Optimal Placement of Geomagnetically Induced Current Monitors for Geoelectric-Field State Estimation in Electric Power Systems

Jonathan M. Snodgrass
Texas A&M University
snodgrass@tamu.edu

Nicole C. LoGiudice
Texas A&M University
nicolelogiudice30@tamu.edu

Thomas J. Overbye
Texas A&M University
overbye@tamu.edu

Abstract

Geomagnetic disturbances (GMDs) drive geomagnetically induced currents (GICs) in the bulk electric power system, saturating transformers and contributing to wide-area outages. The geoelectric field responsible for these currents is highly uncertain and rarely observed directly. Because the quasi-DC relationship between the surface field and transformer-neutral GICs is linear, a weighted least squares (WLS) estimator can reconstruct the regional field from a limited set of neutral-current measurements—using the grid itself as a sensor. The accuracy of that estimate depends on where the GIC monitors are placed, yet existing studies place them arbitrarily. We cast monitor placement as a convex A-optimal experiment-design problem: the binary selection is relaxed to continuous weights, the covariance trace is minimized as a semidefinite program, and the relaxed solution is rounded to a fixed budget, which yields a high-quality feasible monitor set rather than a certified global optimum. On the Illini 42-bus case across 36 synthetic storms, optimized placements achieve lower estimation cost than arbitrary ones at every monitor budget.

Keywords: Geomagnetic disturbances, geomagnetically induced currents, state estimation, optimal sensor placement, convex optimization

This paper was accepted for presentation at the 60th Hawaii International Conference on System Sciences (HICSS), Waikoloa Village, HI, January 2027; after the conference it will be available open access at https://hdl.handle.net/10125/39610, and at overbye.engr.tamu.edu/publications.

1. Introduction

Geomagnetic disturbances (GMDs) drive geomagnetically induced currents (GICs) in the bulk electric power system that can push transformers into saturation, raise their reactive-power demand, inject harmonics, and—in severe events—contribute to voltage collapse. A central difficulty in assessing and operating through a GMD is that the geoelectric field responsible for the GICs is highly uncertain and is directly measured at only a handful of locations. A recent line of work reframes the problem by treating the power grid itself as a distributed sensor: because the quasi-DC relationship between the surface electric field and transformer-neutral GICs is linear, a weighted least squares (WLS) state estimator can invert a limited set of neutral-current measurements to recover the regional electric field, and from it the GIC in every transformer and transmission line (Klauber et al., 2020; LoGiudice et al., 2024). The fidelity of that reconstruction, however, depends on which transformers are instrumented. This paper addresses that design question directly. Rather than treating the estimator's measurement weighting as a quantity fixed in advance, we treat it as a decision variable and solve a convex optimization that selects monitor locations so as to minimize the variance of the estimated field.

GMDs belong to the class of high-impact, low-frequency (HILF) events that occur rarely but stress the bulk power system out of all proportion to their frequency (North American Electric Reliability Corporation (NERC), 2010); the February 2021 Texas cold-weather event and the 2003 Northeast blackout are familiar weather- and grid-driven examples. GMDs are the space-weather member of this class, and although a truly extreme storm has not yet struck a modern interconnected grid, the historical record is sobering:

the March 1989 storm collapsed the Hydro-Québec system, leaving roughly six million customers without power (Boteler, 2019); the October 2003 "Halloween" storms caused a transformer-related blackout in Malmö, Sweden (Pulkkinen et al., 2005); and the May 2024 "Gannon" storm, the largest in two decades, produced measurable GICs and transformer harmonics (Clilverd et al., 2025).

The physical chain begins at the Sun: a coronal mass ejection perturbs the magnetosphere and drives ionospheric and auroral electrojet currents that produce a time-varying magnetic field at the Earth's surface (Boteler & Pirjola, 2017; Pirjola, 2002). By Faraday's law of induction, in differential (Maxwell–Faraday) form,

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{1}$$

the changing field $\mathbf{B}$ induces a geoelectric field $\mathbf{E}$ within the conductive Earth, on the order of a few volts per kilometer at the surface. On its own this is negligible, but high-voltage transmission lines span hundreds of kilometers and are tied to ground at each end through wye-connected transformer neutrals, completing a path for quasi-DC current; a long line in a uniform $5\ \mathrm{V/km}$ field can carry hundreds of amperes of GIC.

Once a GIC enters the network through a transformer's grounded neutral connection, it acts as a slowly varying DC bias on the $60\ \mathrm{Hz}$ exciting current, shifting the transformer operating point on its B–H curve and driving one half-cycle deep into saturation (Molinski, 2002). The resulting exciting current is strongly distorted; because the offset destroys the half-wave symmetry that ordinarily restricts power-system currents to odd harmonics, half-cycle saturation also injects *even* harmonics, most notably the second (Saleh et al., 2023). The consequences are twofold: the enlarged exciting current raises reactive-power consumption and depresses the voltage profile, while the harmonics can overheat capacitor banks and cause protective relays to misoperate (Clilverd et al., 2025; Molinski, 2002).

The core question in GMD analysis is operational—a storm is forecast, what should be done?—and answering it is complicated by two dominant uncertainties. The first is the Earth's conductivity: converting a measured magnetic field to a surface electric field relies on a magnetotelluric relationship through the Earth's surface impedance, evaluated with layered or three-dimensional conductivity models whose errors propagate directly into the inferred field (Fernberg, 2012; Kelbert et al., 2019). The second is the power-system data required for GIC analysis, which goes well beyond a conventional power flow: the transformer configuration and grounding arrangement, the substation grounding impedance, the DC winding and line resistances, and the geographic coordinates of every bus (which set how strongly the field couples to each line) are all needed (Kazerooni et al., 2017; Overbye et al., 2012). Because of this uncertainty, most studies adopt a benchmark field; the North American Electric Reliability Corporation (NERC) standard TPL-007 specifies an $8\ \mathrm{V/km}$ reference field with a $12\ \mathrm{V/km}$ hotspot, scaled by latitude and conductivity (North American Electric Reliability Corporation (NERC), 2016, 2023). A uniform benchmark, however, does not capture the spatial structure of a real event, whose magnitude and direction can vary substantially over short distances (Dehghanian et al., 2024), and Carrington-event reconstructions suggest localized fields of 10–$20\ \mathrm{V/km}$ (Winter et al., 2017).

