

Principled Data Completion of Network Constraints for Day Ahead Auctions in Power Markets

Ioan Alexandru Puiu, *University of Oxford*, Raphael Andreas Hauser, *University of Oxford*

Abstract—Network constraints play a key role in the price finding mechanism for European Power Markets, but historical data is very sparse and usually insufficient for many quantitative applications. We reconstruct the constraints data, known as the Power Transmission Distribution Factors (PTDFs) and Remaining Available Margins (RAMs), by first recovering the underlying time dependent signals known as the Generation Shift Keys (GSKs) and Phase Angles (PAs), and the electricity grid characteristics, via a mathematical optimisation problem. This is solved by exploiting marginal convexity in certain subspaces via alternating minimisation. The GSKs and PAs are then mapped to the PTDFs and RAMs, using the grid structure. Our reconstruction achieves in-sample relative errors of 22.3% and 15.9% for the PTDFs and RAMs respectively, while the out-of-sample ones are 25.4% and 18.1% respectively. We further show that our model outperforms the naive approach, and that the reconstructed GSKs and PAs recover specific structure.

Index Terms—Central Western European Power Markets, Day Ahead Auction, network constraints, mathematical optimisation, data completion, data reconstruction, Flow based market coupling, GSK, PTDF, RAM, PSDF.

I. INTRODUCTION AND BACKGROUND INFORMATION

POWER markets are frequently organised as Day Ahead Auctions (DAA) due to the difficulty of storing electricity, coupled with uncertainty in demand and production patterns. European Electricity Markets have been concerned with clearing the market at a unique zonal price, where zones usually correspond to countries. We consider the Flow Based Market Coupling (FBMC) approach, which enables zonal market clearing, by using physical principles for tracking the flow loads in the grid. This has become increasingly popular for multiple reasons, when compared to the Available Transfer Capacity (ATC) approach [2], which was used prior to May 2015. First, it offers increased trading volume, and the European Commission targets a more efficient use of energy sources, which they believe can be achieved easier by allowing for higher inter-zonal electricity exchange [18]. Secondly, by allowing for higher inter-zonal trade, market clearing becomes less constrained resulting in larger Social Welfare (SW) on average [2]. Finally, as a result of more accurate grid modelling when compared to ATC, the FBMC approach was observed to offer increased operational safety [2]. FBMC was pioneered in the Central Western European (CWE) Power Market, which we use as a case study for our methodology.

The core of the FBMC approach is a set of linear constraints enforced at each time step on the zonal (country) production

levels, which ensure that the electricity grid always operates within the safety limits [26], [3]. This set of constraints governs the feasibility domain of power production and flow patterns, and thus ultimately determines zonal prices and their potential decoupling. However, only a very small subset of these constraints is published *ex-post* [6], which is generally insufficient for domain reconstruction or thorough statistical analysis that accounts for the data complexity. Nevertheless, for risk management reasons, market participants must understand their risk exposure in a wide range of scenarios, some of which are far from the “*expected outcomes*” and closer to tail events. Under such circumstances the active set of constraints can be totally different than the ones published in the past, and thus the full set of constraints is needed. Furthermore, the constraint set is sharp, in the sense that a small change in the dispatch vector can result in totally different active constraints. This, coupled with the fact that small changes in constraints may have large impact on the outcomes of market participants with bids close to the margin, makes it imperative for market participants to understand the full set of constraints as well as possible. To further complicate matters for market participants, these constraints change rather erratically in time and thus naive techniques are inappropriate for sound quantitative decision making. For the DAA participants, this translates to reduced ability to manage risk, and ultimately to increased market price inefficiencies.

To the best of our knowledge, obtaining the full set of constraints was not addressed before in the literature. Thus, we propose a methodology that uses the physical principles of the FBMC approach and publicly available data to reconstruct the full set of linear constraints, known as the combination of Power Transmission Distribution Factors (PTDFs) and Remaining Available Margins (RAMs). To this end, we first recover the electricity grid structure, as well as the underlying time-dependent signals known as the Generation Shift Keys (GSKs) and Phase Angles (PAs). The GSKs and PAs are then transformed to the PTDFs and RAMs using maps determined by the graph structure. We achieve the reconstruction by formulating an optimisation model based on the principles of FBMC, regularised by the available prior knowledge. We exploit the structure of the optimisation problem by solving via alternating minimisation: stochastic optimisation is used for the non-linear and computationally intensive parts of the objective, and convex quadratic programming for the marginally convex components. We are able to show good results for the reconstruction of PTDFs and RAMs in out-of-sample tests, outperforming the naive approach. Further, we analyse the structure of the recovered grid characteristics, as well as of the

GSKs and PAs signals, and we show that we recover specific clustering structure, suggesting good reconstruction quality.

Our approach provides a different view to the ones available in the current literature, which are directly or indirectly focused on the TSOs's point of view: assuming knowledge of certain data, parameters and choices, explain or analyse the FBMC methodology [27], [29]. By contrast, our novel methodology takes the view of the market participant, considering its interests, available information, and unknowns. Participants' position in the electricity market are totally different than the TSOs': market participants have different goals, do not know what the operators know, and they do not know what the operators forecast. Calculations of the GSKs based on fundamental equations require TSOs' prediction of the market outcomes in advance, as explained in [27]. However, these are not known by the market participants, and should not be mistaken for their own predictions of the market, which are likely to be obtained by using different methodologies or even different data. Further, market participants fully replicating TSOs' methodology may require too much computational and labour effort for their end goal: they are not concerned with managing the grid, but rather with optimising their risk-return exposure in the market. While some of these predictions may be available one day ahead, it is insufficient for market participants' planning horizon, which is a much longer period, and the corresponding data is unavailable. Since the market participants do not have the TSOs' prediction of the market outcome, they cannot use a purely model-based TSO approach, motivating the need for our proposed method.

The novel contribution of our approach is fourfold: (i) a graph model containing the quantitative physical characteristics is obtained, which is useful for any fundamental model that relies on the graph structure, (ii) the underlying GSKs and PAs signals are recovered, which give insight into the transmission system operators' behaviour, (iii) full reconstruction of the PTDFs and RAMs, allowing for full feasibility domain analysis by market participants, and (iv) our approach enables simulation and forecasting based on the reconstructed signals. Forecasting based on the available sparse PTDF and RAM data is extremely challenging. However, one can use the GSKs and PAs data recovered by our model, which is much lower dimensional, changes slower in time, and allows for explicitly incorporating impact of changes in the grid structure, or other fundamental factors.

The rest of the paper is structured as follows: background information is presented in Section II and a brief discussion of the Market Data is given in Section III. Sections IV & V present the main optimisation model and its algorithmic solution. Numerical results are presented in Section VI, and conclusions are drawn in Section VII.

II. BACKGROUND INFORMATION

Since our method is model-based, relying on the fundamentals of Flow-Based Market Coupling approach, in this section we present the main concepts and equations that are necessary for our method. The interested reader is referred to [25], [27], [28], [3] for more detailed reviews and to [8] for the official

documentation of the methodology. Each set of power injections and extractions decisions determines the flows across the lines of the network, via Kirchhoff's Law. Transmission system operators (TSO) monitor grid transmission lines (and other elements), and ensure that power flows do not exceed the safety limits by imposing the linear constraints given by the (PTDF, RAM) combination [8]. The constraints calculations in FBMC are based on the DC flow model presented in detail in [25], which involves a linear approximation based on three assumptions: (i) there are no losses on transmission lines, i.e. the resistance on lines is zero, (ii) the voltage profile is flat, i.e. all nodes have the same voltage, and (iii) the voltage angle differences are small for neighbouring nodes.

