Load Curve Data Cleansing and Imputation Via Sparsity and Low Rank

Gonzalo Mateos, Member, IEEE, and Georgios B. Giannakis, Fellow, IEEE

Abstract—The smart grid vision is to build an intelligent power network with an unprecedented level of situational awareness and controllability over its services and infrastructure. This paper advocates statistical inference methods to robustify power monitoring tasks against the outlier effects owing to faulty readings and malicious attacks, as well as against missing data due to privacy concerns and communication errors. In this context, a novel load cleansing and imputation scheme is developed leveraging the low intrinsic-dimensionality of spatiotemporal load profiles and the sparse nature of "bad data." A robust estimator based on principal components pursuit (PCP) is adopted, which effects a twofold sparsity-promoting regularization through an ℓ_1 -norm of the outliers, and the nuclear norm of the nominal load profiles. Upon recasting the non-separable nuclear norm into a form amenable to decentralized optimization, a distributed (D-) PCP algorithm is developed to carry out the imputation and cleansing tasks using networked devices comprising the so-termed advanced metering infrastructure. If D-PCP converges and a qualification inequality is satisfied, the novel distributed estimator provably attains the performance of its centralized PCP counterpart, which has access to all networkwide data. Computer simulations and tests with real load curve data corroborate the convergence and effectiveness of the novel D-PCP algorithm.

Index Terms—Advanced metering infrastructure, distributed algorithms, load curve cleansing and imputation, principal components pursuit, smart grid.

I. INTRODUCTION

T HE U.S. power grid has been recognized as the most important engineering achievement of the 20th century [27], yet it presently faces major challenges related to efficiency, reliability, security, environmental impact, sustainability, and market diversity issues [26]. The crystallizing vision of the smart grid (SG) aspires to build a cyber-physical network that can address these challenges by capitalizing on state-of-the-art information technologies in sensing, control, communication, optimization, and machine learning. Significant effort and investment are being committed to architect the necessary infrastructure by installing advanced metering systems, and establishing data communication networks throughout the grid.

The authors are with the Department of ECE and the Digital Technology Center, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: mate0058@umn.edu; georgios@umn.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSG.2013.2259853

Accordingly, algorithms that optimally exploit the pervasive sensing and control capabilities of the envisioned advanced metering infrastructure (AMI) are needed to make the necessary breakthroughs in the key problems in power grid monitoring and energy management. This is no easy endeavor though, in view of the challenges posed by increasingly distributed network operations under strict reliability requirements, also facing malicious cyber-attacks.

Statistical inference techniques are expected to play an increasingly instrumental role in power system monitoring [10], not only to meet the anticipated "big data" deluge as the installed base of phasor measurement units (PMUs) reaches out throughout the grid, but also to robustify the monitoring tasks against the "outlier" effects owing to faulty readings, malicious attacks, and communication errors, as well as against missing data due to privacy concerns and technical anomalies [8]. In this context, a *load cleansing and imputation* scheme is developed in this paper, building on recent advances in sparsity-cognizant information processing [12], low-rank matrix completion [9], and large-scale distributed optimization [4].

Load curve data refers to the electric energy consumption periodically recorded by meters at points of interest across the power grid, e.g., end-user premises, buses, and substations. Accurate load profiles are critical assets aiding operational decisions in the envisioned SG system [7], and are essential for load forecasting [25]. However, in the process of acquiring and transmitting such massive volumes of information for centralized processing, data are oftentimes corrupted or lost altogether. In a smart monitoring context for instance, incomplete load profiles emerge due to three reasons: r1) PMU-instrumented buses are few; r2) SCADA data become available at a considerably smaller time scale than PMU data; and r3) regional operators are not willing to share all their variables [24]. Moreover, a major requirement for grid monitoring is robustness to outliers, i.e., data not adhering to nominal models [1], [23]. Sources of so-termed "bad data" include meter failures, as well as strikes, unscheduled generator shutdowns, and extreme weather conditions [7], [11]. Inconsistent data can also be due to malicious (cyber-) attacks that induce abrupt load changes, or counterfeit meter readings [18].

In light of the aforementioned observations, the first contribution of this paper is on modeling spatiotemporal load profiles, accounting for the structure present to effectively impute missing data and devise robust load curve estimators stemming from convex optimization criteria (Section II). Existing approaches to load curve cleansing have relied on separate processing per time series [7], [11], [22], and have not capitalized on spatial correlations to improve performance. The aim

Manuscript received September 21, 2012; revised February 20, 2013; accepted April 21, 2013. Date of publication September 13, 2013; date of current version November 25, 2013. This work was supported by the Inst. of Renewable Energy and the Environment (IREE) under Grant RL-0010-13, University of Minnesota. Part of this paper was presented at the 3rd International Conference on Smart Grid Communications, Tainan, Taiwan, November 5–8, 2012. Paper no. TSG-00615-2012.

is for minimal-rank cleansed load data, while also exploiting outlier sparsity across buses and time. An estimator tailored to these specifications is principal components pursuit (PCP) [5], [6], [29], which is outlined in Section III. PCP minimizes a tradeoff between the least-squares (LS) model fitting error and a twofold sparsity-promoting regularization, implemented through an ℓ_1 -norm of the outliers and the nuclear norm of the nominal load profiles. While PCP has been widely adopted in computer vision [5], for voice separation in music [13], and unveiling network anomalies [20], its benefits to power systems engineering and monitoring remains so far largely unexplored. The second contribution pertains to developing a distributed (D-) PCP algorithm, to carry out the imputation and cleansing task using a network of interconnected devices as part of the AMI (Section IV). This is possible by leveraging a general algorithmic framework for sparsity-regularized rank minimization put forth in [20]. Upon recasting the (non-separable) nuclear norm present in the PCP cost into a form amenable to decentralized optimization (Section IV-A), the D-PCP iterations are obtained in Section IV-C via the multi-block alternating-directions method of multipliers (ADMM) solver [3], [4]. In a nutshell, per iteration each smart meter exchanges simple messages with its (directly connected) neighbors in the network, and then solves its own optimization problem to refine its current estimate of the cleansed load profile. In the context of power systems, the ADMM has been recently adopted to carry out dynamic network energy management [16], and distributed robust state estimation [14].

Computer simulations corroborate the convergence and optimality of the novel D-PCP algorithm, and demonstrate its effectiveness in cleansing and imputing real load curve data (Section V). Concluding remarks and directions for future research are outlined in Section VI, while a few algorithmic details are deferred to the Appendix.

II. MODELING AND PROBLEM STATEMENT

This section introduces the model for (possibly) incomplete and grossly corrupted load curve measurements, acquired by geographically-distributed metering devices monitoring the power grid. The communication network model needed to account for exchanges of information among smart meters is described as well. Lastly, the task of load curve cleansing and imputation is formally stated.

A. Spatiotemporal Load Curve Data Model

Let the $N \times 1$ vector $\mathbf{y}(t) := [y_{1,t}, \ldots, y_{N,t}]'$ (' stands for transposition) collect the spatial load profiles measured by smart meters monitoring N network nodes (buses, residential premises), at a given discrete-time instant $t \in [1, T]$. Consider the $N \times T$ matrix of observations $\mathbf{Y} := [\mathbf{y}(1), \ldots, \mathbf{y}(T)]$. The *n*-th row $(\mathbf{y}_n)'$ of \mathbf{Y} is the time series of energy consumption (load curve) measurements at node *n*, while the *t*-th column $\mathbf{y}(t)$ of \mathbf{Y} represents a snapshot of the networkwide loads taken at time *t*. To model missing data, consider the set $\Omega \subseteq \{1, \ldots, N\} \times \{1, \ldots, T\}$ of index pairs (n, t) defining a sampling of the entries of \mathbf{Y} . Introducing the matrix sampling operator $\mathcal{P}_{\Omega}(\cdot)$, which sets the entries of its matrix argument not indexed by Ω to zero and leaves the rest unchanged, the (possibly) incomplete spatiotemporal load curve data in the presence of outliers can be modeled as

$$\mathcal{P}_{\Omega}(\mathbf{Y}) = \mathcal{P}_{\Omega}(\mathbf{X} + \mathbf{O} + \mathbf{E})$$
(1)

where **X**, **O**, and **E** denote the nominal load profiles, the outliers, and small measurement errors, respectively. For nominal observations $y_{n,t} = x_{n,t} + e_{n,t}$, one has $o_{n,t} = 0$.