These difficulties motivate a different question: rather than computing the field externally and pushing it through a network model, can the grid estimate the field directly? The linear, quasi-DC coupling makes this possible. Treating the regional field as the unknown state and the neutral-current readings as measurements, a WLS estimator recovers the field that best explains the data and, with it, the GIC in every branch (Juvekar et al., 2021; Klauber et al., 2020; LoGiudice et al., 2024). A few monitors per region—at least two, preferably three or more for redundancy and bad-data detection—make each region observable, and a companion manuscript extends the approach to a spatially continuous field (LoGiudice et al., 2026). The estimator is fast and scalable, making it attractive for real-time situational awareness during a storm.

This capability is immediately useful where monitors already exist, but it raises a planning question for those that do not. As utilities become more aware of GMD risk, programs to deploy additional transformer-neutral GIC monitors are underway, and operators must decide *where* to install them. Existing state-estimation studies sidestep this: monitors are placed arbitrarily or only to satisfy observability, with no principle for ranking one location above another by how much it improves the estimate. The broader GIC literature has studied where to place *blocking* or mitigation devices (Zhu & Overbye, 2015) but not where to place *measurements*; indeed, the originating estimator names optimal GIC-meter placement as future work (Klauber et al., 2020). Accuracy-driven placement is, however, mature in adjacent settings: optimal phasor measurement unit (PMU) and meter placement supply transferable formalisms (Gou, 2008; Xu & Abur, 2004), and the optimal design of experiments supplies the

optimality criteria and convex-relaxation machinery for choosing a fixed number of measurements that minimize estimation variance (Boyd & Vandenberghe, 2004; Joshi & Boyd, 2009; Nagata et al., 2021; Pukelsheim, 2006; Yu et al., 2018). To our knowledge, none has been applied to GIC monitoring, where the state is a distributed geoelectric field rather than bus voltages and angles. This paper closes that gap.

The contributions of this paper are as follows.

- We formulate GIC-monitor placement as an A-optimal experiment-design problem, treating the measurement-selection weights in the WLS field estimator as decision variables and minimizing the trace of the estimation-error covariance.
- We relax the binary selection to a continuous weight vector and cast the resulting problem as a convex semidefinite program (SDP), recovering a physical, fixed-budget monitor set by top-K rounding; the formulation also accommodates monitors that are already installed.
- We evaluate optimized placements against arbitrary placements across a parameterized family of synthetic storm conditions on the Illini 42-bus synthetic GMD case (Hutchins & Overbye, 2016), comparing estimation cost against a χ^2 acceptance bound and examining the relaxed weights as a ranking of the most informative monitor locations.

The remainder of the paper is organized as follows. Section 2 reviews GIC modeling, geoelectric-field state estimation, and optimal experiment design. Section 3 develops the measurement model, the placement optimization, and the rounding procedure. Section 4 describes the test system and storm scenarios, Section 5 presents results, and Section 6 concludes.

2. Background and Related Work

2.1. GIC Modeling

In the quasi-DC regime relevant to GMDs, a power network is modeled as a resistive circuit: transmission lines by their series DC resistances, transformer windings by their winding resistances, and each substation by a grounding resistance tying the network to the remote Earth (Boteler & Pirjola, 2017; Lehtinen & Pirjola, 1985). The geoelectric field appears as DC voltage sources in series with the lines; integrating the field along a line gives the induced voltage, and the resulting nodal current injections drive GICs through the grounded neutrals (Horton et al., 2012; Overbye et al., 2012). Substation grounding resistance governs how much current the network draws from the Earth at each node (Bui et al., 2013), while the transformer configuration governs how neutral current maps to the effective per-phase current that determines reactive-power loss. These dependencies make GIC analysis more data-intensive than a conventional power flow, but the governing relationships remain linear—which is what makes both rapid GIC calculation and the estimator below tractable.

2.2. Geoelectric-Field State Estimation

State estimation infers unknown quantities from a limited, noisy set of measurements by minimizing a weighted residual between measured and modeled values (Monticelli, 1999). For GMDs the estimated quantity is not bus voltages and angles but the regional geoelectric field. Because the field is continuous over geography, it is made estimable by discretizing the footprint into regions, each carrying an eastward and a northward component—conceptually similar to gridded reanalysis in the atmospheric sciences, where sparse measurements and a physical model jointly estimate a field over space. The measurements are transformer-neutral GICs and, when available, magnetometer-derived field values. Given a linear map from the field to these measurements, a WLS estimator returns the best-fit field, and the cost at the solution—compared against a chi-squared threshold—indicates whether the assumed field structure is statistically consistent with the data (LoGiudice et al., 2024). This work builds directly on that estimator and the GIC-inclusive literature preceding it (Juvekar et al., 2021; Klauber et al., 2020).

2.3. Measurement Placement for the Estimator

In prior studies the GIC measurements feeding the estimator are placed without optimization—arbitrarily, distributed across regions for balance, or only to make the system observable. Observability-driven placement, familiar from PMU siting (Gou, 2008; Xu & Abur, 2004), guarantees a unique estimate but says nothing about its *quality*: two observable placements of the same size can yield very different variances. What is missing is a principle that ranks candidate locations by how much each reduces the uncertainty of the estimated field—precisely the question addressed by the optimal design of experiments.

2.4. Optimal Experiment Design

The optimal design of experiments selects measurements to make an estimate as informative as possible, summarizing the estimation-error covariance by a scalar criterion: A-optimality minimizes its trace (total variance), D-optimality its determinant, and E-optimality its largest eigenvalue (Boyd & Vandenberghe, 2004; Pukelsheim, 2006). Selecting a fixed number of measurements is combinatorial, but relaxing the binary selection to a continuous weight yields a convex problem that is solved efficiently and rounded back to a discrete subset (Joshi & Boyd, 2009; Yu et al., 2018); A-optimal placement in particular has well-developed convex and greedy algorithms (Aarset, 2024; Nagata et al., 2021). We adopt the A-optimal criterion—total variance is the natural target for field reconstruction—and apply this relaxation-and-rounding strategy to GIC-monitor placement, which has not previously appeared in the GMD literature.