Let the power grid be represented by a graph \mathcal{G} , with a transmission lines set \mathcal{L} of cardinality $|\mathcal{L}| = L$, and node set \mathcal{N} with $|\mathcal{N}| = N$ nodes. Let the zone set for the considered graph be \mathcal{Z} with $|\mathcal{Z}| = Z$, and the node set of each zone z be \mathcal{N}_z with $|\mathcal{N}_z| = N_z$. The nodal PTDF is fully determined by the physical characteristics of the power grid [26] and is given by

$$P(b, A) = B\tilde{A}(\tilde{A}^T B\tilde{A})^{-1}, \quad (1)$$

where $\tilde{A} = A\Omega_{ref}^T$, with $A \in \mathbb{R}^{L \times N}$ being the line-node incidence matrix of the power grid defined as

$$A_{ji} = \begin{cases} +1 & \text{if line } j \text{ originates in node } i \\ -1 & \text{if line } j \text{ terminates at node } i \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

and $\Omega_{ref} \in \mathbb{R}^{(N-1) \times N}$ is a projection matrix obtained by deleting row i_0 of the identity matrix $I_N \in \mathbb{R}^{N \times N}$, with i_0 some chosen reference node [25]. Further, the diagonal susceptance matrix $B = \text{diag}(b)$ aggregates the line susceptances b_j , $\forall j \in \mathcal{L}$. The nodal power changes $\Delta p_N = p_N - g^0 \in \mathbb{R}^N$ are then mapped through the Nodal PTDF to line power flow changes $\Delta p_L = p_L - F^{ref} \in \mathbb{R}^L$ by

$$\Delta p_L = P(b, A)\Delta p_N, \quad (3)$$

where Δp_N and Δp_L are defined with respect to (g^0, F^{ref}) , known as the *base case*, which we describe later. Negative and positive signs in p_N indicate nodal consumption and production respectively. The interpretation of $P_{ji}(b, A)$ is the change in flow along line j that is caused by injection of one unit of power at node i for extraction at the reference node i_0 .

Note that the directionality of A and the choice of the reference node i_0 are arbitrary but necessary [25], with the only requirement that these are fixed once and for all. Similarly, the phase angles α^L are mapped to power flow changes through the Phase Shift Distribution Factor (PSDF) [25], defined as

$$PSDF(b, A) := B - B\tilde{A}(\tilde{A}^T B\tilde{A})^{-1}\tilde{A}^T B, \quad (4)$$

which yields the linear mapping $\Delta^\alpha p_L = PSDF(b, A)\alpha^L$, where $\Delta^\alpha p_L$ is the line power changes given by the line phase angles α^L . Note that this linearisation is accurate only for sufficiently small α^L .

To enforce the constraints, in each zone, the TSOs define the *base case* which includes a nodal power pattern [8]. We here denote this pattern as $g_{z,t}^0 \in \mathbb{R}^{N_z}$. By collating these vectors for all zones $z \in \mathbb{Z}$, one obtains $g_t^0 = \big\|_{z \in \mathbb{Z}} (g_{z,t}^0) \in \mathbb{R}^N$,

where $\|^\vee(\cdot)$ represents vertical concatenation. The *base case* is defined *ex-ante* and is the TSO's best guess of the nodal production and consumption patterns sitting at day $D-2$, for each hour t of day D . This is implicitly used in the DAA welfare maximisation problem for price finding and order acceptance at day $D-1$ [20]. Note that unless otherwise specified, from here onward, t denotes hourly time indexing.

Constraining the DAA auction at nodal level is not practical, as power is not traded on a per node (power-plant) basis. Instead, nodes are collated into zones, and power injection at each node within a zone (usually country) is treated as equivalent and traded at a zonal price. However, it is not possible to reduce the power injection decisions at all nodes within a zone to a single equivalent zonal power injection decision, because flows within each zone must also be kept feasible. Nevertheless, the impact can be approximated: for each hourly delivery period, the nodal power levels within zone z are assumed to be described by a vector

$$p_{z,t}^N = g_{z,t}^0 + \Delta P_{z,t} g_{z,t}^1, \quad (5)$$

where $\Delta P_{z,t} \in \mathbb{R}$ is the overall power deviation from the *base case* scenario and $g_{z,t}^1 \in \mathbb{R}^{N_z}$ is the Generation Shift Key vector representing the power injection at each node in zone z at time t , under a shift scenario where one extra unit of power is injected across zone z and extracted at the reference node i_0 . Similarly to $g_{z,t}^0$, the GSK is also defined at day $D-2$, but obeys $1^T g_{z,t}^1 = 1$, and we define the extended version of $g_{z,t}^1$ for all nodes \mathcal{N} as $\bar{g}_{z,t}^1$ with, $(\bar{g}_{z,t}^1)_i = 0, \forall i \notin \mathcal{N}_z$, instead of restraining $\bar{g}_{z,t}^1$ to the N_z elements in zone z .

Thus, the deviations of the power generation pattern from the base case are modelled as a 1-dimensional set: for each zone and time instance, once the base case and the GSK are fixed, the power production levels vector in (5), is fully defined up to a one dimensional parameter, $\Delta P_{z,t}$. In other words, the set of possible values of $p_{z,t}^N$ lie on a line parametrized by the scalar $\Delta P_{z,t}$. This means that once the base case and the GSKs are fixed, the only parameter that influences the estimated flow values in every line is $\Delta P_{z,t}$. This is clearly a massive simplification. However, this can be improved by considering a set of outage scenarios \mathcal{S} for the same hour slot, because although $p_{z,t}^N$ is fixed for fixed $\Delta P_{z,t}$, $P_s(b, A)$ is different for each s , and thus so is the line flows vector $P_s(b, A)p_{z,t}^N$ for different s . This means that multiple directions along which line flows levels can change are estimated and bounded via constraints. Most commonly, an outage s refers to a line j , which effectively accounts for removing entry b_j and row $A_{j,:}$, and we denote the resulting quantities as b^s and A^s . We restrict our attention only to line outages here. For each $s \in \mathcal{S}$ and $t \in \mathcal{T}$ the *line-zonal* PTDF can be now obtained as [1]

$$PTDF_{s,t} = P_s(b, A)G_{s,t}^1, \quad (6)$$

where the subscript s denotes explicit scenario dependence, $P_s(b, A) := P(b^s, A^s)$, and $G_{s,t}^1 = \|\|_{z \in \mathcal{Z}}^H(\bar{g}_{s,t}^1, z) \in \mathbb{R}^{N \times Z}$ is the horizontal concatenation of the GSKs for each zone under scenario s at time t . Note that when horizontally concatenating

vectors we view them as a matrix with one column, and thus $G_{s,t}^1$ is a matrix. The constraints can now be written as

$$r_{s,t} \leq PTDF_{s,t} \delta p_t^Z \leq R_{s,t}, \quad (7)$$

where the RAMs are defined as

$$r_{s,t} = -F_{s,t}^{max} - F_{s,t}^{ref} + FRM_{s,t} + FAV_{s,t} \quad (8)$$

$$R_{s,t} = F_{s,t}^{max} - F_{s,t}^{ref} - FRM_{s,t} - FAV_{s,t} \quad (9)$$

and $\delta p_t^Z := \|\|_{z \in \mathcal{Z}}^V(\Delta P_{z,t}) \in \mathbb{R}^Z$ is the vector of zonal power changes, with respect to the base case values given by $1^T g_{z,t}^0$ [6]. By writing $PSDF_s(b, A) := PSDF(b^s, A^s)$, the reference flows are expressed as

$$F_{s,t}^{ref} = P_s(b, A)g_{s,t}^0 + V_0^2 PSDF_s(b, A)\alpha_t^L, \quad (10)$$

where V_0 is the flat voltage profile [25]. The FRMs are introduced to account for TSO's grid modelling errors, which mainly includes the linearisation simplification but also flow measurements uncertainties [6]. They are computed for each line as a multiple of the standard deviation of the error series defined as the difference between measured and computed flows. The final adjustment values (FAV) incorporates TSO information not accounted for in the DC model [6]. However, this term is often omitted in reality and in the literature [16], and thus we neglect this term for simplicity. Finally $F_{s,t}^{max} \in \mathbb{R}^L$ is the vector of maximum allowed flows, which is generally constant in time but can vary when grid conditions change.

Ensuring that every single transmission line $l \in \mathcal{L}$ satisfies the feasibility constraints (7) is computationally intensive and somewhat redundant [6]. Thus, the TSOs monitor a much smaller set of *Critical Network Elements* (CNEs), also known as *Critical Branches* (CBs), denoted as $\mathcal{L}_{cne} \subset \mathcal{L}$ with cardinality $|\mathcal{L}_{cne}| = L_{cne}$ [1]. We stick with the newer, CNEs notation. This set is determined by considering the transmission lines whose flow is significantly affected by possible remedial actions taken by the TSOs when managing the grid [8]. Let $\Omega_{cne} \in \mathbb{R}^{L_{cne} \times L}$ be the canonical projection from the space of all lines, \mathcal{L} , to the critical ones, \mathcal{L}_{cne} . The constraints can then be re-written by multiplying equation (7) by Ω_{cne} , and while the form is preserved, the number of constraints reduces dramatically.