Remark 1 (Model Under-Determinacy): The model is inherently under-determined, since even for the (most favorable) case of full data, i.e., $\Omega \equiv \{1, ..., N\} \times \{1, ..., T\}$, there are twice as many unknowns in **X** and **O** as there is data in **Y**. Estimating **X** and **O** becomes even more challenging when data are missing, since the number of unknowns remains the same, but the amount of data is reduced.

In any case, estimation of $\{\mathbf{X}, \mathbf{O}\}\$ from $\mathcal{P}_{\Omega}(\mathbf{Y})$ is an illposed problem unless one introduces extra structural assumptions on the model components to reduce the effective degrees of freedom. To this end, two cardinal properties of X and O will prove instrumental. First, common temporal patterns among the energy consumption of a few broad classes of loads (e.g., industrial, residential, seasonal) in addition to their (almost) periodic behaviors render most rows and columns of X linearly dependent, and thus X typically has low-rank. Second, outliers (or attacks) only occur sporadically in time and affect only a few buses, yielding a sparse matrix O. Smoothness of the nominal load curves is related to the low-rank property of X, which was adopted in [7] to motivate a smoothing splines-based algorithm for cleansing. However, while in [7] smoothness was enforced per load profile time-series, i.e., per row $(\mathbf{x}_n)'$ of **X**; the low-rank property of X also captures the spatial dependencies introduced by the network. Approaches capitalizing on outlierand "bad data-" sparsity can be found in, e.g., [15], [14] and [22].

B. Communication Network Model

Suppose that on top of the energy measurement functionality, the N networked smart meters are capable of performing simple local computations, as well as exchanging messages among directly connected neighbors. Single-hop communication models are appealing due to their simplicity, since one does not have to incur the routing overhead. The AMI network is naturally abstracted to an undirected graph $G(\mathcal{N}, \mathcal{L})$, where the vertex set $\mathcal{N} := \{1, \dots, N\}$ corresponds to the network nodes, and the edges (links) in \mathcal{L} represent pairs of nodes that are connected via a physical communication channel. Node $n \in \mathcal{N}$ communicates with its single-hop neighboring peers in \mathcal{J}_n , and the size of the neighborhood will be henceforth denoted by $|\mathcal{J}_n|$. The graph G is assumed connected, i.e., there exists a (possibly multihop) path that joins any pair of nodes in the network. This requirement ensures that the network is devoid of multiple isolated (connected) components, and allows for the data collected by, e.g., smart meter n, namely the n-th row $(\mathbf{y}_n)'$ of Y, to eventually reach every other node in the network. This way, even when only local interactions are allowed, the flow of information can percolate the network.

The importance of the network model will become apparent in Section IV.

C. Load Curve Cleansing and Imputation

The load curve cleansing and imputation problem studied here entails identification and removal of outliers (or "bad data"), in addition to completion of the missing entries from the nominal load matrix, and denoising of the observed ones. To some extent, it is a joint estimation-interpolation (prediction)-detection problem. With reference to (1), given generally incomplete, noisy and outlier-contaminated spatiotemporal load data $\mathcal{P}_{\Omega}(\mathbf{Y})$, the cleansing and imputation tasks amount to estimating the nominal load profiles **X** and the outliers **O**, by leveraging the low-rank property of **X** and the sparsity in **O**. Collaboration between metering devices (collecting networkwide data) is considered here, rather than local processing of load curves per bus.

Note that load cleansing and imputation are different from *load forecasting* [25], which amounts to predicting future load demand based on historical data of energy consumption and the weather conditions. Actually, cleansing and imputation are critical preprocessing tasks utilized to enhance the quality of load data, that would be subsequently used for load forecasting and optimum power flow [2].

III. PRINICIPAL COMPONENTS PURSUIT

An estimator matching nicely the specifications of the load curve cleansing and imputation problem stated in Section II-C, is the so-termed (stable) principal components pursuit (PCP) [5], [6], [29], that will be outlined here for completeness. PCP seeks estimates $\{\hat{\mathbf{X}}, \hat{\mathbf{O}}\}$ as the minimizers of

(P1)
$$\min_{\{\mathbf{X},\mathbf{O}\}} \frac{1}{2} \|\mathcal{P}_{\Omega}(\mathbf{Y} - \mathbf{X} - \mathbf{O})\|_{F}^{2} + \lambda_{*} \|\mathbf{X}\|_{*} + \lambda_{1} \|\mathbf{O}\|_{1}$$

where the ℓ_1 -norm $\|\mathbf{O}\|_1 := \sum_{n,t} |o_{n,t}|$ and the nuclear norm $\|\mathbf{X}\|_* := \sum_i \sigma_i(\mathbf{X}) \ (\sigma_i(\mathbf{X}) \text{ denotes the } i\text{-th singular value of}$ **X**) are utilized to promote sparsity in the number of outliers (nonzero entries) in O, and the low rank of X, respectively. The nuclear and ℓ_1 -norms are the closest convex surrogates to the rank and cardinality functions, which albeit the most natural criteria they are in general NP-hard to optimize. The tuning parameters $\lambda_1, \lambda_* \geq 0$ control the tradeoff between fitting error, rank, and sparsity level of the solution. When an estimate $\hat{\sigma}_{v}^{2}$ of the observation noise variance is available, guidelines for selecting λ_* and λ_1 have been proposed in [29]. The location of nonzero entries in Ô reveals "bad data" across both buses and time, while their amplitudes quantify the magnitude of the anomalous event. Thus, identification of outliers is critical not only to cleanse load profiles and enhance the performance of load forecasting algorithms, but also to, e.g., identify faulty meters, and out of the ordinary load consumptions that may deserve special treatment. Clearly, it does not make sense to flag outliers in data that has not been observed, namely for $(n, t) \notin \Omega$. In those cases (P1) yields $\hat{o}_{n,t} = 0$ since both the Frobenious and ℓ_1 -norms are separable across the entries of their matrix arguments.

Being convex (P1) is computationally appealing, and it has been shown to attain good performance in theory and practice. For instance, in the absence of noise and when there is no missing data, identifiability and exact recovery conditions were reported in [5] and [6]. Even when data are missing, it is possible to recover the low-rank component under some technical assumptions [5]. Theoretical performance guarantees in the presence of noise are also available [29].

Remark 2 (Online Processing for Dynamic Grid Monitoring): Problem (P1) offers a batch estimator for cleansing and imputation of load curve data, which are collected over a horizon of T time instants and only processed a posteriori. In dynamic environments, real-time adaptive algorithms are desirable since they are capable of efficiently processing meter measurements and identifying outliers "on the fly." Online processing enables tracking of "bad data" in nonstationary environments, typically arising due to, e.g., network topology changes and missing data. An online variant of (P1) was puth forth in [21], along with an efficient real-time algorithm with performance guarantees.

