$$

Since σ^2 does not depend on π, the MLE is the least-squares estimator $\hat{\pi}_{\text{MLE}} = \arg\min_\pi \|A - P_\pi B P_\pi^\top\|_F^2$. Using $\|P_\pi B P_\pi^\top\|_F^2 = \|B\|_F^2$ and that $\|A\|_F^2$ is constant in π,

$$\|A - P_\pi B P_\pi^\top\|_F^2 \;=\; \|A\|_F^2 + \|B\|_F^2 - 2\,\operatorname{tr}\big(P_\pi^\top A\, P_\pi B\big),$$

so $\hat{\pi}_{\text{MLE}} = \arg\max_\pi \operatorname{tr}(P_\pi^\top A\, P_\pi B)$, which is precisely the QAP objective $f(\pi)$ solved exactly by the Fourier B&B.



(a) Fiedler seriation

(b) SABRE seriation

(c) Order derivation (AOD)

(d) Least squares (ALS)

Fig. 1: Success rate p_{succ} vs. shadowing level σ/α under LNS (cyan) and vsJPLM d=0.4 (orange). Node counts (from top down) $n \in \{9, 11, 15, 21, 31, 41, 51\}$, except ALS with $n \in \{9, 11, 15\}$.

C. Correlated shadowing: the singular-covariance MLE

We vectorize the $m = \binom{n}{2}$ upper-triangular entries: $a = \mathrm{vec}(A)$, $b = \mathrm{vec}(B)$, and $\mu_\pi = \mathrm{vec}(P_\pi B P_\pi^\top) = \widetilde{P}_\pi\, b$, where $\widetilde{P}_\pi$ is the permutation that π induces on the m links, so the model reads $a = \mu_{\pi^\star} + N$ with $N \sim \mathcal{N}(0, K)$ *singular* of rank $n-1 < m$. We sketch the consequences here and refer to [5] for the derivation.

The rank deficiency, not the particular K, is the whole mechanism. Every realization of N lies in the column space $\mathrm{ran}(K)$, so the orthogonal projector $Q = I - KK^{+}$ onto the complementary *noise-free* subspace $\ker K$, of dimension $\frac{(n-1)(n-2)}{2}$, annihilates the noise identically: $Q\,N = 0$ for every realization. A singular normal is supported on the affine flat $\mu_\pi + \mathrm{ran}(K)$, so its likelihood vanishes unless $Q(a-\mu_\pi) = 0$: the MLE is feasibility followed by a generalized least-squares fit. But the residual is *exactly* zero at the truth — its mean part trivially, its noise part by $Q\,N = 0$ — so the fit is vacuous, the minimizer is independent of the norm, and the MLE collapses to the Q-projection estimator

$$\hat{\pi} = \arg\min_{\pi} \left\| Q(\pi)(a-\mu_\pi) \right\|^2 = \arg\min_{\pi} \left\| Q_{\mathrm{pos}}\,(\widetilde{P}_\pi^\top a - b) \right\|^2. \tag{3}$$

The second form uses a single canonical kernel: with K_{pos} the covariance and Q_{pos} its kernel projector for the natural order $1, \ldots, n$, hypothesis π merely re-indexes the same physical shadowing, $Q(\pi) = \widetilde{P}_\pi\, Q_{\mathrm{pos}}\, \widetilde{P}_\pi^\top$, so the covariance is never re-estimated per hypothesis and π acts on the data alone. Recovery is exact and independent of the noise level (the infinite Le Cam threshold of Remark 1).

Expanded against a basis $v_1, \ldots, v_{(n-1)(n-2)/2}$ of $\ker K_{\mathrm{pos}} = \mathrm{ran}\, Q_{\mathrm{pos}}$, the residual

$$r_\ell(\pi) = v_\ell \cdot (\widetilde{P}_\pi^\top a - b) = \sum_{i<j} v_\ell[\{i,j\}]\, a_{\{\pi(i),\pi(j)\}} - v_\ell \cdot b \tag{4}$$

is manifestly second-order, a function of the pairs $\{\pi(i), \pi(j)\}$. Braces denote unordered pairs: m-vectors are indexed by the links $\{i,j\}$, and $a_{\{\pi(i),\pi(j)\}}$ is well

defined by the symmetry of A. The problem is thus solved by a constraint-propagation branch-and-bound, which we call the *kernel least-squares* algorithm (AKLS): the truth satisfies every $r_\ell = 0$, so a partial assignment is pruned as soon as one r_ℓ is bounded away from zero over its completions.

Theorem 1. *For every singular correlated shadowing model of Section III-C, the residuals* (4) *vanish, almost surely and at every noise level, only at the true ordering and its reversal, so the Q-projection* (3) *is the MLE and identifies the line up to reversal.*

The proof is given in [5].