3. Methodology

This section develops the measurement model that links the geoelectric field to GIC and magnetometer observations, states the WLS field estimator, and then poses monitor placement as a convex experiment-design problem over that estimator. The overall data flow is summarized in Figure 1: the grid and storm models yield a storm-independent sensitivity matrix; a trace-minimizing convex program selects monitor weights; and the selected monitors feed the WLS estimator that reconstructs the field; magnetometer-derived measurements, when available, are treated as fixed and enter only at the estimation stage. Equations (2)–(8) follow the formulation of LoGiudice et al. (2024); the placement formulation in Section 3.3 is new.

3.1. State Vector and Measurement Model

Let the geoelectric field be discretized into r regions, each described by an eastward component E_x and a northward component E_y. The state vector collects these components,

$$\boldsymbol{E} = \begin{bmatrix} E_{x1} & E_{y1} & \cdots & E_{xr} & E_{yr} \end{bmatrix}^{\top} \in \mathbb{R}^{2r}, \quad (2)$$

and the formulation accommodates any linear field model (for example, a single uniform region, or a rotated directional basis) by redefining $\boldsymbol{E}$ and the matrices below.

The field induces voltages along the lines and, through the resistive network, approximately constant current injections at the grounded transformers. With n buses and substations, the bus voltages are

$$\boldsymbol{V} = \boldsymbol{G}^{-1}\boldsymbol{H}\boldsymbol{E}, \quad (3)$$

where $\boldsymbol{G} \in \mathbb{R}^{n\times n}$ is the network conductance matrix (line and grounding conductances on the diagonal, negative mutual conductances off-diagonal) and $\boldsymbol{H} \in \mathbb{R}^{n\times 2r}$ maps the regional field components to nodal current injections, accounting for line length, orientation, region, and conductance. When the network is not fully grounded, $\boldsymbol{G}$ is singular and a symmetric pseudoinverse is used in place of $\boldsymbol{G}^{-1}$.

Each GIC monitor-instrumented transformer contributes a row to $\boldsymbol{\Phi}_n \in \mathbb{R}^{l\times n}$, with nonzero entries at its grounded neutrals weighted by the effective conductance between them, so that the l measured neutral currents are

$$\boldsymbol{I}_{\mathrm{GIC}} = \boldsymbol{\Phi}_n\boldsymbol{V} = \boldsymbol{\Phi}_n\boldsymbol{G}^{-1}\boldsymbol{H}\boldsymbol{E} = \boldsymbol{A}_{\mathrm{GIC}}\boldsymbol{E}, \quad (4)$$

which defines the GIC sensitivity matrix $\boldsymbol{A}_{\mathrm{GIC}} = \boldsymbol{\Phi}_n\boldsymbol{G}^{-1}\boldsymbol{H} \in \mathbb{R}^{l\times 2r}$. This matrix depends only on the network; it is independent of the storm.

Magnetometers, when available, provide direct field measurements $\boldsymbol{Z}_E \in \mathbb{R}^{2m}$ at m sites, mapped to

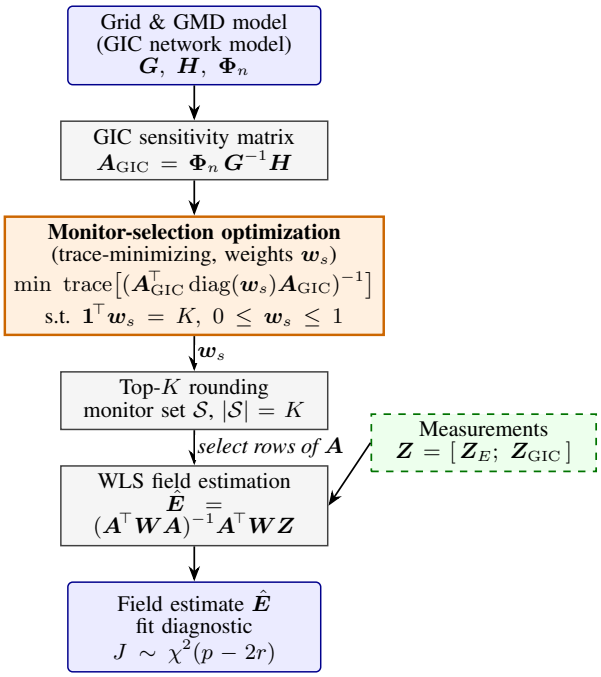

Figure 1. Conceptual block diagram of the GIC-monitor placement framework, in which a storm-independent sensitivity matrix feeds a convex program whose relaxed weights are rounded to a K-monitor set for WLS estimation of the regional field.

the field components by $\boldsymbol{A}_E \in \mathbb{R}^{2m\times 2r}$. Ground magnetometer networks supplying measurements of exactly this kind have been deployed and applied to grid monitoring during GMDs (Shetye et al., 2021), so such data can reasonably be treated as an asset already in place. Stacking the magnetometer and GIC measurements gives the full linear model

$$\boldsymbol{Z} = \begin{bmatrix} \boldsymbol{Z}_E \\ \boldsymbol{Z}_{\mathrm{GIC}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{A}_E \\ \boldsymbol{A}_{\mathrm{GIC}} \end{bmatrix} \boldsymbol{E} = \boldsymbol{A}\boldsymbol{E}. \tag{5}$$

Here $\boldsymbol{Z} \in \mathbb{R}^p$ collects $p = 2m + l$ measurements. Magnetometers are not required: the estimator can operate on GIC data alone, with magnetometer measurements added as available.

3.2. Weighted Least Squares Estimator

Each measurement is assigned a variance, assembled into the diagonal covariance $\boldsymbol{R} \in \mathbb{R}^{p\times p}$, and the weight matrix is its inverse, $\boldsymbol{W} = \boldsymbol{R}^{-1}$, so that more trusted measurements receive larger weight. The estimator minimizes

$$J = (\boldsymbol{Z} - \boldsymbol{A}\boldsymbol{E})^\top \boldsymbol{W}(\boldsymbol{Z} - \boldsymbol{A}\boldsymbol{E}), \tag{6}$$

whose minimizer is

$$\hat{\boldsymbol{E}} = (\boldsymbol{A}^\top \boldsymbol{W}\boldsymbol{A})^{-1}\boldsymbol{A}^\top \boldsymbol{W}\boldsymbol{Z}, \tag{7}$$

provided $\boldsymbol{A}$ has full column rank, i.e., the system is observable. The cost at the solution provides a statistical check on the assumed field model: under Gaussian measurement noise,

$$J \sim \chi^2(p - 2r), \tag{8}$$

and a value below the 95th percentile indicates that the model is consistent with the data. The information matrix $\boldsymbol{A}^\top \boldsymbol{W}\boldsymbol{A}$ in (7) is the quantity whose inverse—the estimation-error covariance—we will minimize through monitor placement.