Regarding batch algorithms, a PCP solver based on the accelerated proximal gradient method was put forth in [17], while the ADMM was employed in [28]. Implementing these centralized algorithms presumes that networked metering devices continuously communicate their local load measurements $y_{n,t}$ to a central monitoring and data analytics station, which uses their aggregation in $\mathcal{P}_{\Omega}(\mathbf{Y})$ to reject outliers and impute missing data. While for the most part this is the prevailing operational paradigm nowadays, there are limitations associated with this architecture. For instance, collecting all this information centrally may lead to excessive overhead in the communication network, especially when the rate of data acquisition is high at the meters. Moreover, minimizing (or avoiding altogether) the exchanges of raw measurements may be desirable for privacy and cyber-security reasons, as well as to reduce unavoidable communication errors that translate to missing data. Performing the optimization in a centralized fashion raises robustness concerns as well, since the central data analytics station represents an isolated point of failure. These reasons motivate devising fully-distributed iterative algorithms for PCP, embedding the load cleansing and imputation functionality to the AMI. This is the subject of the next section.

IV. DISTRIBUTED CLEANSING AND IMPUTATION

A distributed (D-)PCP algorithm to solve (P1) using a network of smart meters (modeled as in Section II-B) should be understood as an iterative method, whereby each node carries out simple local (optimization) tasks per iteration k = 1, 2, ...,and exchanges messages only with its directly connected neighbors. The ultimate goal is for each node to form local estimates $\mathbf{x}_n[k]$ and $\mathbf{o}_n[k]$ that coincide with the *n*-th rows of $\hat{\mathbf{X}}$ and $\hat{\mathbf{O}}$ as $k \to \infty$, where { $\hat{\mathbf{X}}, \hat{\mathbf{O}}$ } is the solution of (P1) obtained when all data $\mathcal{P}_{\Omega}(\mathbf{Y})$ are centrally available. Attaining the centralized performance with distributed data is impossible if the network is disconnected.

To facilitate reducing the computational complexity and memory storage requirements of the D-PCP algorithm sought, it is henceforth assumed that an upper bound rank($\hat{\mathbf{X}}$) $\leq \rho$ is a priori available [recall $\hat{\mathbf{X}}$ is the estimated low-rank cleansed load profile obtained via (P1)]. As argued next, the smaller the value of ρ , the more efficient the algorithm becomes. Small values of ρ are well motivated due to the low intrinsic dimensionality of the spatiotemporal load profiles (cf. Section II-A). Because rank($\hat{\mathbf{X}}$) $\leq \rho$, (P1)'s search space is effectively reduced and one can factorize the decision variable as $\mathbf{X} = \mathbf{PQ'}$, where **P** and **Q** are $N \times \rho$ and $T \times \rho$ matrices, respectively. Adopting this reparametrization of **X** in (P1) and making explicit the distributed nature of the data (cf. Section II-B), one arrives at an equivalent optimization problem

(P2)
$$\min_{\{\mathbf{P},\mathbf{Q},\mathbf{O}\}} \sum_{n=1}^{N} \left[\frac{1}{2} \| \mathcal{P}_{\Omega_n} (\mathbf{y}_n - \mathbf{Q}\mathbf{p}_n - \mathbf{o}_n) \|_2^2 + \frac{\lambda_*}{N} \| \mathbf{P}\mathbf{Q}' \|_* + \lambda_1 \| \mathbf{o}_n \|_1 \right]$$

which is non-convex due to the bilinear term $\mathbf{PQ'}$, and where $\mathbf{P} := [\mathbf{p}_1, \dots, \mathbf{p}_N]'$. The number of variables is reduced from 2NT in (P1), to $\rho(N + T) + NT$ in (P2). The savings can be significant when ρ is small, and both N and T are large. Note that the dominant NT-term in the variable count of (P2) is due to \mathbf{O} , which is sparse and can be efficiently handled even when both N and T are large.

Remark 3 (Challenges Facing Distributed Implementation): Problem (P2) is still not amenable for distributed implementation due to: c1) the non-separable nuclear norm present in the cost function; and c2) the global variable \mathbf{Q} coupling the per-node summands.

Challenges c1)–c2) are dealt with in the ensuing sections.

A. A Separable Low-Rank Regularization

To address c1), consider the following alternative characterization of the nuclear norm (see, e.g., [20])

$$\|\mathbf{X}\|_{*} := \min_{\{\mathbf{P}, \mathbf{Q}\}} \frac{1}{2} \left(\|\mathbf{P}\|_{F}^{2} + \|\mathbf{Q}\|_{F}^{2} \right), \quad \text{s. to } \mathbf{X} = \mathbf{P}\mathbf{Q}'.$$
(2)

The optimization (2) is over all possible bilinear factorizations of \mathbf{X} , so that the number of columns ρ of \mathbf{P} and \mathbf{Q} is also a variable. Leveraging (2), the following reformulation of (P2) provides an important first step towards obtaining the D-PCP algorithm:

(P3)
$$\min_{\{\mathbf{P},\mathbf{Q},\mathbf{O}\}} \sum_{n=1}^{N} \left[\frac{1}{2} \| \mathcal{P}_{\Omega_n} (\mathbf{y}_n - \mathbf{Q}\mathbf{p}_n - \mathbf{o}_n) \|_2^2 + \lambda_1 \|\mathbf{o}_n\|_1 + \frac{\lambda_*}{2N} \left(N \|\mathbf{p}_n\|_2^2 + \|\mathbf{Q}\|_F^2 \right) \right].$$

As asserted in [20], adopting the separable Frobenius-norm regularization in (P3) comes with no loss of optimality relative to (P1), provided rank($\hat{\mathbf{X}}$) $\leq \rho$. By finding the global minimum of (P3) [which could have considerably less variables than (P1)], one can recover the optimal solution of (P1). This could be challenging however, since (P3) is non-convex and it may have stationary points which need not be globally optimum.

Interestingly, it is possible to certify global optimality of a stationary point { $\bar{\mathbf{P}}$, $\bar{\mathbf{Q}}$, $\bar{\mathbf{O}}$ } of (P3). Specifically, one can establish that if $\|\mathcal{P}_{\Omega}(\mathbf{Y} - \bar{\mathbf{P}}\bar{\mathbf{Q}}' - \bar{\mathbf{O}})\| < \lambda_*$, then { $\hat{\mathbf{X}} := \bar{\mathbf{P}}\bar{\mathbf{Q}}'$, $\hat{\mathbf{O}} := \bar{\mathbf{O}}$ } is the globally optimal solution of (P1) ([20], Prop. 1). The qualification condition $\|\mathcal{P}_{\Omega}(\mathbf{Y} - \bar{\mathbf{P}}\bar{\mathbf{Q}}' - \bar{\mathbf{O}})\| < \lambda_*$ captures tacitly the role of ρ . In particular, for sufficiently small ρ the residual $\|\mathcal{P}_{\Omega}(\mathbf{Y} - \bar{\mathbf{P}}\bar{\mathbf{Q}}' - \bar{\mathbf{O}})\|$ becomes large and consequently the condition is violated [unless λ_* is large enough, in which case a sufficiently low-rank solution to (P1) is expected]. The condition on the residual also implicitly enforces $\operatorname{rank}(\hat{\mathbf{X}}) \leq \rho$, which is necessary for the equivalence between (P1) and (P3).