Remark 1 (infinite Le Cam threshold). For every singular model of Section III-C, distinct orderings induce mutually singular observation laws, so the Le Cam two-point method returns only the trivial bound, and by Theorem 1 the recovery threshold in σ/α is infinite [24, §2.3, Thm. 2.2]. This is a property of the exact kernel: with an approximate one, recovery degrades gracefully rather than staying exact (Fig. 2b).

V. Recovery Experiments

Figure 1 is the centerpiece: it sweeps the exact-recovery rate against the noise level σ/α for all four algorithms — Fiedler, SABRE, AOD, and ALS — under both uncorrelated (LNS) and correlated (vsJPLM) shadowing, across a range of node counts. We then compare the exact solvers' runtimes.

The cases $n = 3, 5, 7$ have been studied in great detail in [4], so we focus mainly on odd $n \geq 9$. As in [4], for correlated shadowing $d = 6.4$ leads to worse success rates than $d = 0.4$; we therefore present $d = 0.4$ (very high correlation), the favorable regime that flatters AOD and Fiedler seriation. We cap ALS at $n = 15$: it may run even at $n = 51$ under low noise, but the runtime can explode on a bad sample, which makes a systematic evaluation at the $N_{\text{runs}} = 10^4$ runs we use for ALS impossible beyond $n = 15$ on our hardware.

A. Heuristics and ALS

Figure 1 compares the success rates of four algorithms, three of them polynomial-time; the fourth is the least-squares estimator ALS, solved exactly by the Fourier B&B. The two seriation algorithms have similarly shaped performance curves. Under correlated shadowing SABRE dominates Fiedler at every n; under LNS the two cross over near n=15, SABRE leading below it and Fiedler above. Against order derivation under LNS the picture is mixed: seriation is better at the lower node counts, AOD at $n = 21, 31, 41, 51$. ALS dominates all other algorithms under LNS for $n = 9, 11, 15$, where its runtime is practical. Under correlated shadowing AOD is best up to $\sigma/\alpha \approx 4.3$, when ALS takes the lead ($n = 9, 11, 15$ only) — ALS is not the ML estimator there, which is why AOD can beat it.

B. Exact Solver Runtime Comparison

At matched $(n, \sigma/\alpha)$ the Fourier B&B is up to $60\times$ faster than the GLB solver [11] at n=9 (Fig. 2a). Both runtimes grow with noise, but the Fourier B&B has the steeper SNR dependence ($\approx 10\times$ swing against GLB's $\approx 1.9\times$), so the speedup $t_{\text{GLB}}/t_{\text{B\&B}}$ (dashed) peaks at $\approx 60\times$ near σ/α=0.75 and decays to $\approx 9\times$ at σ/α=7.0. The Fiedler and AOD heuristics are noise-independent at $\approx 7\,\mu$s, AOD the fastest of all methods compared. The exact solvers agree in all but 3 cells, where a GLB exactness defect (sub-optimal return on $\sim 10^{-4}$ instances) invalidates the local speedup. The SDP relaxation [10] agrees on sampled instances but is intractable at this size, and its margin widens with n: $\approx 10^4\times$ the Fourier B&B already at n=7, and $\approx 10^5\times$ at n=9, where a single solve takes $\approx$5–8 min.

C. MLE for correlated shadowing

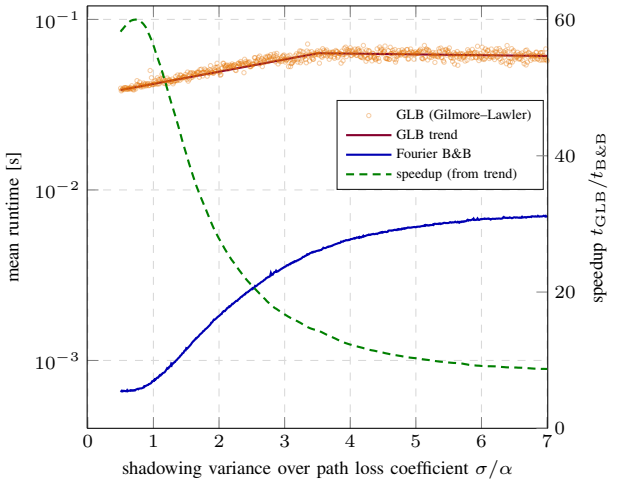
1) Perturbation: AKLS, the branch-and-bound solver of our kernel MLE, restores the order of the line topology perfectly at $\sigma/\alpha = 7.0$. By Theorem 1 this recovery is exact at *every* noise level; what the experiment adds is the practical fact that the solver actually *terminates* and returns $\pi^\star$ up to reversal, which we have tested up to $n = 15$.