The critical lines constraints version of (7) are then used in conjunction to the Welfare Maximisation Problem (WMP) to obtain the optimal δp_t^Z values which ultimately incorporate the accepted orders and determine the zonal prices [20].

III. ANALYSIS OF THE INPUT DATA

While the general framework of our model is applicable to any FBMC DAA market, it is worth noting that every single power market has its unique features [12], and thus any fundamental model needs to be adapted accordingly. Here we use the CWE DAA Power Market as a case study, and for added clarity we describe its data structure before formulating the model. The CWE Power Market is composed of five zones of interest: Austria, Belgium, Germany, France and the Netherlands. There are four major data sets: (i) the constraints data including the line-zonal PTDF and RAM [7], (ii) the

power production data [15], (iii) the grid data [14], [22], and (iv) other time series data, typically recorded per zone, which also include their day-ahead estimates [15]. Unless otherwise specified, the data is available at [15]. We restrict our attention to the first 210 days of 2019, which gives $|\mathcal{T}| = 5040$ data points for each hourly time series. The reason for this is three-fold:

- 1) data sparsity is roughly the same at every t , and thus considering more time instances does not increase the density of the data set, and thus does not directly increase the ability to recover the GSKs and the PAs, which are time dependent signals.
- 2) If the grid structure was fixed in reality, considering more time instances would increase the ability to recover the physical characteristics of the grid. However, in reality the grid is continuously changing, and thus considering windows that are too large might be detrimental when not tracking time dependent changes of the grid. This issue merits much attention on its own, but is outside the scope of this study.
- 3) Most importantly, the resulting optimisation problem is very computationally intense. As will become clear later, considering a larger can data-set drastically increases the computational time.

We next describe data preprocessing in more detail and analyse the data briefly.

A. Constraints Data

The data set, which we denote as $\mathcal{D}_{\mathcal{LST}}$, consists of hourly recorded $PTDF$, RAM , F_{ref} , F_{max} , and FRM values, on transmission lines corresponding to each zone of interest and is available at [7]. A brief description of the data features is given below:

- 1) at each time, we observe data only of a small subset of the critical lines \mathcal{L}_{cne} , corresponding to active constraints in the optimisation algorithm [20].
- 2) At each fixed time, for each critical branch, $l \in \mathcal{L}_{cne}$, we potentially observe multiple corresponding *Contingencies* (C), also known as *Critical Outages* (CO), $s \in \mathcal{S}$. This gives the following data structure for each line-scenario-time combination $\mathcal{D}_{l,s,t} = \{PTDF, RAM, F_{ref}, F_{max}, FRM, FAV\}$. Since we do not observe all (l, s, t) pairs for $l \in \mathcal{L}_{cne}$, $s \in \mathcal{S}$, and $t \in \mathcal{T}$, we denote the available ones as \mathcal{LST} , and $\mathcal{ST} := \{(s, t) : (l, s, t) \in \mathcal{LST}\}$.
- 3) The chosen line orientation is not given in the data, and we account for this in our model,
- 4) the constraint data is independent of the grid data, and thus we map this to the grid data described in subsection C, by matching the node names specified in the constraints data, with the ones in the grid data.

We observe 106 CNEs and 152 Contingencies in the considered time interval, and successfully map 102 and 133 respectively. This gives $|\mathcal{LS}| = 336$ line-scenario combinations observed from the data, which is much smaller than the total number of possible combinations, $|\mathcal{L}_{cne}||\mathcal{S}|$. This is because

we do not observe every possible line-scenario pair (l, s) in the data.

The distributions of number of recorded data points for each l or (l, s) are not uniform but tends to resemble a power law, as observed in Figure 1. We thus observe that most of the

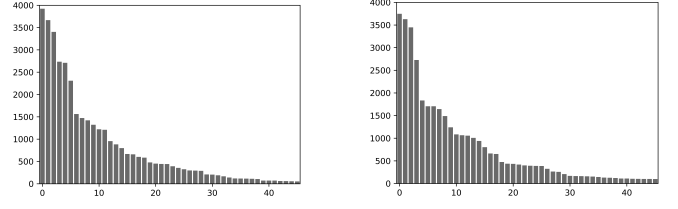


Fig. 1. Number of recorded data points for the first 45 lines (left) and scenarios (right), sorted in descending order of recorded information amount.

information is recorded on a small fraction of lines as well as (l, s) combinations. Adding the numbers in Figure 1 and dividing by $|\mathcal{LS}||\mathcal{T}|$, we see that the data amounts for only 2.3% of all CNEC (i.e (l, s)) combinations tracked by the TSOs. Clearly, the observed data is very sparse and insufficient for many applications.

B. Power Production Data

The generation units data can be found at [15] and is comprised of hourly production levels for named power plants. We record 35 Austrian, 42 Belgian, 173 German, 142 French and 45 Dutch power plants. A data extraction algorithm was written to download the data. The data accounts for most of the non-renewable production in these countries, but does not include renewables. For each zone, we denote the plant set \mathcal{GU}_z , with $|\mathcal{GU}_z| = N_p^z$, the vector of production at each time by ϕ_t^z and the concatenation over zones as $\mathcal{GU} = \cup_{z \in \mathcal{Z}} \mathcal{GU}_z$, with $|\mathcal{GU}| = N_p$, and $\phi_t = ||_{z \in \mathcal{Z}}^V (\phi_t^z)$. Further $\Omega_N \in \mathbb{R}^{N_p \times N}$ is the canonical projection from all nodes to power plant nodes. Note that no geographical coordinates, ID, or any other indication of power plants' connection to the grid is given. Thus, we locate these mostly by hand since standard algorithms generally give unsatisfactory results. This is clearly a limitation and a source of error, is a solution that enables us to use the methodology presented in this paper. We later compensate for this by allowing some optimisation over the location of the power plants and their connection to the grid, based on observed data, as presented in Section IV-B. This is generally the case for any other task performed by hand: while this input data is important and prone to error, during the optimisation process we allow for updating our best guess based on observed data.

C. Grid Data

Most of the consumption grid nodes and edges are obtained by aggregating information from SciGRID [22], APG [4], Elia [13] and ENSTOE [14], where the latter is used mostly for corrections by hand and visualisation purposes. We obtain the grid data as follows:

- 1) SciGRID [22] data is used as a starting point. This contains node and line information along with the geographical coordinates. Resistances (R), reactances (X),

and susceptances (b) can be computed by using the given per-unit length constants, and as per the procedure described in [11]. We discard the data for Austria and Belgium due to incompleteness.

- 2) We use the data available at [4] and [13] to add the node and line information for Austria and Belgium. This requires geographical mapping. Since the Dutch grid from SciGRID is incomplete, we complete this by hand using [14]. Any missing zone-connecting or critical lines are added by hand as per [14]. Any lines with unknown length are assigned $1.1 \cdot \text{geod}_l$, where geod_l is the geodesic distance between the nodes connected by line l . The susceptances are then computed as per [11]. This gives the final consumption grid.
- 3) Following the geographical locating described in subsection III-B the power plants are connected to the closest node to yield the final CWE grid model.

We denote the resulting susceptance vector by b_0 . The graph structure is shown in Figure 2. Note that the topology of $\mathcal{G} = (V, E)$ fully defines A , and the only physical characteristic needed are the line susceptances, b . To the best of our knowledge, this procedure makes use of the best available public information, but this will still introduce some error. Nevertheless, as we shall see in Section IV, we only use this for initialisation and regularisation in our optimisation model, and thus inaccuracies in the graph structure (A, b) can be tolerated.

D. Other time series data

The rest of the data used in our model are: (i) zonal time series for the demand $\{D_t^z\}_{t \in \mathcal{T}}$, (ii) import and export giving the Net Export series $\{NEX_t^z\}_{t \in \mathcal{T}}$, (iii) solar $\{\phi_t^{s,z}\}_{t \in \mathcal{T}}$ and wind $\{\phi_t^{w,z}\}_{t \in \mathcal{T}}$ renewable energy production, and (iv) the zonal prices $\{S_t^z\}_{t \in \mathcal{T}}$, all of which can be found at [15]. The day-ahead versions of $\{D_t^z\}_{t \in \mathcal{T}}$, $\{\phi_t^{s,z}\}_{t \in \mathcal{T}}$ and $\{\phi_t^{w,z}\}_{t \in \mathcal{T}}$ are also available at [15].