3.3. Monitor Placement as Convex Experiment Design

We now choose which transformers to instrument. Let there be n candidate GIC-monitor sites, and let $\boldsymbol{w}_s \in \{0,1\}^n$ be a binary selection vector, with $(w_s)_i = 1$ if site i is instrumented. Selecting K monitors to minimize the total estimation variance (the A-optimal criterion) gives the combinatorial program

$$\underset{\boldsymbol{w}_s \in \{0,1\}^n}{\text{minimize}} \quad \operatorname{trace}\Big[\big(\boldsymbol{A}_{\mathrm{GIC}}^\top \operatorname{diag}(\boldsymbol{w}_s)\, \boldsymbol{A}_{\mathrm{GIC}}\big)^{-1}\Big] \tag{9a}$$

$$\text{subject to} \quad \mathbf{1}^\top \boldsymbol{w}_s = K, \tag{9b}$$

$$(w_s)_i \in \{0,1\}, \quad i = 1,\dots,n, \tag{9c}$$

where $\operatorname{diag}(\boldsymbol{w}_s)$ acts as a measurement-selection weighting on the GIC sensitivity matrix. (Magnetometer terms, treated as fixed, can be added to the information matrix as a constant $\boldsymbol{A}_E^\top \boldsymbol{W}_E \boldsymbol{A}_E$; here the design ranges over GIC monitors only.) Problem (9) requires evaluating $\binom{n}{K}$ subsets and is intractable at utility scale.

Relaxing the binary constraint to the interval $0 \le (w_s)_i \le 1$ makes the feasible set convex while preserving the objective, which is convex in $\boldsymbol{w}_s$ (Boyd & Vandenberghe, 2004; Joshi & Boyd, 2009). The matrix-inverse objective is then linearized exactly by introducing an auxiliary matrix $\boldsymbol{X} \in \mathbb{S}^{2r}$ that upper-bounds the error covariance and applying a Schur complement, which converts the condition $\boldsymbol{X} \succeq (\boldsymbol{A}_{\mathrm{GIC}}^\top \operatorname{diag}(\boldsymbol{w}_s)\boldsymbol{A}_{\mathrm{GIC}})^{-1}$ into a linear matrix inequality. The result is the semidefinite program

$$\underset{\substack{\boldsymbol{w}_s \in \mathbb{R}^n \\ \boldsymbol{X} \in \mathbb{S}^{2r}}}{\text{minimize}} \quad \operatorname{trace}(\boldsymbol{X}) \tag{10a}$$

$$\text{subject to} \quad \begin{bmatrix} \boldsymbol{X} & \boldsymbol{I} \\ \boldsymbol{I} & \boldsymbol{A}_{\mathrm{GIC}}^\top \operatorname{diag}(\boldsymbol{w}_s)\, \boldsymbol{A}_{\mathrm{GIC}} \end{bmatrix} \succeq 0, \tag{10b}$$

$$\mathbf{1}^\top \boldsymbol{w}_s = K, \tag{10c}$$

$$0 \le (w_s)_i \le 1, i = 1,\dots,n. \tag{10d}$$

which is solvable in polynomial time by standard conic solvers. Minimizing $\operatorname{trace}(\boldsymbol{X})$ subject to the matrix inequality recovers the relaxed A-optimal objective of (9). If certain monitors are already installed, their weights are simply fixed to one, $(w_s)_i = 1$, and the program optimizes the placement of the remaining budget around them.

The cost of solving (10) is set not by the number of buses but by the field discretization and the number of candidate monitors. The matrix variable $\boldsymbol{X} \in \mathbb{S}^{2r}$ and the information matrix $\boldsymbol{A}_{\mathrm{GIC}}^\top \operatorname{diag}(\boldsymbol{w}_s)\boldsymbol{A}_{\mathrm{GIC}}$ are both $2r \times 2r$, where $2r$ is twice the number of field regions—typically only a handful—so the semidefinite cone stays small even for an interconnection-sized grid. The selection vector $\boldsymbol{w}_s \in \mathbb{R}^n$, with its budget and

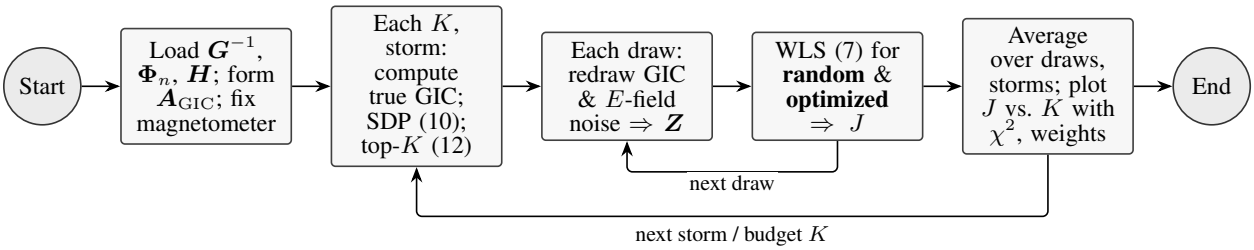

Figure 2. Procedure for comparing optimized and random monitor placement across budgets and storm conditions, with an inner Monte Carlo loop over measurement-noise draws.