A natural worry is that this 100% rate is an artifact of the perfectly collinear line, on which K is exactly singular. To probe it we deliberately break the line: each of the n=5 nodes is displaced off it by an independent $\mathcal{N}(0, v)$ in each coordinate, which makes the covariance full-rank and destroys the exact kernel. Figure 2b ($\sigma/\alpha = 7$) plots the recovery rate against v for the analytic singular kernel of the nominal line and for two kernels *built* by sampling, where the near-null space is read off an empirical covariance formed from s noise realizations. The exact-line assumption is therefore not essential.

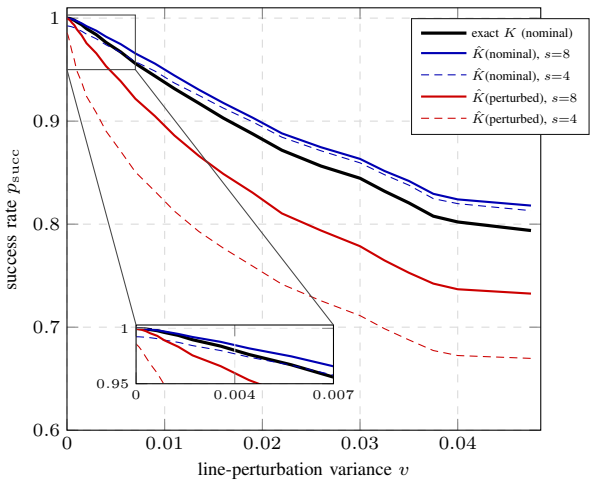
2) Model Comparison: Since recovery depends only on the kernel $\ker K$, we compare the singular models of Section III-C by their kernels rather than by their full covariances; the detailed comparison is given in [5]. The kernel's *dimension* is fixed by the collinear topology and the segment construction, not by the covariance, so it is the same for all of them; what distinguishes them is its *orientation*. That kernel is $D^{-1}\text{ran}(E)^\perp$, so it sees C only through the per-link variances that set D: all un-scaled models share one kernel whatever their correlation function, and variance scaling is what rotates it.

VI. Conclusion

We treated wireless node ranking as maximum-likelihood estimation of the ordering of a line topology from RSSI measurements. Under uncorrelated log-normal shadowing the MLE is the least-squares QAP, which the non-abelian Fourier B&B (Kondor bound) solves exactly, up to $60\times$ faster than the GLB. We also identified node ranking with *seriation* — a link between the wireless and mathematical-statistics communities that appears to be new — and used the Fiedler and SABRE seriation estimators as polynomial-time benchmarks. The central result concerns *correlated* shadowing: when the covariance is singular, as for the (variance-scaled) joint path-loss model, the MLE reduces to a projection onto the covariance kernel and recovers the ordering exactly, independently of the shadowing



(a) Exact-solver runtime over shadowing level σ/α at n=9 under uncorrelated (LNS, Gaussian) noise; mean over N_{runs}=10^4 Fourier and N_{glb}=20 GLB instances per level, log-scaled. The solid purple line is a linear trend fit to the GLB runtime; the speedup (dashed, right, linear axis) is that trend over the moving-average-filtered Fourier B&B runtime.



(b) Recovery rate vs. line-perturbation variance v (n=5, vsJPLM d=0.4, $\sigma/\alpha = 7$): the exact singular kernel of the nominal line against approximate kernels $\hat{K}$ sampled from $s \in \{4, 8\}$ realizations, on the nominal line and on the perturbed array, whose full-rank covariance admits no exact kernel.

Fig. 2: Left: exact-solver runtime at n=9 under LNS. Right: kernel sample efficiency under line perturbation (n=5, vsJPLM).

variance (Theorem 1). Our solver confirms this 100% recovery up to $\sigma/\alpha = 7$, and an approximate, sampled kernel recovers nearly as well when the line is perturbed.

Future work. Several directions remain open: a polynomial-time approximation of the kernel MLE, presumably pairing the projection Q with a seriation heuristic (Fiedler, SABRE); a classification of the recovery kernels; field experiments on a real line of nodes; and an extension beyond the line, since a grid or other structured array built from the same segment construction likewise has a singular covariance with a nontrivial kernel.

References

[1] R. Kondor, "A Fourier space algorithm for solving quadratic assignment problems," in *Proc. ACM–SIAM Symp. Discrete Algorithms (SODA)*, pp. 1017–1028, 2010.

[2] M. Clausen, "Fast generalized Fourier transforms," *Theoretical Computer Science*, vol. 67, no. 1, pp. 55–63, 1989.