IV. MATHEMATICAL OPTIMISATION MODEL

The core idea of the model is to fit the GSKs, PAs and grid characteristics to the historical *PTDF* and F^{ref} data described in III-A, via equations (6) and (10). Direct fitting to the RAMs is not possible, but note that once F^{ref} is recovered, the RAMs can be obtained as per (8). When FRM and FAV are unknown, a simple solution is to set them to the mean historical value or zero. We only take $FAV = 0$ and compute the FRMs based on equation (8), since we observe r or R , F^{ref} and F^{max} .

However, due to the sparsity of the data set $\mathcal{D}_{\mathcal{LST}}$, even for fixed (b, A), multiple solutions exist for the GSKs and PAs. To simplify the analysis and increase data density, we assume that the GSK and base case are scenario independent, which is in agreement with literature proposals of defining the GSKs [23], [9], [26]. This means that at every time instant, we only consider a single GSK, instead of one GSK per scenario at every time. Thus we have that $G_{s,t}^1 = G_t^1$ and $g_{s,t}^0 = g_t^0$, which immediately implies that the s dependence of the $\{PTDF_{s,t}\}_{(s,t) \in \mathcal{ST}}$ and $\{F_{s,t}^{ref}\}_{(s,t) \in \mathcal{ST}}$ data comes exclusively from A^s and b^s . This reduces the number of

time dependent variables and helps preventing over-fitting. However, the corresponding optimisation problem remains under-determined and thus we introduce a regularisation for g_t^1 in IV-A. We then analyse the expressivity of $P_s(b, A)$ map, which is crucial to reconstruct the PTDFs in IV-B, and we present the main optimisation model in IV-C.

Before we proceed, we mention that the GSKs are only defined on the production nodes, since the TSOs have virtually no control over the demand. Thus, we define $g_{z,t}^1 \in \mathbb{R}^{N_p^z}$, $\bar{g}_{z,t}^1 \in \mathbb{R}^{N_p}$ and $G_t^1 \in \mathbb{R}^{N_p \times Z}$ only over the production nodes. We write the equivalent of equation (5) for production nodes only as

$$\phi_t^z = g_{z,t}^0 + \Delta \tilde{P}_{z,t} g_{z,t}^1, \quad (11)$$

where $g_{z,t}^0 \in \mathbb{R}^{N_p^z}$ $\Delta \tilde{P}_z$ are also re-defined for production nodes only. We will frequently use equation (11), instead of (5), but this does not reduce the amount of available information, since no nodal demand data is available. Note that $\Delta \tilde{P}_z$ denotes the total power deviation from the base case production levels.

A. A regularisation for the GSKs

To obtain our desired GSKs regularisations, we use the JPW [24] variant of the Sensitivity Coefficients Regression (SCR) model, which was proposed to define the GSKs in Austria [24] and Germany [23]. One could imagine using the fundamental approach for finding the GSKs based on the strategies defined in [9]. However, it is worth noting that this is much more complicated to do from the market participant's perspective, than from the TSOs' perspective. This is because the strategies for defining the GSKs require selecting a subset of power plants, that are then assigned non-zero values based on some rule, that varies between TSOs. The problem is that market participants do not know what subset the operators choose, and thus this would require combinatorial optimisation, which is very computationally intense. For this reason, we use the model proposed in [24]. For each zone z and time τ , the model solves the linear regression problem denoted as $\mathcal{R}(z, p, \tau)$,

$$\phi^p = \theta + \mathbf{n}\eta + \mathbf{l}\lambda + \mathbf{a}\iota + \mathbf{R}\rho + \mathbf{D}\delta + \mathbf{X}\xi + \epsilon, \quad (12)$$

for each plant p in zone z , where ϵ is the error term, $\phi^p \in \mathbb{R}^{T_p^{(\tau)}}$ is the vector of production levels for plant p , with $T_p^{(\tau)} = |\mathcal{T}_p^{(\tau)}|$ and $\mathcal{T}_p^{(\tau)} := \{t : (\phi_t^z)_p \neq 0, t \in [\tau - w_z, \tau]\}$ being the index set of all data points for which plant p has nonzero production and with times shifted backwards at most w_z time units with respect to the time of interest, τ .

The net position vector $\mathbf{n} \in \mathbb{R}^{T_p}$ can be computed based on the import ($I_{z,t}$) and export levels ($E_{z,t}$) for every zone $n_{z,t} + E_{z,t} - I_{z,t}$, while $\mathbf{a}_t := S_t^z$ and $\mathbf{l}_t := D_t^z$ define the vector of zonal energy prices and demands respectively. Wind and solar production levels are considered as exogenous variables and each is assigned a column of $\mathbf{R}^{T_p \times 2}$. The dummy variable matrix $\mathbf{D} \in \mathbb{R}^{T_p \times 3}$ differentiates between peak/offpeak and summer/winter conditions, while $\mathbf{X} \in \mathbb{R}^{T_p \times 3}$ is the interaction matrix of time dummy variables and net position change, such that $\mathbf{X} = \text{diag}(\mathbf{n})\mathbf{D}$. Finally $\theta, \eta, \lambda, \iota \in \mathbb{R}$ and $\rho \in \mathbb{R}^2, \delta \in \mathbb{R}^3, \xi \in \mathbb{R}^3$ are the regression coefficients to be determined.

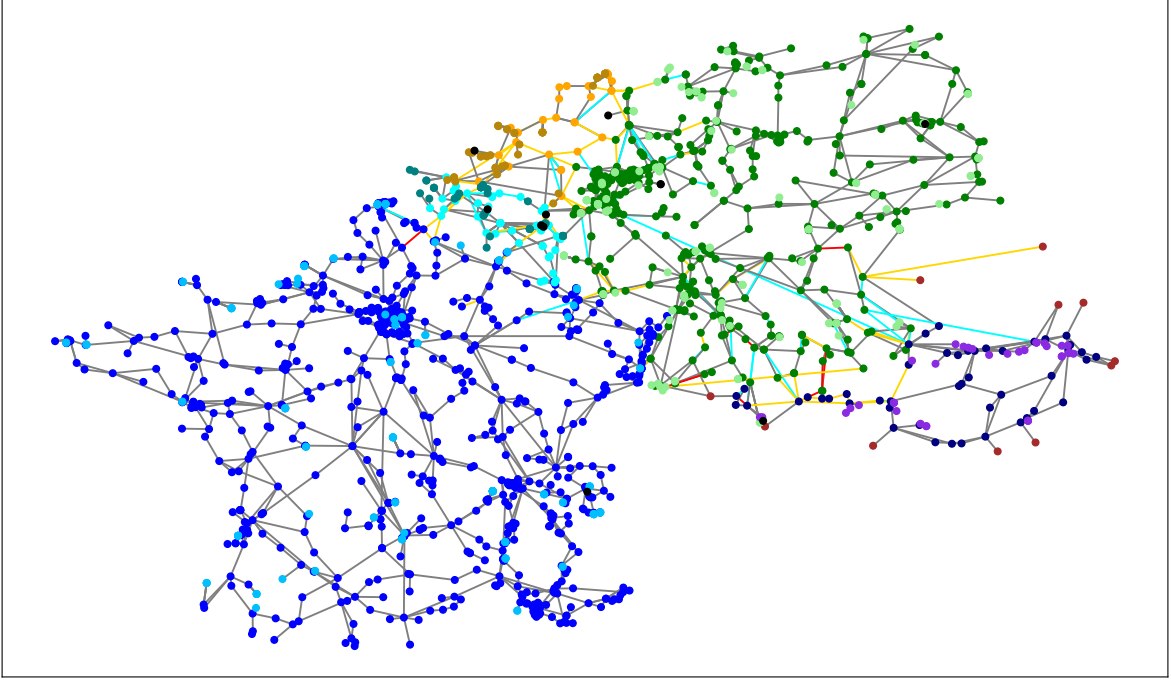


Fig. 2. CWE Grid Model. Each country is assigned two colours, with the more frequent one representing consumption nodes, while the other represents power plants. Critical lines and outages (contingencies) are marked in red and cyan respectively. Lines that are both CNEs and Contingencies (for other CNEs) are marked in yellow.