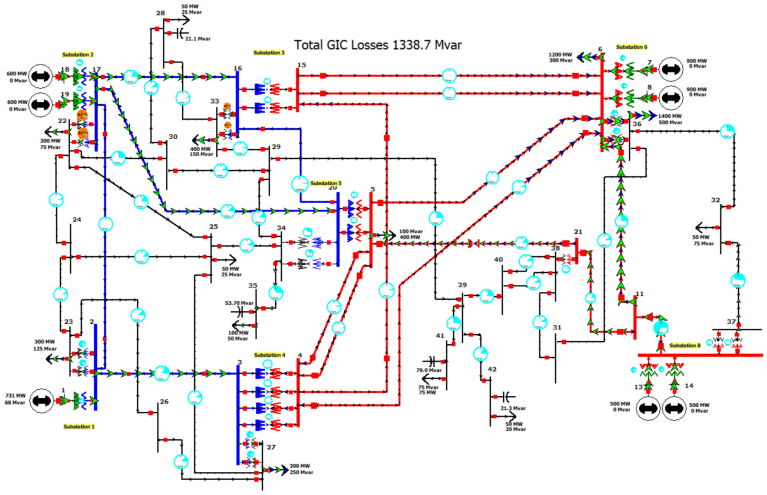

Figure 3. One-line diagram of the Illini 42-bus synthetic GMD case, in which each GIC transformer is a candidate monitor location.

box constraints, grows only linearly in the number of candidate sites n. System size enters just once, in forming the storm-independent sensitivity matrix $\boldsymbol{A}_{\mathrm{GIC}} = \boldsymbol{\Phi}_n \boldsymbol{G}^{-1} \boldsymbol{H}$, which is a single sparse network solve of the kind production GIC tools already perform at continental scale; once built, $\boldsymbol{A}_{\mathrm{GIC}}$ is reused across every budget and storm. The candidate set is moreover far smaller than the bus count: only wye-grounded transformers carry measurable neutral GIC and can host a monitor, so n counts GIC transformers rather than buses, and it can be pruned further—for example, by dropping transformers whose rows of $\boldsymbol{A}_{\mathrm{GIC}}$ have negligible norm, or by restricting the design to a region of interest. With utilities realistically deploying only a small number of monitors per region, the effective problem entering (10) remains modest even on a large system, which suggests the formulation transfers well beyond the 42-bus case studied here.

3.4. Recovering a Physical Monitor Set

The relaxed program returns fractional weights $w_s^* \in [0,1]^n$ rather than a binary selection. For a conference-scale, fast, and utility-deployable procedure we round by ranking: the weights are sorted in descending order,

$$w^*_{(1)} \geq w^*_{(2)} \geq \cdots \geq w^*_{(n)}, \tag{11}$$

and the K largest-weighted sites are installed,

$$\mathcal{S} = \left\{ i : w_i^* \geq w^*_{(K)} \right\}. \tag{12}$$

Top-K rounding is inexpensive and gives operators a direct ranking of candidate locations. Because it can split weight across highly correlated candidates, a more rigorous sparsification—for example, an ℓ_1 regularizer that drives the weights toward $\{0,1\}$ directly—is a natural refinement and is left to future work.

The relationship between (9) and (10) warrants a precise statement. Because (10) optimizes over a superset of the binary designs, its optimal value

is a lower bound on the A-optimal value attainable by any K-element selection, and the rounded set $\mathcal{S}$ of (12) is a feasible point whose objective can only be larger. The size of that gap is not certified here, so the procedure should be read as returning a high-quality feasible design rather than a proven global optimum; the term *A-optimal* names the design criterion being minimized—the trace of the estimation-error covariance—and not a guarantee about the selection that survives rounding (Joshi & Boyd, 2009). The looseness is moreover not uniform in K: where the relaxed weights saturate toward $\{0, 1\}$ the rounding is nearly exact, and where they remain diffuse it is at its worst, a pattern that the results of Section 5 make visible. Certifying the gap—by exhaustive enumeration on a case of this size, or through the ℓ_1 formulation above, which dispenses with the rounding step altogether—is left to future work.

The complete experimental procedure, in which the relaxed program is solved once per storm condition and the selected and randomly chosen monitor sets are then compared under repeated measurement noise, is shown in Figure 2.

4. Test System and Storm Scenarios

We evaluate the placement framework on the Illini 42-bus synthetic GMD case, a fictitious but realistic test system with publicly available GIC data that can be downloaded from the Texas A&M University electric grid test case repository (Kunkolienkar et al., 2024).[1] The case spans eight substations and includes the transformer configurations, DC winding resistance, neutral grounding resistance (if applicable), substation grounding resistances, and bus coordinates required for GIC analysis; its one-line diagram, annotated with the per-transformer GIC contributions that sum to the system reactive-power loss, is shown in Figure 3. Each GIC transformer in the case is treated as a candidate monitor site, and we test monitor budgets $K \in \{2, 4, 8, 10, 12, 14, 18, 20, 22\}$.

Because this study isolates the placement question, the geoelectric field is modeled as a single region: the entire footprint is described by one eastward and one northward component, so the state in (2) has dimension $2r = 2$. This is the coarsest field model the estimator admits; it keeps the comparison transparent but, as Section 5 shows, it also means that monitors in different parts of the footprint see a field that a single uniform state cannot fully reconcile.

A single magnetometer is included in every estimate, representing a regional field measurement already available to the operator rather than a sensor being sited here. Its location is drawn once from a seeded uniform distribution over the bounding box of the transformer coordinates and then held fixed across every budget, storm condition, and noise draw, so that GIC-monitor placement is the only quantity that varies; drawing the site rather than choosing it also keeps it from being positioned to favor either placement strategy. It contributes two direct field measurements $\boldsymbol{Z}_E$. With one magnetometer ($m = 1$) and a single region ($r = 1$), the total measurement count is $p = 2 + K$ and the chi-squared degrees of freedom $p - 2r$ reduce to K. Treating the magnetometer as given, rather than placing magnetometers and GIC monitors jointly, keeps the present study focused on the transformer-instrumentation decision.

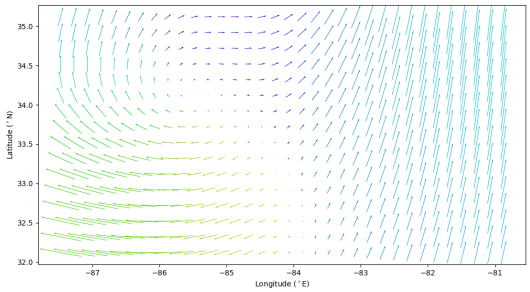

Figure 4. Example synthetic geoelectric field, one of the 36 storm conditions. Arrow length encodes the field magnitude and arrow color the field direction.