[3] H. Frey and D. Schneider, "Determining the center of a line topology with signal strength measurements under correlated log-normal shadowing: Analyzing the three node case," in *Proc. ACM MSWiM*, 2021, pp. 229–238. DOI: 10.1145/3479239.3485711.

[4] D. Schneider and H. Frey, "Determining the ordering of a line topology under correlated shadowing and fast fading," in *Proc. ACM MSWiM*, Montréal, Oct.–Nov. 2023, pp. 1–10. DOI: 10.1145/3616388.3617538.

[5] D. Schneider and H. Frey, "Wireless node ranking via Fourier-driven maximum-likelihood estimation (extended version)," arXiv preprint, in preparation.

[6] X. Zhu, X. Wu, and G. Chen, "Relative localization for wireless sensor networks with linear topology," *Computer Communications*, vol. 36, no. 15, pp. 1581–1591, 2013.

[7] X. Zhu and G. Chen, "Spatial ordering derivation for one-dimensional wireless sensor networks," in *Proc. IEEE Int. Symp. Parallel and Distributed Processing with Applications (ISPA)*, 2011. (Original "order derivation"; extended journal version is [6].)

[8] W. S. Robinson, "A method for chronologically ordering archaeological deposits," *American Antiquity*, vol. 16, no. 4, pp. 293–301, 1951.

[9] J. E. Atkins, E. G. Boman, and B. Hendrickson, "A spectral algorithm for seriation and the consecutive ones problem," *SIAM J. Comput.*, vol. 28, no. 1, pp. 297–310, 1998.

[10] Q. Zhao, S. E. Karisch, F. Rendl, and H. Wolkowicz, "Semidefinite programming relaxations for the quadratic assignment problem," *J. Combinatorial Optimization*, vol. 2, pp. 71–109, 1998.

[11] P. C. Gilmore, "Optimal and suboptimal algorithms for the quadratic assignment problem," *J. SIAM*, vol. 10, no. 2, pp. 305–313, 1962.

[12] M. Gudmundson, "Correlation model for shadow fading in mobile radio systems," *Electronics Letters*, vol. 27, no. 23, pp. 2145–2146, Nov. 1991.

[13] P. Agrawal and N. Patwari, "Correlated link shadow fading in multi-hop wireless networks," *IEEE Trans. Wireless Commun.*, vol. 8, no. 8, pp. 4024–4036, Aug. 2009.

[14] E. M. Loiola, N. M. de Abreu, P. O. Boaventura-Netto, P. Hahn, and T. Querido, "A survey for the quadratic assignment problem," *Eur. J. Oper. Res.*, vol. 176, no. 2, pp. 657–690, 2007.

[15] Z. Lotker, M. M. de Albeniz, and S. Pérennès, "Range-free ranking in sensors networks and its applications to localization," in *Ad-Hoc, Mobile, and Wireless Networks (ADHOC-NOW)*, LNCS vol. 3158, pp. 158–171. Springer, 2004.

[16] R. D. Komguem, R. Stanica, M. Tchuente, and F. Valois, "Node ranking in wireless sensor networks with linear topology," in *Proc. IFIP Wireless Days*, Porto, 2017.

[17] M. O. Ergin, V. Handziski, A. Behboodi, and A. Wolisz, "Determining node sequence in a linear configuration," in *Proc. IEEE IPIN*, Busan, Oct. 2014.

[18] M. O. Ergin, *Node Position Discovery in Wireless Sensor Networks*, Ph.D. dissertation, Technische Universität Berlin, 2018.

[19] M. M. Abdellatif, J. M. Oliveira, and M. Ricardo, "The self-configuration of nodes using RSSI in a dense wireless sensor network," *Telecommunication Systems*, vol. 62, no. 4, pp. 695–709, 2016.

[20] F. Fogel, R. Jenatton, F. Bach, and A. d'Aspremont, "Convex relaxations for permutation problems," in *Proc. NIPS*, 2013. Extended in *SIAM J. Matrix Anal. Appl.*, vol. 36, no. 4, 2015.

[21] R. Kondor and K. M. Borgwardt, "The skew spectrum of graphs," in *Proc. Int. Conf. Machine Learning (ICML)*, Helsinki, 2008.

[22] Y. Issartel, C. Giraud, and N. Verzelen, "Minimax optimal seriation in polynomial time," arXiv:2405.08747 [math.ST], 2024.

[23] A. Natik and A. Smith, "Consistency of spectral seriation," arXiv:2112.04408 [math.ST], 2021.

[24] A. B. Tsybakov, *Introduction to Nonparametric Estimation*. Springer, 2009.