Since for each zone, the GSKs can be expressed as

$$(g_{z,t}^1)_p = \frac{(\Delta\phi_t)_p}{\Delta n_{z,t}}; \quad \forall p \in \mathcal{GU}_z, \quad (13)$$

where $\Delta\phi_t^z = \phi_t^z - g_{z,t}^0$ and $\Delta P_{z,t} = \sum_{p \in \mathcal{GU}_z} (\Delta\phi_t^z)_p$, the SCR model sets the GSK entries to be proportional to the infinitesimal version of (13) and thus [23]

$$(g_{z,\tau}^1)_p \sim \frac{\partial \phi_{p,\tau}}{\partial n_{z,\tau}} = \eta + \mathbf{D}_{\tau,:} \boldsymbol{\xi} =: v_{z,p,\tau}. \quad (14)$$

The SCR then uses the fact that $\mathbf{1}^T g_{z,\tau}^1 = 1$ to obtain the GSK values by normalisation as

$$(g_{z,\tau}^1)_p = \frac{v_{z,p,\tau}}{\sum_{p \in \mathcal{GU}_z} v_{z,p,\tau}}. \quad (15)$$

Since data points for which $(\phi_t^z)_p = 0$ are excluded, the number of data points for some regression problems $\mathcal{R}(z, p, \tau)$ might not be sufficient. In this case, or if $v_{z,p,\tau}$ is below a (small) threshold value, the corresponding entry in $v_{z,p,\tau}$ (and thus in $g_{z,\tau}^1$) is set to zero [23].

One can now apply the procedure for all $\tau \in \mathcal{T}$ to obtain the zonal GSKs which we denote as $G_z^1(\mathcal{T})$ and we denote this procedure as $SCR(z, \mathcal{T})$. It is suggested in [24] that better results are obtained when using the day-ahead (D-1) version of the regression coefficients. For us, this poses a problem for the left hand side of (12), since the D-1 version of ϕ_t^z is $g_{z,t}^0$, which is unknown. However, we propose to achieve this by initialising $g_{z,t}^0 \leftarrow \phi_t^z$, obtaining $g_{z,t}^1$ via $SCR(z, \mathcal{T})$, updating $g_{z,t}^0$ via (11) and then iterating again until we reach converge, which we measure by the norm of the difference of $g_{z,t}^1$ between two consecutive iterations. This

gives us the GSK regularisation that we denote with a hat from here on, i.e. as $\hat{g}_{z,t}^1$, \hat{g}_t^1 and \hat{G}_t^1 respectively. By using

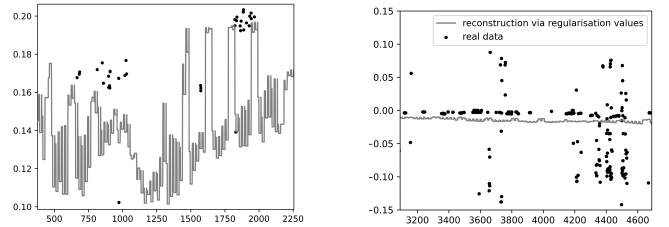


Fig. 3. PTDF values on an Austrian line (left) and a German-Dutch crossing line (right) showing the reconstruction of the PTDF based on the regularisation values of b_0 and $\{\hat{g}_t^1\}_{t \in \mathcal{T}}$.

only the regularisation values b_0 and $\{\hat{g}_t^1\}_{t \in \mathcal{T}}$, we can already obtain a PTDF reconstruction via (6). As seen in Figure 3, this generally gives reasonable qualitative reconstruction suggesting that we obtained appropriate regularisations, since the PTDF data was not used up to this point. Nevertheless, the overall reconstruction error is rather large, about 4.5 larger than the mean, and the regularisations' reconstruction performs poorly on many lines, particularly due to the inability to capture large variations, or be centred appropriately, as shown in the right part of Figure 3. Further, no reasonable reconstruction of the RAMs is available at this point. This shows the need for our main optimisation model, introduced in IV-C.

B. Expressivity of the Nodal PTDF map

When running the main optimisation problem with Ω_N^T , on many lines, we observe poor ability to reconstruct (expres-

sivity) the PTDFs' large variation, regardless of the chosen GSK. One reason for poor expressivity is that some power plants cannot be differentiated by the grid structure. This is because power plant location was obtained mostly by hand and is thus prone to errors. Further, by artificially connecting the power plants to the closest consumption node, multiple power plants inject power to the same single node, and thus are indistinguishable. To compensate for these errors, one could find the optimal location of each power plant, amongst its k closest neighbours. However, this approach introduces combinatorial complexity, and thus we only consider a convex relaxation of this problem. To this end we replace Ω_N^T with the map $M(x) \in \mathbb{R}^{N \times N_p}$ defined as

$$M(x)_{:,p} = \sum_{p \in \mathcal{GU}} \Omega^p x_p, \quad (16)$$

with $\Omega^p = [e_{v_1}, e_{v_2}, \dots, e_{v_k}] \in \mathbb{R}^{N \times k}$, $\forall v_i \in \mathcal{N}_k(p)$, where e_q is the indicator vector for position q , $\mathcal{N}_k(p)$ is the set of k -nearest neighbours of plant p (including itself), and $x := \|\cdot\|_{p \in \mathcal{GU}}^V x_p$. Thus, x^p tells us how the power for plant p is distributed across neighbouring nodes. By optimising over x , and enforcing that $0 \leq x_p \leq 1$ and $1^T x = 1$, we allow finding the most likely power injection pattern of each plant to the k closest neighbours, instead of pre-defining a single fixed connection to the closest node, and by choosing a sufficiently small $k > 1$ we can avoid over-fitting. Note that each entry of vector x_p corresponds to a node and can be seen as the probability that power plant p is located at that node.

C. The optimisation model

We now introduce the main optimisation model, explaining every term individually. The optimisation variables are b, x , the time series $\{g_t^1\}_{t \in \mathcal{T}}$, $\{\alpha_t^L\}_{t \in \mathcal{T}}$, and the constraints direction, which we model as binary variables $\nu_{l,s,t} \in \{-1, 1\}$ for every $(l, s, t) \in \mathcal{LST}$. We introduce optimisation terms to (i) fit to the PTDF data, (ii) fit to the F^{ref} data, (iii) regularise b , (iv) regularise g_t^1 , and the corresponding physical constraints. We further regularise α_t^L implicitly by taking this to be constant over non-overlapping time windows of length w_α , and $M(x)$ by choosing a small value of k . Note that we do not optimise over A since this is NP hard, and because the number of variables would become very large, and thus we write $P_s(b) := P_s(b, A)$. This is reasonable since we expect to have a good initial estimate of the graph topology, but it is much less likely that we have accurate physical characteristics of this (such as b).

We start by describing the PTDF term. Our choice of reference node might be mismatched by the TSO, and as a result we can recover $PTDF_{s,t}$ only up to a rank one matrix. This can be remedied by introducing the following mean centering: for each line j , let $\mathcal{ST}_j = \{(s, t) : (j, s, t) \in \mathcal{LST}\}$, $\mathcal{L}_{s,t} := \{j : (j, s, t) \in \mathcal{LST}\}$ and let $Y \in \mathbb{R}^{L_{cne} \times Z}$ with

$$Y_{j,:} := \frac{1}{|\mathcal{ST}_j|} \sum_{(s,t) \in \mathcal{ST}_j} (PTDF_{s,t})_{j,:}, \quad (17)$$

which is well defined since $\mathcal{ST}_j \neq \emptyset$ for all $j \in \mathcal{L}_{cne}$, and $PTDF_{s,t} \in \mathbb{R}^{|\mathcal{L}_{s,t}| \times Z}$ is the concatenation of all the available PTDFs at (s, t) . Likewise,

$$\tilde{g}_j(b, s, t) := \frac{1}{N-1} (P_s(b))_{j,:} \mathbf{1}_N, \quad (18)$$

where $\mathbf{1}_N \in \mathbb{R}^N$. The centred PTDF data is then

$$PTDF_{s,t}^c = \nu_{s,t} \odot PTDF_{s,t} - \Omega_{s,t} Y, \quad (19)$$

where \odot denotes the element-wise product, $\Omega_{s,t} \in \mathbb{R}^{|\mathcal{L}_{s,t}| \times L_{cne}}$ is a projection map from the critical lines to the one observed for (s, t) . Similarly, the centred version of the nodal PTDF becomes

$$P_s^c(b) := P_s(b) - \tilde{g}(b, s, t) \mathbf{1}_N^T. \quad (20)$$

We further regularise g_t^1 by imposing constant value over a (small) time window w_g via $\theta_t^g = \lfloor t/w_g \rfloor$. The PTDF optimisation term is then

$$f_P(b, g^1(\mathcal{T}), x) = \sum_{s,t \in \mathcal{ST}} \|PTDF_{s,t}^c - \hat{\Omega}_{s,t} P_s^c(b) M(x) G_{\theta_t^g}^1\|_2^2, \quad (21)$$