B. Local Variables and Consensus Constraints

To decompose the cost in (P3), in which summands inside the square brackets are coupled through the global variable \mathbf{Q} [cf. c2) under Remark 3], introduce auxiliary variables $\{\mathbf{Q}_n\}_{n=1}^N$ representing local estimates of \mathbf{Q} per smart meter n. To obtain a separable PCP formulation, use these estimates along with *consensus* constraints

(P4)
$$\min_{\{\mathbf{P}_{n},\mathbf{Q}_{n},\mathbf{O}\}} \sum_{n=1}^{N} \left[\frac{1}{2} \| \mathcal{P}_{\Omega_{n}}(\mathbf{y}_{n} - \mathbf{Q}_{n}\mathbf{p}_{n} - \mathbf{o}_{n}) \|_{2}^{2} + \lambda_{1} \|\mathbf{o}_{n}\|_{1} + \frac{\lambda_{*}}{2N} \left(N \|\mathbf{p}_{n}\|_{2}^{2} + \|\mathbf{Q}_{n}\|_{F}^{2} \right) \right]$$

s. to $\mathbf{Q}_{n} = \mathbf{Q}_{m}, \quad m \in \mathcal{J}_{n}, n \in \mathcal{N}.$

Notice that (P3) and (P4) are equivalent optimization problems, since the network graph $G(\mathcal{N}, \mathcal{L})$ is connected by assumption. Even though consensus is a fortiori imposed only within neighborhoods, it extends to the whole (connected) network and local estimates agree on the global solution of (P3). To arrive at the desired D-PCP algorithm, it is convenient to reparametrize the consensus constraints in (P4) as

$$\mathbf{Q}_{n} = \bar{\mathbf{F}}_{n}^{m}, \ \mathbf{Q}_{m} = \tilde{\mathbf{F}}_{n}^{m}, \ \text{and} \ \bar{\mathbf{F}}_{n}^{m} = \tilde{\mathbf{F}}_{n}^{m}, \quad m \in \mathcal{J}_{n}, \ n \in \mathcal{N}$$
(3)

where $\{\bar{\mathbf{F}}_{n}^{m}, \tilde{\mathbf{F}}_{n}^{m}\}_{n \in \mathcal{N}}$, are auxiliary optimization variables that will be eventually eliminated (cf. Remark 4).

C. The D-PCP Algorithm

To tackle (P4), associate Lagrange multipliers $\bar{\mathbf{M}}_n^m$ and $\bar{\mathbf{M}}_n^m$ with the first pair of consensus constraints in (3). Introduce the quadratically *augmented* Lagrangian function [3]

$$\mathcal{L}_{c}(\mathcal{V}_{1}, \mathcal{V}_{2}, \mathcal{V}_{3}, \mathcal{M}) = \sum_{n=1}^{N} \left[\frac{1}{2} \| \mathcal{P}_{\Omega_{n}}(\mathbf{y}_{n} - \mathbf{Q}_{n}\mathbf{p}_{n} - \mathbf{o}_{n}) \|_{2}^{2} + \lambda_{1} \|\mathbf{o}_{n}\|_{1} + \frac{\lambda_{*}}{2N} \left(N \|\mathbf{p}_{n}\|_{2}^{2} + \|\mathbf{Q}_{n}\|_{F}^{2} \right) \right] + \sum_{n=1}^{N} \sum_{m \in \mathcal{J}_{n}} \left(\left\langle \bar{\mathbf{M}}_{n}^{m}, \mathbf{Q}_{n} - \bar{\mathbf{F}}_{n}^{m} \right\rangle + \left\langle \tilde{\mathbf{M}}_{n}^{m}, \mathbf{Q}_{m} - \tilde{\mathbf{F}}_{n}^{m} \right\rangle \right) + \frac{c}{2} \sum_{n=1}^{N} \sum_{m \in \mathcal{J}_{n}} \left(\left\| \mathbf{Q}_{n} - \bar{\mathbf{F}}_{n}^{m} \right\|_{F}^{2} + \left\| \mathbf{Q}_{m} - \tilde{\mathbf{F}}_{n}^{m} \right\|_{F}^{2} \right)$$

$$(4)$$

where c > 0 is a penalty parameter, and the primal variables are split into three groups $\mathcal{V}_1 := \{\mathbf{Q}_n\}_{n=1}^N, \mathcal{V}_2 := \{\mathbf{p}_n\}_{n=1}^N$ and $\mathcal{V}_3 := \{\mathbf{o}_n, \bar{\mathbf{F}}_n^m, \tilde{\mathbf{F}}_n^m\}_{n \in \mathcal{N}}^{m \in \mathcal{J}_n}$. For notational brevity, collect all Lagrange multipliers in $\mathcal{M} := \{\bar{\mathbf{M}}_n^m, \tilde{\mathbf{M}}_n^m\}_{n \in \mathcal{N}}^{m \in \mathcal{J}_n}$. Note that the remaining constraints in (3), namely $\mathcal{C}_F := \{\bar{\mathbf{F}}_n^m = \tilde{\mathbf{F}}_n^m, m \in \mathcal{J}_n, n \in \mathcal{N}\}$, have not been dualized.

To minimize (P4) in a distributed fashion, (a multi-block variant of) the ADMM will be adopted here. The ADMM is an

iterative augmented Lagrangian method especially well suited for parallel processing [3], [4], which has been proven successful to tackle the optimization tasks stemming from general distributed estimators of deterministic and (non-)stationary random signals; see, e.g., [14], [20] and references therein. The proposed solver entails an iterative procedure comprising four steps per iteration $k = 1, 2, ..., [n \in \mathcal{N}, m \in \mathcal{J}_n \text{ in (5)-(6)}]$

[S1] Update dual variables:

$$\bar{\mathbf{M}}_{n}^{m}[k] = \bar{\mathbf{M}}_{n}^{m}[k-1] + c\left(\mathbf{Q}_{n}[k] - \bar{\mathbf{F}}_{n}^{m}[k]\right)$$
(5)

$$\tilde{\mathbf{M}}_{n}^{m}[k] = \tilde{\mathbf{M}}_{n}^{m}[k-1] + c\left(\mathbf{Q}_{m}[k] - \tilde{\mathbf{F}}_{n}^{m}[k]\right).$$
(6)

[S2] Update first group of primal variables:

$$\mathcal{V}_1[k+1] = \arg \min_{\mathcal{V}_1} \mathcal{L}_c(\mathcal{V}_1, \mathcal{V}_2[k], \mathcal{V}_3[k], \mathcal{M}[k]).$$
(7)

[S3] Update second group of primal variables:

$$\mathcal{V}_2[k+1] = \arg \min_{\mathcal{V}_2} \mathcal{L}_c(\mathcal{V}_1[k+1], \mathcal{V}_2, \mathcal{V}_3[k], \mathcal{M}[k]).$$
(8)

[S4] Update third group of primal variables:

$$\mathcal{V}_3[k+1] = \arg \min_{\mathcal{V}_3 \in \mathcal{C}_F} \mathcal{L}_c(\mathcal{V}_1[k+1], \mathcal{V}_2[k+1], \mathcal{V}_3, \mathcal{M}[k])$$
(9)

which amount to a block-coordinate descent method cycling over $\mathcal{V}_1 \rightarrow \mathcal{V}_2 \rightarrow \mathcal{V}_3$ to minimize \mathcal{L}_c , and dual variable updates [3]. At each step while minimizing the augmented Lagrangian, the variable groups not being updated are treated as fixed, and are substituted with their most up to date values. Different from the standard two-block ADMM [3], [4], the multi-block variant here cycles over three groups of primal variables [19].