Measurement noise is added independently in each scenario: GIC readings are perturbed by zero-mean Gaussian noise with standard deviation 0.3 A and magnetometer field values with standard deviation 0.6 V/km. For each budget and storm condition the relaxed program (10) is solved once (using the splitting conic solver, SCS) and rounded to the top-K monitor set; the WLS estimator is then run over 20 Monte Carlo noise draws for both the optimized set and a uniformly random set of the same size, and the resulting costs are averaged over the draws and over the 36 conditions.

5. Results

Figure 5 and Table 1 report the dimensionless average WLS estimation cost J of (6) for the optimized and random monitor sets as the GIC-monitor budget K grows, together with the 5th- and 95th-percentile chi-squared bounds. The optimized placement is cheaper than the random placement at every budget. The advantage is dramatic at small budgets—at $K = 2$ the optimized cost is 2.23 against 5.66×10^2 for random,

[1] https://electricgrids.engr.tamu.edu/electric-grid-test-cases/illinigmd-42-hemp/

and at $K = 4$ it is 5.72 against 4.51×10^4—and remains substantial at larger budgets, for example 2.66×10^4 versus 1.44×10^5 at $K = 8$ and 3.78×10^5 versus 5.11×10^5 at $K = 20$.

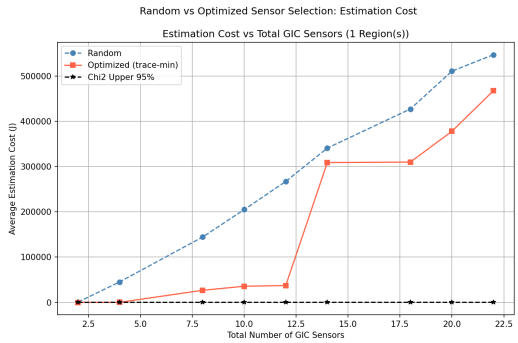

Figure 5. Average WLS estimation cost J versus the number of GIC monitors for random and optimized placement, shown with the 95th-percentile χ^2 bound.

Table 1. Average WLS estimation cost J for random (J_{rand}) versus optimized (J_{opt}) GIC-monitor placement on the Illini 42-bus case, with 5th- and 95th-percentile χ^2 bounds at degrees of freedom $\mathrm{dof} = p - 2r = K$.

K	J_{rand}	J_{opt}	$\chi^2_{0.95}$	$\chi^2_{0.05}$	dof
2	5.66×10^2	2.23×10^0	5.99	0.10	2
4	4.51×10^4	5.72×10^0	9.49	0.71	4
8	1.44×10^5	2.66×10^4	15.51	2.73	8
10	2.05×10^5	3.55×10^4	18.31	3.94	10
12	2.67×10^5	3.71×10^4	21.03	5.23	12
14	3.41×10^5	3.09×10^5	23.68	6.57	14
18	4.27×10^5	3.10×10^5	28.87	9.39	18
20	5.11×10^5	3.78×10^5	31.41	10.85	20
22	5.47×10^5	4.68×10^5	33.92	12.34	22

The chi-squared bounds put these numbers in context. At $K = 2$ and $K = 4$ the optimized cost falls inside the acceptance band between the 5th and 95th percentiles (0.10–5.99 and 0.71–9.49, respectively), so an optimized two- or four-monitor set leaves the single-region field model statistically consistent with the data; the corresponding random costs lie one to four orders of magnitude above the upper bound, indicating a placement whose data the model cannot explain.

Both costs nonetheless rise with K and, beyond the smallest budgets, climb above the acceptance band. This is a consequence of the single-region field model rather than of the estimator itself. The geoelectric field is spatially continuous, but here it is represented by one eastward/northward pair; as more monitors are added, they sample neutral currents driven by a field that varies across the footprint, and no single uniform state can reconcile all of them, so the weighted residual—and hence J—grows with the number of measurements. The optimized selection slows this growth by favoring the most informative, least redundant monitors, but it cannot remove a modeling bias intrinsic to estimating a varying field as one region.

The relaxed weights underlying these selections are shown in Figure 6, sorted in descending order with the top-K cutoff marked, averaged across the 36 storm conditions. At small budgets the weight is spread rather than concentrated: at $K = 2$ the largest weight is only about 0.30 and the remainder decay gradually, so no single monitor dominates and the optimizer hedges across mutually correlated neutral-current measurements. As the budget grows, the highest-ranked sites saturate toward one—by $K = 14$ and above, the leading weights sit near 1.0 before dropping sharply past the cutoff—so the ranking becomes crisp and top-K rounding cleanly recovers the high-weight set.

This progression also helps explain the optimized cost curve in Figure 5, which rises gently through $K = 12$ and then jumps near $K = 14$. When the relaxed weights are diffuse, top-K rounding must promote sites whose weights the relaxation deliberately held below one; forcing these correlated, lower-value monitors into the set is where rounding departs most from the relaxed optimum, consistent with the optimized cost leaving the chi-squared band as K increases. A sparsification that drives the weights toward $\{0, 1\}$ directly, rather than rounding after the fact, is the natural remedy and is discussed below.

The choice of baseline warrants comment. Random placement serves as the reference here because it stands in for what the GIC state-estimation literature actually does: monitors are placed arbitrarily, distributed across regions for balance, or chosen only to secure observability (Juvekar et al., 2021; Klauber et al., 2020; LoGiudice et al., 2024). Averaged over noise draws and over the 36 conditions, it therefore measures the expected performance of an uninformed choice—the position an operator occupies with no placement principle at all—rather than serving as an artificially weak comparator. Sharper baselines nonetheless exist: greedy forward selection on the A-optimal criterion, a ranking by the row norms of $\boldsymbol{A}_{\mathrm{GIC}}$ that instruments the transformers carrying the largest neutral currents, and a geographically spread placement of the kind planners already favor (Aarset, 2024; Nagata et al., 2021). Benchmarking against these would separate the contribution of the convex formulation itself from the more basic gain of avoiding a poor placement.