where $g^1(\mathcal{T}) = \{g_{\theta_t^g}^1\}_{t \in \mathcal{T}}$, $\hat{\Omega}_{s,t} = \Omega_{s,t} \Omega_{cne}$. Note that $PTDF_{s,t}^c$ has to be computed just once, while $P_s^c(b)$ must be updated at every iteration for b . The PTDF reconstruction for the CNEs can be then obtained for every (s, t) as

$$\overline{PTDF}_{s,t} = P_s^c(b) M(x) G_{\theta_t^g}^1 + Y. \quad (22)$$

We can obtain the F^{ref} optimisation term in a similar way. However, note that since α^L are centred around zero, $PSDF_s(b)$ centring is not necessary. Further, since the number of lines is very large, we only allow a small fraction of lines to have nonzero entries in α^L . To achieve this we redefine $\alpha_L \in \mathbb{R}^{L_\alpha}$ to be the vector of PAs, and $\Omega_\alpha \in \mathbb{R}^{L \times L_\alpha}$ to be the projection map from the set of nonzero PAs to all lines. We further regularise α^L by imposing constant value over the time window w_α via $\theta_t^\alpha = \lfloor t/w_\alpha \rfloor$, to obtain

$$f_F(b, \{\alpha_{\theta_t^\alpha}^L\}_{t \in \mathcal{T}}) := \sum_{s,t \in \mathcal{ST}} \|F_{s,t}^{ref,c} - \hat{\Omega}_{s,t} PSDF_s(b) \Omega_\alpha \alpha_{\theta_t^\alpha}^L\|_2^2, \quad (23)$$

where $F_{s,t}^{ref,c} = \nu_{s,t} \odot F_{s,t}^{ref} - \Omega_{s,t} \mu(F^{ref})$, with $\mu(F^{ref})$ being the mean reference flow values over the critical lines. By comparing (23) to (10), one can observe that we neglected the term given by g_t^0 , which can be written as a function of only g_t^1 via (11). In numerical experiments, we found that this term tends to only introduce noise, without improving the quality of fit or generalisation error. We note that this does not mean that the base case g_t^0 does not have an impact for the calculation of the reference flow, but rather that mean centring is the only approach that allowed us to capture a causal relationship. This may be because changes in the PAs may account for most of the flow variability, when compared to changes in the base case, making the task of capturing causal relationship between g_t^0 and F^{ref} very challenging. In other words, once the mean production vector is known, it is much more important to know how the TSOs use the grid controls (the PAs), than knowing the exact values of the production vector. Since the PAs can be controlled by the TSOs, and g_t^0 are not, this is a

desirable property for safe operation of the grid, as it means that the operators have more control over the critical flows than otherwise. We also enforce that $-\pi/6 \leq \alpha_{\theta_t^\alpha}^L \leq \pi/6$ as per [25], and that there is no pre-congestion, that is, the reference flows do not exceed the maximum allowed flows on the critical lines

$$\overline{lb} \leq P_S^{\text{cne}}(b)\alpha_{\theta_t^\alpha}^L \leq \overline{ub}, \quad (24)$$

for every θ_t^α , where $P_S^{\text{cne}} = \|\nu_{s \in \mathcal{S}} \Omega_{\text{cne}} P_s(b)\|$, the flow bounds are $\overline{ub} = \|\nu_{s \in \mathcal{S}}^V (F^{\text{max}} - \mu(F^{\text{ref}}))\|$, and $\overline{lb} = \|\nu_{s \in \mathcal{S}}^V (-F^{\text{max}} - \mu(F^{\text{ref}}))\|$, with F^{max} being the vector of maximum recorded $F_{s,t}^{\text{max}}$ values for each line. Assuming no pre-congestion is reasonable, because case this is very undesirable, and the TSOs choose remedial actions such that pre-congestion is minimised or removed [5].

The regularisation terms can simply be written as

$$f_G(\{g_{\theta_t^g}^1\}_{t \in \mathcal{T}}) = \sum_{t \in \mathcal{T}} \|\hat{g}_t^1 - g_{\theta_t^g}^1\|_2^2, \quad f_b(b) = \|b_0 - b\|_2^2. \quad (25)$$

By collecting all the terms, the final optimisation problem is:

$$\begin{aligned} \min_{\omega} & f_P(b, \{g_{\theta_t^g}^1\}_{t \in \mathcal{T}}, x) + \lambda_F f_F(b, \{\alpha_{\theta_t^\alpha}^L\}_{t \in \mathcal{T}}) + \\ & \lambda_G f_G(\{g_{\theta_t^g}^1\}_{t \in \mathcal{T}}) + \lambda_b f_b(b), \\ \text{s.t. } & b \geq \delta_b, \text{ and } \nu_{s,t} \in \{-1, 1\}^{|\mathcal{L}_{s,t}|} \\ & 1^T g_{z, \theta_t^g}^1 = 1, \text{ and } -\delta_g^z \leq g_{z, \theta_t^g}^1 \leq \delta_g^z, \forall z \in \mathcal{Z}, t \in \mathcal{T}, \\ & 1^T x^p = 1, \text{ and } 0 \leq x^p \leq 1, \forall p \in \mathcal{GU}, \\ & \overline{lb} \leq P_S^{\text{cne}}(b)\alpha_{\theta_t^\alpha}^L \leq \overline{ub}, \text{ and } -\frac{\pi}{6} \leq \alpha_{\theta_t^\alpha}^L \leq \frac{\pi}{6}, \forall t \in \mathcal{T}, \end{aligned}$$

where $\delta_b > 0$ is a small number, $0 < \delta_g^z \leq 1$ is the maximum allowed GSK value per entry, $\lambda, \in \mathbb{R}_+$ are regularisation parameters, and the full variable set is denoted as $\omega = (b, \{g_{\theta_t^g}^1, \alpha_{\theta_t^\alpha}^L\}_{t \in \mathcal{T}}, x, \nu_{s,t \in \mathcal{ST}})$. Note that b is fixed in principle, and once found it can be directly used for future predictions, but $g_{\theta_t^g}^1$ and $\alpha_{\theta_t^\alpha}^L$ being time dependent can not. However, these can be used for statistical analysis or forecasting.

V. NUMERICAL OPTIMISATION

The optimisation problem is highly nonlinear and non-convex over the variable set ω , and very large, with about 4 million variables, but the vast majority of variables are the $g_{\theta_t^g}^1$, which are decoupled in time. Thus, we propose an alternating minimisation approach, by solving for one variable at a time, while keeping the others fixed, and using the most recent updates. This has the following steps: (i) solve for each $\nu_{s,t}$ independently, (ii) solve for b via stochastic optimisation, (iii) solve for x via stochastic optimisation, (iv) solve for $\{g_{\theta_t^g}^1\}_{t \in \mathcal{T}}$, for each time independently, via Convex Quadratic Programming (CQP), (v) solve for $\{\alpha_{\theta_t^\alpha}^L\}_{t \in \mathcal{T}}$, for each time independently, via CQP, and (vi) iterate until convergence or time limit is reached. This approach does not guarantee the algorithm reaches a local minimum, since the resulting problems for b and $\nu_{s,t}$ are not marginally convex. However, it makes the problem computationally feasible and leverages the marginally convex structure of the objective function in g_t^1 and $\alpha_{\theta_t^\alpha}^L$ as well as time decoupling. Note that we are guaranteed to reduce the objective function value at every optimising step

for $\nu_{s,t}$, $g_{\theta_t^g}^1$, and $\alpha_{\theta_t^\alpha}^L$, but not for b and x , and thus, great care needs to be taken when optimising over these. We next describe in detail optimising over each set of variables.

A. Optimising over $\nu_{s,t}$

We optimise over each entry $\nu_{l,s,t} \in \{-1, 1\}$ such that $(l, s, t) \in \mathcal{LST}$ independently, by choosing the value that decreases the objective function value most. With all the other quantities pre-computed, this is fast since we only have to compare two cases for each variable $\nu_{s,t}$. Further, since the data points are independent, this approach gives the global optimal solution when $b, \{g_{\theta_t^g}^1, \alpha_{\theta_t^\alpha}^L\}_{t \in \mathcal{T}}$ and x are fixed.