Reformulating the estimator (P1) to its equivalent form (P4) renders the augmented Lagrangian in (4) highly decomposable. The separability comes in two flavors, both with respect to the variable groups $\mathcal{V}_1 - \mathcal{V}_3$, as well as across the network nodes $n \in \mathcal{N}$. This leads to highly parallelized, simplified recursions to be run by the networked smart meters. Specifically, it is shown in the Appendix that the aforementioned ADMM steps [S1]–[S4] give rise to the D-PCP iterations tabulated under Algorithm 1. Per iteration, each device updates: [S1] a local matrix of dual prices $\mathbf{S}_n[k]$; [S2]–[S3] local cleansed load estimates $\mathbf{Q}_n[k+1]$ and $\mathbf{p}_n[k+1]$ obtained as solutions to respective unconstrained quadratic problems (QPs); and [S4] its local outlier vector, through a sparsity-promoting soft-thresholding operation. The (k + 1)-st iteration is concluded after smart meter n transmits $\mathbf{Q}_n[k+1]$ to its single-hop neighbors in \mathcal{J}_n .

By inspection of Algorithm 1, one realizes that the required memory space per smart meter n is only $T + T\rho + T\rho + \rho$ scalars, to store \mathbf{y}_n , $\mathbf{S}_n[k]$, $\mathbf{Q}_n[k]$, and $\mathbf{p}_n[k]$, respectively. This number of scalars can be upper bounded by $3T\rho$, which will be typically smaller that NT required by a centralized solver to store the data \mathbf{Y} . (Note that ρ is typically small due to the low rank assumption, while N will be large in pragmatic networks.) Regarding communication cost, $\mathbf{Q}_n[k+1]$ is a $T \times \rho$ matrix and its transmission does not incur significant overhead for small ρ . Observe also that $\mathcal{P}_{\Omega}(\mathbf{y}_n)$ need not be exchanged which is desirable to preserve data secrecy, and the communication cost is independent of N.

Algorithm 1: D-PCP at smart meter $n \in \mathcal{N}$

input
$$\mathbf{y}_n, \mathbf{\Omega}_n, \lambda_*, \lambda_1$$
, and c .

initialize $\mathbf{S}_n[0] = \mathbf{0}_{T \times \rho}$, and $\mathbf{Q}_n[1]$, $\mathbf{p}_n[1]$ at random.

for
$$k = 1, 2, ...$$
 do

Receive $\{\mathbf{Q}_m[k]\}\$ from neighbors $m \in \mathcal{J}_n$.

[S1] Update local dual variables:

$$\mathbf{S}_n[k] = \mathbf{S}_n[k-1] + c \sum_{m \in \mathcal{J}_n} (\mathbf{Q}_n[k] - \mathbf{Q}_m[k]).$$

[S2] Update first group of local primal variables:

$$\mathbf{A}_{n}[k+1] := \{ (\mathbf{p}_{n}[k]\mathbf{p}_{n}'[k]) \otimes \mathbf{\Omega}_{n} + (\lambda_{*}/N + 2c|\mathcal{J}_{n}|)\mathbf{I}_{\rho T} \}^{-1} \\ \mathbf{Q}_{n}[k+1] = \mathrm{unvec}(\mathbf{A}_{n}[k+1]\{(\mathbf{p}_{n}[k] \otimes \mathbf{\Omega}_{n})(\mathbf{y}_{n} - \mathbf{o}_{n}[k]) \\ - \mathrm{vec}(\mathbf{S}_{n}[k]) + \mathrm{vec}(c\sum_{m \in \mathcal{J}_{n}} (\mathbf{Q}_{n}[k] + \mathbf{Q}_{m}[k])) \}).$$

[S3] Update second group of local primal variables:

$$\mathbf{p}_n[k+1] = \{\mathbf{Q}'_n[k+1]\mathbf{\Omega}_n\mathbf{Q}_n[k+1] + \mathbf{I}_\rho\}^{-1} \\ \times \mathbf{Q}'_n[k+1]\mathbf{\Omega}_n(\mathbf{y}_n - \mathbf{o}_n[k])$$

[S4] Update third group of local primal variables:

 $\mathbf{o}_n[k+1] = \mathcal{S}_{\lambda_1}(\mathbf{\Omega}_n(\mathbf{y}_n - \mathbf{Q}_n[k+1]\mathbf{p}_n[k+1])).$

Transmit $\mathbf{Q}_n[k+1]$ to neighbors $m \in \mathcal{J}_n$.

end for

return
$$\mathbf{Q}_n[\infty], \mathbf{p}_n[\infty], \mathbf{o}_n[\infty]$$
.

Before moving on, a clarification on the notation used in Algorithm 1 is due. To define matrix Ω_n in [S2]–[S4], observe first that the local sampling operator can be expressed as $\mathcal{P}_{\Omega_n}(\mathbf{z}) = \boldsymbol{\omega}_n \odot \mathbf{z}$, where \odot denotes Hadamard product, and the binary masking vector $\boldsymbol{\omega}_n \in \{0, 1\}^T$ has entries equal to 1 if the corresponding entry of \mathbf{z} is observed, and 0 otherwise. It is then apparent that the Hadamard product can be replaced with the usual matrix-vector product as $\mathcal{P}_{\Omega_n}(\mathbf{z}) = \Omega_n \mathbf{z}$, where $\Omega_n :=$ diag($\boldsymbol{\omega}_n$). Operators \otimes and vec[·] denote Kronecker product and matrix vectorization, respectively. Finally, the soft-thresholding operator is $\mathcal{S}_{\lambda_1}(\cdot) := \operatorname{sign}(\cdot) \max(|\cdot| - \lambda_1, 0)$.

Remark 4 (Elimination of Redundant Variables): Careful inspection of Algorithm 1 reveals that the redundant auxiliary variables $\{\bar{\mathbf{F}}_n^m, \tilde{\mathbf{F}}_n^m, \tilde{\mathbf{M}}_n^m\}_{n \in \mathcal{N}}^{m \in \mathcal{J}_n}$ have been eliminated. Each smart meter, say the *n*-th, does not need to separately keep track of all its non-redundant multipliers $\{\bar{\mathbf{M}}_n^m\}_{m \in \mathcal{J}_n}^m$, but only update their respective sums $\mathbf{S}_n[k] := 2 \sum_{m \in \mathcal{J}_n} \bar{\mathbf{M}}_n^m[k]$.

When employed to solve non-convex problems such as (P4), ADMM so far offers no convergence guarantees. However, there is ample experimental evidence in the literature which supports convergence of ADMM, especially when the non-convex problem at hand exhibits "favorable" structure [4]. For instance, (P4) is bi-convex and gives rise to the strictly convex optimization subproblems each time \mathcal{L}_c is minimized with respect to one of the group variables, which admit unique closed-form solutions per iteration [cf. (7)–(9)]. This observation and the linearity of the constraints suggest good convergence properties for the D-PCP algorithm. Extensive numerical tests including those presented in Section V demonstrate that this is indeed the case. While a formal convergence proof is the subject of ongoing investigation, the following proposition asserts that upon convergence, the D-PCP algorithm attains consensus and global optimality. For a proof (omitted here due to space limitations), see ([20], Appendix C).

Proposition 1: Suppose iterates $\{\mathbf{Q}_n[k], \mathbf{p}_n[k], \mathbf{o}_n[k]\}_{n \in \mathcal{N}}$ generated by Algorithm 1 converge to $\{\bar{\mathbf{Q}}_n, \bar{\mathbf{p}}_n, \bar{\mathbf{o}}_n\}_{n \in \mathcal{N}}$. If $\{\hat{\mathbf{X}}, \hat{\mathbf{O}}\}$ is the optimal solution of (P1), then $\bar{\mathbf{Q}}_1 = \bar{\mathbf{Q}}_2 = \dots = \bar{\mathbf{Q}}_N$. Also, if $\|\mathcal{P}_{\Omega}(\mathbf{Y} - \bar{\mathbf{P}}\bar{\mathbf{Q}}'_1 - \bar{\mathbf{O}})\| < \lambda_*$, then $\{\hat{\mathbf{X}} = \bar{\mathbf{P}}\bar{\mathbf{Q}}'_1, \hat{\mathbf{O}} = \bar{\mathbf{O}}\}$.