Taken together, the results show that the

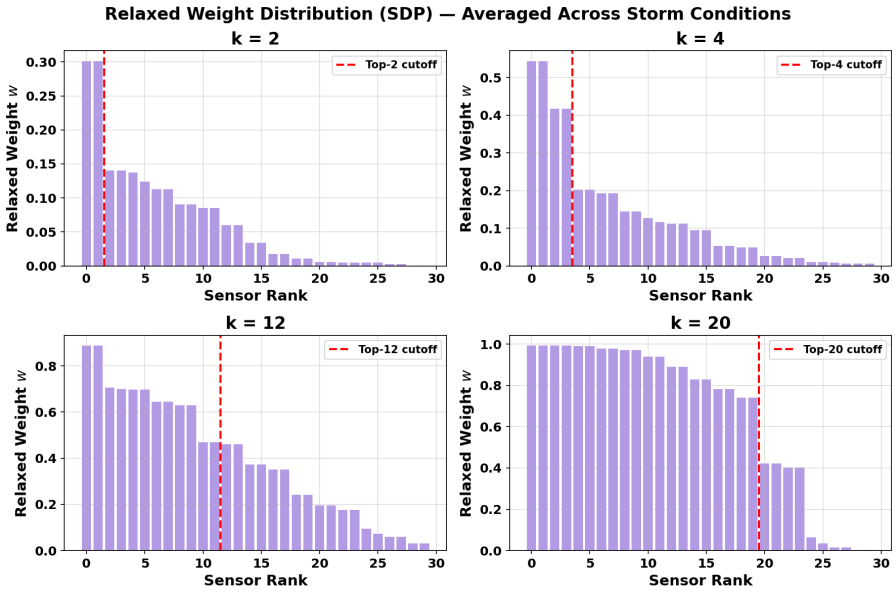

Figure 6. Relaxed monitor-selection weights w_s for several budgets K, sorted in descending order with the top-K cutoff marked and averaged over the 36 storm conditions.

experiment-design formulation provides both a principled ranking of candidate monitor locations and a consistent accuracy advantage over arbitrary placement. That advantage is largest precisely in the small-budget regime—adding a handful of monitors to a region—which is the situation utilities most often face.

6. Conclusion and Future Work

We posed the placement of transformer-neutral GIC monitors as an A-optimal experiment-design problem over a WLS geoelectric-field estimator, treating the estimator's measurement-selection weights as decision variables and minimizing the trace of the estimation-error covariance. Relaxing the binary selection to continuous weights yields a semidefinite program that is solved in polynomial time and rounded to a fixed monitor budget by ranking, and the formulation accommodates already-installed monitors by fixing their weights to one. On the Illini 42-bus synthetic GMD case, evaluated across 36 synthetic storm conditions, the optimized placement achieved a lower estimation cost than arbitrary placement at every monitor budget, and at the smallest budgets it kept the WLS fit within the chi-squared acceptance band where random placement did not. The relaxed weights additionally rank candidate transformers by how much each improves the estimate, giving operators a fast, interpretable guide to where a limited number of new monitors should go.

The study has clear limitations. It uses a single synthetic test system and models the geoelectric field as one region; the latter is the source of the cost growth observed at larger budgets and bounds how far the present results generalize. Top-K rounding can also split weight across highly correlated candidates and carries no certificate that the recovered set attains the combinatorial optimum, and the reported costs are averages over 20 Monte Carlo noise draws per condition.

These limitations set the agenda for future work. Replacing top-K rounding with an ℓ_1-regularized objective would drive the weights toward $\{0, 1\}$ directly, producing a sparse, installation-ready monitor set without a separate rounding step. Estimating the field over multiple regions—or as a spatially continuous quantity, as a companion manuscript develops—would remove the single-region modeling bias and is expected to change the cost-versus-budget behavior. Further work includes benchmarking the formulation against greedy and sensitivity-ranked placement heuristics, scaling it to larger interconnection-sized systems, jointly placing magnetometers and GIC monitors, and validating the optimized placements against field data.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy's Office of Cybersecurity, Energy Security, and Emergency Response (CESER) through project DE-AC05-76RL01830, through a

sub-award from Pacific Northwest National Lab (PNNL). The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

Aarset, C. (2024). A global optimum-informed greedy algorithm for A-optimal experimental design.

Boteler, D. H., & Pirjola, R. J. (2017). Modeling geomagnetically induced currents. *Space Weather*, *15*(1), 258–276. https://doi.org/10.1002/2016SW001499

Boteler, D. H. (2019). A 21st century view of the March 1989 magnetic storm. *Space Weather*, *17*(10), 1427–1441. https://doi.org/10.1029/2019SW002278

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.

Bui, U., Overbye, T. J., Shetye, K., Zhu, H., & Weber, J. (2013). Geomagnetically induced current sensitivity to assumed substation grounding resistance. *2013 North American Power Symposium (NAPS)*, 1–6. https://doi.org/10.1109/NAPS.2013.6666892

Clilverd, M. A., Rodger, C. J., Mac Manus, D. H., Brundell, J. B., Dalzell, M., Renton, A., Lo, V., Lapthorn, A., Smith, A. W., Malone-Leigh, J., et al. (2025). Geomagnetically induced currents, transformer harmonics, and reactive power impacts of the Gannon storm in May 2024. *Space Weather*, *23*(4), e2024SW004235. https://doi.org/10.1029/2024SW004235

Dehghanian, P., Zhang, A., Fatima, R., Snodgrass, J., Birchfield, A. B., Davis, K. R., & Overbye, T. J. (2024). An integrated assessment of a G3 GMD event on large-scale power grids: From magnetometer data to geomagnetically induced current analysis. *IEEE Transactions on Industry Applications*, *60*(1), 1634–1644. https://doi.org/10.1109/TIA.2023.3328974

Fernberg, P. (2012). *One-dimensional earth resistivity models for selected areas of continental United States and Alaska* (Technical Update No. 1026430). Electric Power Research Institute (EPRI). Palo Alto, CA.

Gou, B. (2008). Generalized integer linear programming formulation for optimal PMU placement. *IEEE Transactions on Power Systems*, *23*(3), 1099–1104. https://doi.org/10.1109/TPWRS.2008.926475

Horton, R., Boteler, D., Overbye, T. J., Pirjola, R., & Dugan, R. C. (2012). A test case for the calculation of geomagnetically induced currents. *IEEE Transactions on Power Delivery*, *27*(4), 2368–2373. https://doi.org/10.1109/TPWRD.2012.2206407

Hutchins, T. R., & Overbye, T. J. (2016). Power system dynamic performance during the late-time (E3) high-altitude electromagnetic pulse. *Proceedings of the 19th Power Systems Computation Conference (PSCC)*.