B. Optimising over b

Due to the large number of scenarios $|\mathcal{S}|$, it is extremely expensive to compute all the $P_s(b)$ values, since these require the computation of $(\tilde{A}_s^T \tilde{B}_s \tilde{A}_s)^{-1}$ for every s at every iteration. Further, the number of times $|\mathcal{T}|$ is also very large, making the computation of $f_P(b)$, $f_F(b)$, $\nabla_b f_F(b)$ and $\nabla_b f_P(b)$ extremely time and memory intensive, and thus we use tensor implementation in *PyTorch*. Computing the Hessian is impractical, and thus we resort to first order methods. We use stochastic optimisation, at each iteration k , by randomly choosing only a subset \mathcal{T}_k of \mathcal{T} (but all corresponding s), because this: (i) dramatically reduces the energy usage per iteration, due to much fewer inversions and $P_s(b)M(x)G_{\theta_t^g}^1$ products, (ii) it reduces the computational time, and (iii) we observe much larger objective function decrease per iteration when compared to the deterministic optimisation step, since the landscape of $f_P(b)$ is noisy, and stochastic methods act as a smoother. To allow step size adjustment, we use ADAM [17] as optimiser, and solve for K_b iterations in every cycle.

C. Optimising over x

Although the problem for x is convex, and we could solve via CQP, the double sum over s and t makes the problem excessively large, yielding infeasible memory requirements when using *Mosek*. Thus, we perform stochastic optimisation using the ADAM optimiser, similarly to the optimisation over b . This makes the computations less memory and energy intensive, but does not guarantee convergence to the optimal solution $x^*(b, G^1)$, even when all the other variables are fixed, since ADAM may not converge even on convex problems [21]. However, we find this to be a good trade-off between computational feasibility and making optimisation progress.

D. Optimising for $g_{\theta_t^g}^1$

We solve for each $g_{\theta_t^g}^1$ independently since the problem is decoupled in time. The resulting optimisation problem reads

$$\begin{aligned} \min_{g_{\theta_t^g}^1} & \sum_{\mathcal{F}_t^g} \|PTDF_{s,t}^c - \hat{\Omega}_{s,t} P_s^c(b) M(x) G_{\theta_t^g}^1\|_2^2 + \\ & \lambda_G \|\hat{g}_t^1 - g_{\theta_t^g}^1\|_2^2 \\ \text{s.t. } & 1^T G_{\theta_t^g}^1 = 1, \forall t \in \mathcal{T}, \quad 0 \leq (G_{\theta_t^g}^1)_{ij} \leq 1, \forall t \in \mathcal{T}, \end{aligned}$$

where $\mathcal{F}_\tau^g := \{(s, t) : s \in \mathcal{S}, \theta_t^g = \tau\}$. The problem is convex, and we solve it using *Mosek*'s API in *Python*, but other convex optimisation solvers could be used. Note that the time spent during this step increases linearly with the size of \mathcal{T} . While this is the best one can expect, this step is performed multiple times and thus it can drastically increase the computational time.

E. Optimising over $\alpha_{\theta_t^\alpha}^L$

Similarly to V-D, the problem is convex, and decoupled in time for every θ_t^α . This reads

$$\begin{aligned} \min_{\alpha_{\theta_t^\alpha}^L} \quad & \sum_{s \in \mathcal{S}, \{t: \theta_t^\alpha = \tau\}} \|F_{s,t}^{ref,c} - PSDF_s(b)\Omega_\alpha \alpha_{\theta_t^\alpha}^L\|_2^2 \\ \text{s.t.} \quad & \bar{lb} \leq P_S^{\text{cne}}(b)\alpha_{\theta_t^\alpha}^L \leq \bar{ub}, \quad -\frac{\pi}{6} \leq \alpha_{\theta_t^\alpha}^L \leq \frac{\pi}{6}, \quad \forall t \in \mathcal{T} \end{aligned}$$

and we solve it using *Mosek*'s API in *Python*.

F. Obtaining new values for $g_{\theta_t^g}^1$ and $\alpha_{\theta_t^\alpha}^L$ after the model fit

Once the solution ω of the main optimisation model is available, one can obtain new values for $g_{\theta_t^g}^1$ and $\alpha_{\theta_t^\alpha}^L$ for $t \notin \mathcal{T}$ ($t > \max_{\tau \in \mathcal{T}} \tau$) by pre-fitting $\nu_{s,t}$ and then solving the corresponding convex optimisation problems V-D and V-E once. The number of variables is reduced dramatically and once the $P_s(b)$ maps were pre-computed, this problem can be solved on a usual personal computer in at most few of minutes.

VI. NUMERICAL RESULTS

We evaluate the model for: (i) in sample quality of fit, (ii) generalisation for $(l, s, t) \notin \mathcal{LST}$, with $l \in \mathcal{L}_{\text{cne}}$, $s \in \mathcal{S}$, $t \in \mathcal{T}$, i.e. for unobserved (l, s, t) containing observed individual elements, and (iii) generalisation for unobserved (l, s, t) with new times $t \notin \mathcal{T}$ and observed $l \in \mathcal{L}_{\text{cne}}$ and $s \in \mathcal{S}$. The first evaluation requires in-sample tests while for the other two we perform out of sample tests. Due to the data sparsity, it is impossible to split the data in a train and test set to perform out-of-sample tests for (ii) and (iii) simultaneously. Thus, we split the data twice, fit the model to each training set, and perform tests on the corresponding test set.

For the first train-test split, we cannot simply remove (say) 20% of the data at random and add it to the test set, because this leaves many (l, s) combinations and t indices with no data. To resolve this, we propose the following heuristic: (i) select an available (l, s, t) point at random, (ii) if the number of points on (l, s) and at time t are above the thresholds m_{LS} and m_T respectively, remove the point from the train set, and add it to the test set, (iii) otherwise do not remove the point but mark it such that it cannot be chosen again, (iv) repeat until the desired test set size is achieved or no more points can be removed. We choose $m_T = 8$, since this is the average data density with respect to time, and allowing a smaller value would make data rich time indices become data scarce. We choose $m_{LS} = 50$, but generally it is m_T that has a strong impact on maximum fraction of data that can be removed. In this case, the maximum fraction of removed points is about 18%, but this corresponds to is highly inter-dependent choices.

Thus, we choose to assign 85% and 15% to the train ($\mathcal{D}_1^{\text{train}}$), and test ($\mathcal{D}_1^{\text{test}}$) set respectively. We fit the model to the train set and compute the out-of-sample error on the test set.

The second train and test sets are obtained by splitting the time indices at random, such that 20% of the times are assigned to $\mathcal{D}_2^{\text{test}-t}$ and the rest to $\mathcal{D}_2^{\text{train}}$. Further, $\mathcal{D}_2^{\text{test}-t}$ is randomly split in $\mathcal{D}_2^{\text{test}-fit}$ and $\mathcal{D}_2^{\text{test}-lst}$ which receive 75% and 25% respectively. We fit the model only by using $\mathcal{D}_2^{\text{train}}$. However, to test the model we need GSKs, PAs and $\nu_{s,t}$ for each time interval, but these are unavailable. We obtain these by using fixed b , x (fit based on $\mathcal{D}_2^{\text{train}}$), and fit for $g_{\theta_t^g}^1$, $\alpha_{\theta_t^\alpha}^L$ and $\nu_{s,t}$ only once on $\mathcal{D}_2^{\text{test}-fit}$. Out of sample tests are performed on $\mathcal{D}_2^{\text{test}-lst}$. This is an indication of generalisation quality to new times, given the time-independent variables b and x , and thus the corresponding test evaluates if they give a good representation of the real grid.

We run the algorithm on a 72 core machine with 768 GB of memory. We find choosing $k = 5$, $w_\alpha = 8$ and $w_g = 6$ to be a good trade-off between quality of fit and generalisation ability. We stop after running the algorithm for 14 complete cycles, because the objective function value decrease per cycle becomes rather small, and the computational cost is high.