V. NUMERICAL TESTS

This section corroborates convergence and gauges performance of the D-PCP algorithm, when tested using synthetic and real load curve data.

A. Synthetic Data Tests

A network of N = 25 smart meters is generated as a realization of the random geometric graph model, meaning nodes are randomly placed on the unit square and two nodes communicate with each other if their Euclidean distance is less than a prescribed communication range of $d_c = 0.4$; see Fig. 1. The time horizon is T = 600. Entries of E are independent and identically distributed (i.i.d.), zero-mean, Gaussian with variance $\sigma^2 = 10^{-3}$; i.e., $e_{l,t} \sim \mathcal{N}(0, \sigma^2)$. Low-rank spatiotemporal load profiles with rank r = 3 are generated from the bilinear factorization model $\mathbf{X} = \mathbf{W}\mathbf{Z}'$, where \mathbf{W} and \mathbf{Z} are $N \times r$ and $T \times r$ matrices with i.i.d. entries drawn from Gaussian distributions $\mathcal{N}(0, 100/N)$ and $\mathcal{N}(0, 100/T)$, respectively. Every entry of **O** is randomly drawn from the set $\{-1, 0, 1\}$ with $\Pr(o_{n,t} = -1) = \Pr(o_{n,t} = 1) = 5 \times 10^{-2}$. To simulate missing data, a sampling matrix $\hat{\mathbf{\Omega}} \in \{0,1\}^{N imes T}$ is generated with i.i.d. Bernoulli distributed entries $o_{n,t} \sim \text{Ber}(0.7)$ (30%) missing data on average). Finally, measurements are generated as $\mathcal{P}_{\Omega}(\mathbf{Y}) = \mathbf{\Omega} \odot (\mathbf{X} + \mathbf{O} + \mathbf{V})$ [cf. (1)], and smart meter n has available the *n*-th row of $\mathcal{P}_{\Omega}(\mathbf{Y})$.

To experimentally corroborate the convergence and optimality (as per Proposition 1) of the D-PCP algorithm, Algorithm 1 is run with c = 1 and compared with the centralized benchmark (P1), obtained using the solver in [28]. Parameters $\lambda_1 = 0.0141$ and $\lambda_* = 0.346$ are chosen as suggested in [29]. For both schemes, Fig. 2 shows the evolution of the global estimation errors $e_X[k] := ||\mathbf{X}[k] - \mathbf{X}||_F / ||\mathbf{X}||_F$ and $e_O[k] := ||\mathbf{O}[k] - \mathbf{O}||_F / ||\mathbf{O}||_F$. It is apparent that the D-PCP algorithm converges to the centralized estimator, and as expected convergence slows down due to the delay associated with the information flow throughout the network. The test is also repeated for network sizes of N = 15 and 35 devices, to illustrate that the time till convergence scales gracefully as the network size increases. Finally, for N = 35 and with



Fig. 1. A simulated network graph with N = 25 nodes.



Fig. 2. Convergence of the D-PCP algorithm for different network sizes. D-PCP attains the same estimation error as the centralized solver.



Fig. 3. Evolution of the consensus error.

 $\bar{\mathbf{Q}}[k] := \sum_{n} \mathbf{Q}_{n}[k]/N$, Fig. 3 depicts the consensus error $e_{c,n}[k] := \|\mathbf{Q}_{n}[k] - \bar{\mathbf{Q}}[k]\|_{F}/\|\bar{\mathbf{Q}}[k]\|_{F}$ for three representative smart metering devices. In all cases the error decays rapidly to zero, showing that networkwide agreement is attained on the estimates $\mathbf{Q}_{n}[k]$.



Fig. 4. Normalized singular values of the 5×336 load profile matrix. A rank-three approximant will retain most of its energy.

B. Real Load Curve Data Test

Here, the D-PCP algorithm is tested on real load curve data. The dataset consists of power consumption measurements (in kW) for a government building, a grocery store, and three schools (N = 5) collected every fifteen minutes during a period of more than five years, ranging from July 2005 to October 2010. Data is downsampled by a factor of four, to yield one measurement per hour. For the present experiment, only a subset of the whole data is utilized for concreteness, where T = 336 was chosen corresponding to 336 hour periods. The normalized singular values of the 5×336 data matrix are 1,0.41,0.29,0.06,0.03; see also Fig. 4. Hence, a rank-three approximant to the data matrix will retain most of its energy. This observation motivates the proposed approach to cleansing and imputation, which capitalizes on the low-rank property of the spatio-temporal load profiles.

For the government building case, a snapshot of the available load curve data spanning the studied two-week period is shown in blue, e.g., in Fig. 5 (bottom). Weekday activity patterns can be clearly discerned from those corresponding to weekends, as expected for most government buildings; but different, e.g., for the load profile of the grocery store in Fig. 6 (bottom).

To run the D-PCP algorithm, an underlying communication graph was generated as in Section V-A. A randomly chosen subset of 30% of the measurements was removed to model missing data. For one of the schools and the government building data, Fig. 5 depicts the cleansed load curves that closely follow the measurements, but are smooth enough to avoid overfitting the abnormal energy peaks on the so-termed "building operational shoulders." Indeed, these peaks are in most cases identified as outliers. The effectiveness in terms of imputation of missing data is illustrated in Fig. 6 (identified outliers are not shown here); note how the cleansed (gray) load curve goes through the (red) missing data points. The relative error in predicting missing data is around 6%, and degrades to 8% when the amount of missing data increases to 50%.



Fig. 5. School and government building load curve data cleansing.



Fig. 6. Government building and grocery store load curve imputation, when 30% of the data are missing.

VI. CONCLUSION

A novel robust load curve cleansing and imputation method is developed in this paper, rooted at the crossroads of sparsitycognizant statistical inference, low-rank matrix completion, and large-scale distributed optimization. The adopted PCP estimator jointly leverages the low-intrinsic dimensionality of spatiotemporal load profiles, and the sparse (that is, sporadic) nature of outlying measurements. A separable reformulation of PCP is shown to be efficiently minimized using the ADMM, and gives rise to fully-decentralized iterations which can be run by a network of smart-metering devices. Comprehensive tests with synthetic and real load curve data demonstrate the effectiveness of the novel load cleansing and imputation approach, and corroborate the convergence and global optimality of the D-PCP algorithm.

An interesting future direction is to devise distributed realtime cleansing and imputation algorithms capable of processing load curve data acquired sequentially in time, and analyze their performance. In addition, it is of interest to rigorously establish convergence of the D-PCP algorithm. Such results could significantly broaden the applicability of ADMM for large-scale optimization over networks, even in the presence of non-convex but highly structured and separable cost functions.

APPENDIX ALGORITHMIC CONSTRUCTION

The goal is to show that [S1]–[S4] can be simplified to the iterations tabulated under Algorithm 1. Focusing first on [S3], (8) decomposes into N ridge-regression sub-problems

$$\mathbf{p}_n[k+1] = \arg\min_{\mathbf{p}} \left\{ \|\mathbf{y}_n - \mathbf{Q}_n[k+1]\mathbf{p} - \mathbf{o}_n[k]\|_2^2 + \lambda_* \|\mathbf{p}\|_2^2 \right\}$$

which admit the closed-form solutions shown in Algorithm 1.