Joshi, S., & Boyd, S. (2009). Sensor selection via convex optimization. *IEEE Transactions on Signal Processing*, *57*(2), 451–462. https://doi.org/10.1109/TSP.2008.2007095

Juvekar, G. P., Klauber, C., Davis, K. R., Overbye, T. J., & Shetye, K. (2021). A GIC-inclusive state estimator for power system awareness during geomagnetic disturbance events. *IEEE Transactions on Power Systems*, *36*(4), 2966–2974. https://doi.org/10.1109/TPWRS.2020.3043734

Kazerooni, M., Zhu, H., Overbye, T. J., & Wojtczak, D. A. (2017). Transmission system geomagnetically induced current model validation. *IEEE Transactions on Power Systems*, *32*(3), 2183–2192. https://doi.org/10.1109/TPWRS.2016.2598265

Kelbert, A., Bedrosian, P. A., & Murphy, B. S. (2019). The first 3D conductivity model of the contiguous United States: Reflections on geologic structure and application to induction hazards. In *Geomagnetically induced currents from the sun to the power grid* (pp. 127–151, Vol. 244). American Geophysical Union. https://doi.org/10.1002/9781119434412.ch8

Klauber, C., Shetye, K. S., Overbye, T. J., & Davis, K. (2020). A GIC estimator for electric grid monitoring during geomagnetic disturbances. *IEEE Transactions on Power Systems*, *35*(6), 4847–4855. https://doi.org/10.1109/TPWRS.2020.3001595

Kunkolienkar, S., Safdarian, F., Snodgrass, J., Birchfield, A. B., & Overbye, T. J. (2024). A description of the Texas A&M University electric grid test case repository for power system studies. *2024 IEEE Texas Power and Energy Conference (TPEC)*.

Lehtinen, M., & Pirjola, R. (1985). Currents produced in earthed conductor networks by geomagnetically-induced electric fields. *Annales Geophysicae*, *3*(4), 479–484.

LoGiudice, N. C., Snodgrass, J. M., & Overbye, T. J. (2024). Estimating the electric field from geomagnetically induced currents. *2024 IEEE Kansas Power and Energy Conference (KPEC)*, 1–4. https://doi.org/10.1109/KPEC61529.2024.10676153

LoGiudice, N. C., Snodgrass, J. M., & Overbye, T. J. (2026). Estimating a spatially continuous electric field from geomagnetically induced currents [Manuscript under review; preprint, arXiv (placeholder — update with arXiv ID/DOI before submission)].

Molinski, T. S. (2002). Why utilities respect geomagnetically induced currents. *Journal of Atmospheric and Solar-Terrestrial Physics*, *64*(16), 1765–1778. https://doi.org/10.1016/S1364-6826(02)00126-8

Monticelli, A. (1999). *State estimation in electric power systems: A generalized approach*. Springer.

Nagata, T., Nonomura, T., Nakai, K., Yamada, K., Saito, Y., & Ono, S. (2021). Data-driven sparse sensor placement based on A-optimal design of experiment with ADMM. *IEEE Sensors Journal*, *21*(13), 15098–15109. https://doi.org/10.1109/JSEN.2021.3073978

North American Electric Reliability Corporation (NERC). (2010). High-impact, low-frequency event risk to the North American bulk power system.

North American Electric Reliability Corporation (NERC). (2016). Benchmark geomagnetic disturbance event description.

North American Electric Reliability Corporation (NERC). (2023). TPL-007-4: Transmission system planned performance for geomagnetic disturbance events.

Overbye, T. J., Hutchins, T. R., Shetye, K., Weber, J., & Dahman, S. (2012). Integration of geomagnetic disturbance modeling into the power flow: A methodology for large-scale system studies. *2012 North American Power Symposium (NAPS)*, 1–7. https://doi.org/10.1109/NAPS.2012.6336365

Pirjola, R. (2002). Review on the calculation of surface electric and magnetic fields and geomagnetically induced currents in ground-based technological systems. *Surveys in Geophysics*, *23*(1), 71–90. https://doi.org/10.1023/A:1014816009303

Pukelsheim, F. (2006). *Optimal design of experiments*. Society for Industrial and Applied Mathematics (SIAM).

Pulkkinen, A., Lindahl, S., Viljanen, A., & Pirjola, R. (2005). Geomagnetic storm of 29–31 October 2003: Geomagnetically induced currents and their relation to problems in the Swedish high-voltage power transmission system. *Space Weather*, *3*(8), S08C03. https://doi.org/10.1029/2004SW000123

Saleh, S. A., Zundel, E. W., Young-Morris, G., Jewett, D., Brown, S., Hill, E. F. S., & Meng, J. (2023). Investigating harmonic distortion in power transformers due to geomagnetically induced current flows. *IEEE Transactions on Industry Applications*, *59*(5), 5432–5442. https://doi.org/10.1109/TIA.2023.3285938

Shetye, K. S., Kumar, R. R., Klauber, C., Mao, Z., Overbye, T. J., Gannon, J., & Henderson, M. (2021). Development and electric grid applications of a magnetometer network. *IEEE Open Access Journal of Power and Energy*, *8*, 77–84. https://doi.org/10.1109/OAJPE.2021.3056593

Winter, L. M., Gannon, J., Pernak, R., Huston, S., Quinn, R., Pope, E., Ruffenach, A., Bernardara, P., & Crocker, N. (2017). Spectral scaling technique to determine extreme Carrington-level geomagnetically induced currents effects. *Space Weather*, *15*(5), 713–725. https://doi.org/10.1002/2016SW001586

Xu, B., & Abur, A. (2004). Observability analysis and measurement placement for systems with PMUs. *IEEE PES Power Systems Conference and Exposition*, 943–946. https://doi.org/10.1109/PSCE.2004.1397683

Yu, J., Zavala, V. M., & Anitescu, M. (2018). A scalable design of experiments framework for optimal sensor placement. *Journal of Process Control*, *67*, 44–55. https://doi.org/10.1016/j.jprocont.2017.03.011

Zhu, H., & Overbye, T. J. (2015). Blocking device placement for mitigating the effects of geomagnetically induced currents. *IEEE Transactions on Power Systems*, *30*(4), 2081–2089. https://doi.org/10.1109/TPWRS.2014.2349733