Moving on to [S4], from the decomposable structure of the augmented Lagrangian [cf. (4)] (9) decouples into per-node scalar Lasso subtasks (note that $\mathbf{Q}_n := [\mathbf{q}_{n,1}, \dots, \mathbf{q}_{n,T}]'$)

$$o_{n,t}[k+1] = \arg\min_{\mathbf{o}} \left\{ \frac{1}{2} \left(y_{n,t} - \mathbf{q}'_{n,t}[k+1]\mathbf{p}_n[k+1] - o \right)^2 + \lambda_1 |o|_1 \right\}, \quad t = 1, \dots, T$$
$$= S_{\lambda_1}(y_{n,t} - \mathbf{q}'_{n,t}[k+1]\mathbf{p}_n[k+1]), t = 1, \dots, T$$

and $\sum_{n=1}^{N} |\mathcal{J}_n|$ additional unconstrained QPs

$$\bar{\mathbf{F}}_{n}^{m}[k+1] = \tilde{\mathbf{F}}_{n}^{m}[k+1] \\
= \arg \min_{\bar{\mathbf{F}}_{n}^{m}} \left\{ -\left\langle \bar{\mathbf{M}}_{n}^{m}[k] + \tilde{\mathbf{M}}_{n}^{m}[k], \bar{\mathbf{F}}_{n}^{m} \right\rangle \\
+ \frac{c}{2} \left[\|\mathbf{Q}_{n}[k+1] - \bar{\mathbf{F}}_{n}^{m}\|_{F}^{2} \\
+ \|\mathbf{Q}_{m}[k+1] - \bar{\mathbf{F}}_{n}^{m}\|_{F}^{2} \right] \right\}$$
(10)

which admit the closed-form solutions

$$\bar{\mathbf{F}}_{n}^{m}[k+1] = \bar{\mathbf{F}}_{n}^{m}[k+1] \\
= \frac{1}{2c} \left(\bar{\mathbf{M}}_{n}^{m}[k] + \tilde{\mathbf{M}}_{n}^{m}[k] \right) \\
+ \frac{1}{2} (\mathbf{Q}_{n}[k+1] + \mathbf{Q}_{m}[k+1]). \quad (11)$$

Note that in formulating (10), $\tilde{\mathbf{F}}_n^m$ was eliminated using the constraints $\bar{\mathbf{F}}_n^m = \tilde{\mathbf{F}}_n^m$ defining \mathcal{C}_F . Using (11) to eliminate $\bar{\mathbf{F}}_n^m[k]$ and $\tilde{\mathbf{F}}_n^m[k]$ from (5) and (6) respectively, a simple induction argument establishes that if the initial Lagrange multipliers obey $\bar{\mathbf{M}}_n^m[0] = -\tilde{\mathbf{M}}_n^m[0] = \mathbf{0}$, then $\bar{\mathbf{M}}_n^m[k] = -\tilde{\mathbf{M}}_n^m[k]$ for all $k \ge 0$, where $n \in \mathcal{N}$ and $m \in \mathcal{J}_n$. The set $\{\tilde{\mathbf{M}}_n^m\}$ of multipliers has been shown redundant, and (11) readily simplifies to

$$\bar{\mathbf{F}}_{n}^{m}[k+1] = \tilde{\mathbf{F}}_{n}^{m}[k+1] = \frac{1}{2}(\mathbf{Q}_{n}[k+1] + \mathbf{Q}_{m}[k+1]).$$
(12)

It then follows that $\bar{\mathbf{F}}_n^m[k] = \bar{\mathbf{F}}_m^n[k]$ for all $k \ge 0$, an identity that will be used later on. By plugging (12) in (5), the (non-redundant) multiplier updates become

$$\bar{\mathbf{M}}_{n}^{m}[k] = \bar{\mathbf{M}}_{n}^{m}[k-1] + \frac{c}{2}[\mathbf{Q}_{n}[k] - \mathbf{Q}_{m}[k]], \quad n \in \mathcal{N}, \ m \in \mathcal{J}_{n}.$$
(13)

If $\bar{\mathbf{M}}_{\underline{n}}^{m}[0] = -\bar{\mathbf{M}}_{\underline{m}}^{n}[0] = \mathbf{0}$, then the structure of (13) reveals that $\bar{\mathbf{M}}_{\underline{n}}^{m}[k] = -\bar{\mathbf{M}}_{\underline{m}}^{n}[k]$ for all $k \geq 0$, where $n \in \mathcal{N}$ and $m \in \mathcal{J}_{n}$.

The minimization (9) in [S4] also decouples in N simpler sub-problems, namely

$$\begin{aligned} \mathbf{Q}_{n}[k+1] &= \arg\min_{\mathbf{Q}} \left\{ \frac{1}{2} \| \mathbf{\Omega}_{n}(\mathbf{y}_{n} - \mathbf{Q}\mathbf{p}_{n}[k] - \mathbf{o}_{n}[k]) \|_{2}^{2} \\ &+ \frac{\lambda_{*}}{2N} \| \mathbf{Q} \|_{F}^{2} + \sum_{m \in \mathcal{J}_{n}} \left\langle \bar{\mathbf{M}}_{n}^{m}[k] + \tilde{\mathbf{M}}_{m}^{n}[k], \mathbf{Q} \right\rangle \\ &+ \frac{c}{2} \sum_{m \in \mathcal{J}_{n}} \left(\| \mathbf{Q} - \bar{\mathbf{F}}_{n}^{m}[k] \|_{F}^{2} + \| \mathbf{Q} - \tilde{\mathbf{F}}_{m}^{n}[k] \|_{F}^{2} \right) \right\} \\ &= \arg\min_{\mathbf{Q}} \left\{ \frac{1}{2} \| \mathbf{\Omega}_{n}(\mathbf{y}_{n} - \mathbf{Q}\mathbf{p}_{n}[k] - \mathbf{o}_{n}[k]) \|_{2}^{2} \\ &+ \frac{\lambda_{*}}{2N} \| \mathbf{Q} \|_{F}^{2} + \langle \mathbf{S}_{n}[k], \mathbf{Q} \rangle \\ &+ c \sum_{m \in \mathcal{J}_{n}} \left\| \mathbf{Q} - \frac{\mathbf{Q}_{n}[k] + \mathbf{Q}_{m}[k]}{2} \right\|_{F}^{2} \right\} \end{aligned}$$
(14)

where in deriving (14) it was used that: i) $\bar{\mathbf{M}}_{n}^{m}[k] = \mathbf{M}_{n}^{m}[k]$ which follows from the identities $\bar{\mathbf{M}}_{n}^{m}[k] = -\tilde{\mathbf{M}}_{n}^{m}[k]$ and $\bar{\mathbf{M}}_{n}^{m}[k] = -\bar{\mathbf{M}}_{m}^{m}[k]$ established earlier; ii) the definition $\mathbf{S}_{n}[k] := 2\sum_{m \in \mathcal{J}_{n}} \bar{\mathbf{M}}_{n}^{m}[k]$; and iii) the identity $\bar{\mathbf{F}}_{n}^{m}[k] = \tilde{\mathbf{F}}_{m}^{n}[k]$ which allows one to merge the identical quadratic penalty terms and eliminate both $\bar{\mathbf{F}}_{n}^{m}[k]$ and $\tilde{\mathbf{F}}_{m}^{n}[k]$ using (12). Problem (14) is again an unconstrained QP, which is readily solved in closed form by, e.g., vectorizing \mathbf{Q} and examining the first-order condition for optimality.

Finally, note that upon scaling by two the recursions (13) and summing them over $m \in \mathcal{J}_n$, the update recursion for $\mathbf{S}_n[k]$ in Algorithm 1 follows readily.

ACKNOWLEDGMENT

The authors would like to thank NorthWrite Energy Group and Prof. Vladimir Cherkassky (Dept. of ECE, University of Minnesota) for providing the data analyzed in Section V-B.

REFERENCES

- [1] A. Abur and A. Gomez-Exposito, *Power System State Estimation: Theory and Implementation.* New York: Marcel Dekker, 2004.
- [2] A. R. Bergen and V. Vittal, *Power System Analysis*. Upper Saddle River, NJ, USA: Prentice-Hall, 2000.
- [3] D. Bertsekas and J. Tsitsiklis, Parallel and Distributed Computation: Numerical Methods. Belmont, MA, USA: Athena-Scientific, 1999.
- [4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, pp. 1–122, 2010.
- [5] E. J. Candes, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," J. ACM, vol. 58, no. 1, pp. 1–37, 2011.
- [6] V. Chandrasekaran, S. Sanghavi, P. R. Parrilo, and A. S. Willsky, "Rank-sparsity incoherence for matrix decomposition," *SIAM J. Optim.*, vol. 21, no. 2, pp. 572–596, 2011.
- [7] J. Chen, W. Li, A. Lau, J. Cao, and K. Wang, "Automated load curve data cleansing in power systems," *IEEE Trans. Smart Grid*, vol. 1, pp. 213–221, Sep. 2010.
- [8] D. Duan, L. Yang, and L. L. Scharf, "Phasor state estimation from PMU measurements with bad data," in *Proc. IEEE Workshop Comp. Adv. Multi-Sensor Adaptive Proc.*, San Juan, Puerto Rico, Dec. 2011.
- [9] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Electrical Eng. Dept., Stanford University, Stanford, CA, USA, 2002.

- [10] H. Gharavi, A. Scaglione, M. Dohler, and X. Guan, "Technical challenges of the smart grid: From a signal processing perspective," *IEEE Signal Process. Mag.*, vol. 29, no. 5, pp. 12–13, Sep. 2012.
- [11] Z. Guo, W. Li, A. Lau, T. Inga-Rojas, and K. Wang, "Detecting X-outliers in load curve data in power systems," *IEEE Trans. Power Syst.*, vol. 27, pp. 875–884, May 2012.
- [12] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. New York: Springer, 2009.
- [13] P.-S. Huang, S. D. Chen, P. Smaragdis, and M. Hasegawa-Johnson, "Singing-voice separation from monaural recordings using robust principal component analysis," in *Proc. Int. Conf. Acoust., Speech, Signal Process.*, Kyoto, Japan, Mar. 2012, pp. 57–60.
 [14] V. Kekatos and G. B. Giannakis, "Distributed robust power system
- [14] V. Kekatos and G. B. Giannakis, "Distributed robust power system state estimation," *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 1617–1626, May. 2013.
- [15] O. Kosut, L. Jia, J. Thomas, and L. Tong, "Malicious data attacks on the smart grid," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 645–658, Dec. 2011.
- [16] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Message passing for dynamic network energy management," Tech. Rep., Apr. 2012.
- [17] Z. Lin, A. Ganesh, J. Wright, L. Wu, M. Chen, and Y. Ma, "Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix," UIUC Tech. Rep. UILU-ENG-09-2214, Jul. 2009.
- [18] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," in *Proc. Conf. Comput. Commun. Security*, Chicago, IL, USA, Nov. 2009, pp. 9–13.
- [19] Z.-Q. Luo, "On the linear convergence of the alternating direction method of multipliers," Tech. Rep., Aug. 2012 [Online]. Available: http://arxiv.org/abs/1208.3922
- [20] M. Mardani, G. Mateos, and G. B. Giannakis, "Decentralized sparsityregularized rank minimization: Algorithms and applications," *IEEE Trans. Signal Process.*, 2013 [Online]. Available: http://arxiv.org/pdf/ 1211.6950.pdf, to be published
- [21] M. Mardani, G. Mateos, and G. B. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low rank," *IEEE J. Sel. Topics Signal Process.*, vol. 7, no. 1, pp. 50–66, Feb. 2013.
- [22] G. Mateos and G. B. Giannakis, "Robust nonparametric regression via sparsity control with application to load curve data cleansing," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 1571–1584, Apr. 2012.
- [23] L. Mili, M. G. Cheniae, and P. J. Rousseeuw, "Robust state estimation of electric power systems," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 41, no. 5, pp. 349–358, May 1994.
- [24] K. M. Rogers, R. D. Spadoni, and T. J. Overbye, "Identification of power system topology from synchrophasor data," in *Proc. Power Syst. Conf. Expo.*, Phoenix, AZ, USA, Mar. 2011.
- [25] M. Shahidehpour, H. Yamin, and Z. Li, Market Operations in Electric Power Systems: Forecasting, Scheduling, and Risk Management. New York: Wiley-IEEE Press, 2002.
- [26] "The smart grid: An introduction," 2008, U.S. Department of Energy [Online]. Available: http://www.oe.energy.gov/SmartGridIntroduction.htm
- [27] W. Wulf, "Great achievements and grand challenges," The Brattle Group, Freeman, Sullivan and Co., and Global Energy Partners, LLC, Tech. Rep. 3/4, Fall 2010 [Online]. Available: http://www.greatachievements.org/

- [28] X. M. Yuan and J. Yang, "Sparse and low-rank matrix decomposition via alternating direction methods," *Pac. J. Optim.*, 2012, to be published.
- [29] Z. Zhou, X. Li, J. Wright, E. Candes, and Y. Ma, "Stable principal component pursuit," in *Proc. Int. Symp. Inf. Theory*, Austin, TX, USA, Jun. 2010, pp. 1518–1522.



Gonzalo Mateos (M'12) received his B.Sc. degree in electrical engineering from Universidad de la República (UdelaR), Montevideo, Uruguay, in 2005 and the M.Sc. and Ph.D. degrees in electrical and computer engineering from the University of Minnesota, Minneapolis, MN, USA, in 2009 and 2011.

Since 2012, he has been a Postdoctoral Research Associate with the Department of Electrical and Computer Engineering and the Digital Technology Center, University of Minnesota. Since 2003, he is

an Assistant with the Department of Electrical Engineering, UdelaR. From 2004 to 2006, he worked as a Systems Engineer at Asea Brown Boveri (ABB), Uruguay. His research interests lie in the areas of communication theory, signal processing, and networking. His current research focuses on distributed signal processing, sparse linear regression, and statistical learning for social data analysis and network health monitoring.



Georgios B. Giannakis (F'97) received his Diploma in electrical engineering from the National Technical University of Athens, Greece, 1981. From 1982 to 1986 he was with the University of Southern California (USC), where he received his M.Sc. in electrical engineering, 1983, M.Sc. in mathematics, 1986, and Ph.D. in electrical engineering, 1986. Since 1999 he has been a professor with the University of Minnesota, Minneapolis, MN, USA, where he now holds an ADC Chair in Wireless Telecommunications in the ECE Department, and serves as director of the

Digital Technology Center.

His general interests span the areas of communications, networking and statistical signal processing-subjects on which he has published more than 350 journal papers, 590 conference papers, 20 book chapters, two edited books and two research monographs (h-index 105). Current research focuses on sparsity in signals and systems, wireless cognitive radios, mobile ad hoc networks, renewable energy, power grid, gene-regulatory, and social networks. He is the (co-) inventor of 21 patents issued, and the (co-) recipient of 8 best paper awards from the IEEE Signal Processing (SP) and Communications Societies, including the G. Marconi Prize Paper Award in Wireless Communications. He also received Technical Achievement Awards from the SP Society (2000), from EURASIP (2005), a Young Faculty Teaching Award, and the G. W. Taylor Award for Distinguished Research from the University of Minnesota. He is a Fellow of EURASIP, and has served the IEEE in a number of posts, including that of a Distinguished Lecturer for the IEEE-SP Society.