A. PTDFs and RAMs reconstruction

To evaluate the quality of our reconstruction, we use multiple error measures. First, we use the average absolute mean error which we denote as $\delta_{\text{abs}}(\cdot)$. To estimate the quality of fit, we consider $\delta_\sigma(\cdot) := \delta_{\text{abs}}(\cdot)/\sigma(\cdot)$, where $\sigma(\cdot)$ is the standard deviation of the data set of interest. We use this ratio because if deviations are large, then it is expected that the ability to fit reduces, and thus we should penalise the model less. Further, we are interested in the performance of the model, and since to our knowledge, our model is the first to reconstruct the PTDFs and RAMs, we can only compare to a benchmark model. The benchmark model assumes that all data variation is noise, and thus it reconstructs the data by simply taking the empirical mean value for every line, and assigns this to every single (s, t) combination with $s \in \mathcal{S}$ and $t \in \mathcal{T}$. We define the benchmark model error as $\delta_{\text{null}}(\cdot)$ and the corresponding relative error with respect to the benchmark model as $\delta_{r\text{Null}}(\cdot) = \delta_{\text{abs}}(\cdot)/\delta_{\text{null}}(\cdot)$. Finally, if one is to use these constraints in the optimisation problem for price finding [20], the relative error with respect to the mean value is more relevant, as this measures the fractional error in our constraints. We denote this as $\delta_\mu(\cdot) := \delta_{\text{abs}}/\mu(\cdot)$, where $\mu(\cdot)$ takes the mean of the corresponding data set. However, since the PTDFs are mean zero centred we define $\mu(\cdot)$ to be the mean of absolute values in this case.

We show the overall errors obtained across the whole data set in Table I. Considering that only about 2.3% of the existing data is available, and thus used as input in our model, we obtain good reconstruction quality. Since δ_μ is rather small, the reconstructed PTDFs and RAMs can be used to reconstruct the feasibility domain at each time step or for forecasting with reasonable error. Further, we obtain good generalisation for both \mathcal{D}_1 and \mathcal{D}_2 tests, suggesting that our model is performing well for constraints reconstruction. The slightly

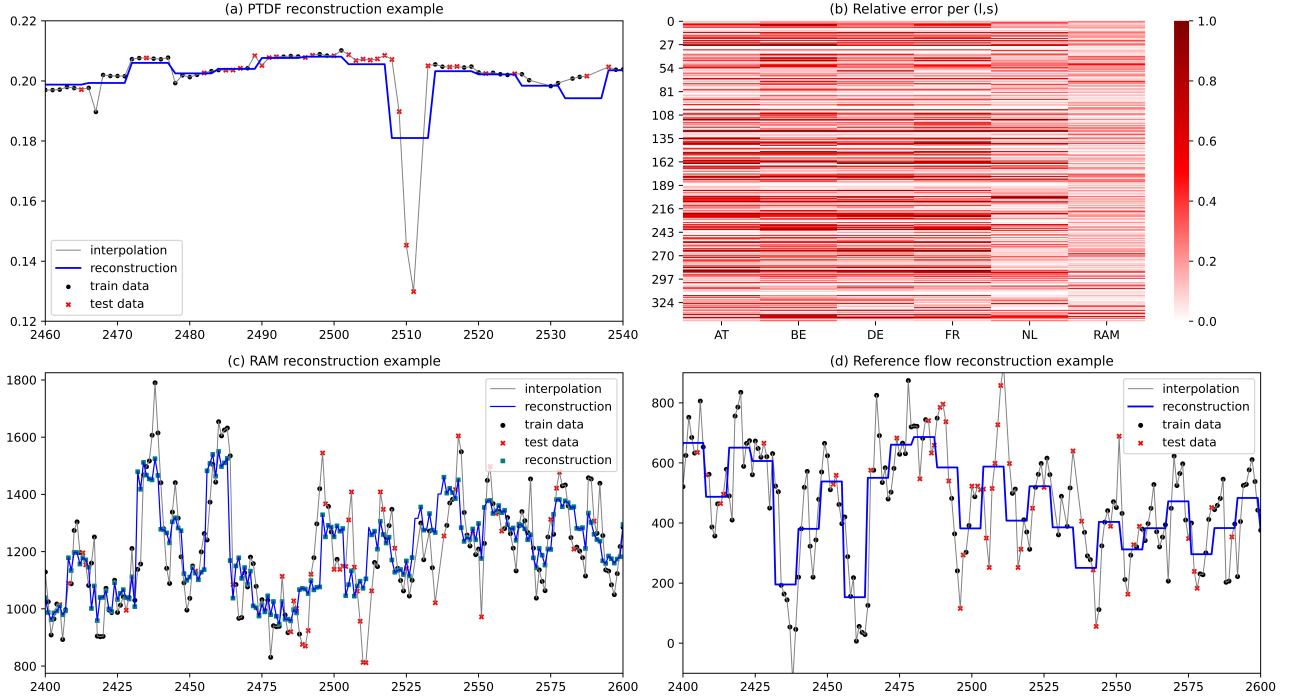


Fig. 4. Detailed reconstruction results. Plots (a), (c) and (d) show the PTDF (French entry), RAM, and F^{ref} reconstruction results for a representative (French-German) line. The blue lines (and dots) indicate the reconstruction, while the black and red dots indicate the training and test data points respectively. The grey line is an interpolation between available data points for better visualisation only. The heat map in (b) indicates the relative RAM error in the sixth column, and relative PTDF error for the five countries, corresponding to the first five columns. These errors are computed with respect to the mean, since due to data scarcity on many (l, s) with very few observations, the standard deviation is very inaccurate or even zero.

better generalisation for \mathcal{D}_1 when compared to \mathcal{D}_2 could be because the grid changes, albeit slowly, over time. Grid

Error measure	δ_{abs}	δ_{σ}	δ_{rNull}	δ_{μ}
PTDF on \mathcal{D}_1^{train}	0.0122	14.4%	36.1%	22.7%
PTDF on \mathcal{D}_1^{test}	0.0119	14.0%	35.2%	22.1%
F^{ref} on \mathcal{D}_1^{train}	88.1	36.1%	39.6%	11.7%
F^{ref} on \mathcal{D}_1^{test}	96.6	39.5%	43.4%	12.8%
RAM on \mathcal{D}_1^{train}	93.1	42.6%	54.9%	16.0%
RAM on \mathcal{D}_1^{test}	94.2	43.1%	55.6%	16.2%
PTDF on \mathcal{D}_2^{train}	0.012	14.2%	35.5%	22.3%
PTDF on $\mathcal{D}_2^{test-lst}$	0.014	16.1%	40.5%	25.4%
F^{ref} on \mathcal{D}_2^{train}	87.5	35.8%	39.3%	11.6%
F^{ref} on $\mathcal{D}_2^{test-lst}$	130.4	53.4%	58.5%	17.3%
RAM on \mathcal{D}_2^{train}	92.7	42.4%	54.7%	15.9%
RAM on $\mathcal{D}_2^{test-lst}$	105.1	48.2%	62.1%	18.1%

TABLE I

TRAIN AND TEST ERRORS FOR THE TWO DATA SPLITS \mathcal{D}_1 AND \mathcal{D}_2 .

elements are added, removed or changed, and thus one might consider to periodically re-calibrate b to new data, starting with the old b in the optimisation process. By analysing δ_{σ} , one can see that we obtain a better quality of fit to the PTDFs when compared to the RAMs. This is not contradictory to what one might infer from Figure 4, where we perform normalisation with respect to the empirical mean (instead of σ) for every (l, s) combination. This is expected, because the PTDFs are simpler to compute, involving fewer variables and simpler physical equations, while the target RAM signal is computed by the TSOs based on their base case estimate, which is much

more complex [30], [31]. Finally, our model is clearly superior for reconstruction of the PTDFs and significantly better for the reconstruction of F^{ref} , when compared to the benchmark model, and thus is superior to the naive approach of treating data on every line as purely random. Therefore we capture meaningful structure in the data, and we further illustrate this by showing more detailed results in Figure 4.

We can see that the reconstruction fits the training signal well, and implies reasonable values for time indices where data is unavailable. By analysing Figure 4 (a), one can observe that we partially capture the sharp change in the PTDF although these points were not part of the training set. We also obtain good reconstruction for the F^{ref} and RAM, although again we cannot fully capture sharp variations. This is because of the introduced time window values chosen for w_{α} and w_g . We can confirm that much better fit is obtained by letting these to one, but in this case the out of sample generalisation is rather poor. Thus, we found $w_{\alpha} = 8$ and $w_g = 6$ to be a good trade-off. Note that by (8), even with fully available F^{ref} values, the RAM cannot be reconstructed without knowing the other terms. To obtain full RAM reconstruction, for each (l, s) pair, we thus always use the last available values at each time index. For this reason, we mark the points where exact F^{max} and FRM are used by blue dots in Figure 4 (c). Finally, Figure 4 (b) shows the δ_{μ} value, computed for each available line scenario combination, that is each $(l, s) \in \mathcal{LS}$. We can observe that our model generally performs well, although some error entries are rather large. This could be because of small amount of information on some lines, but also because of

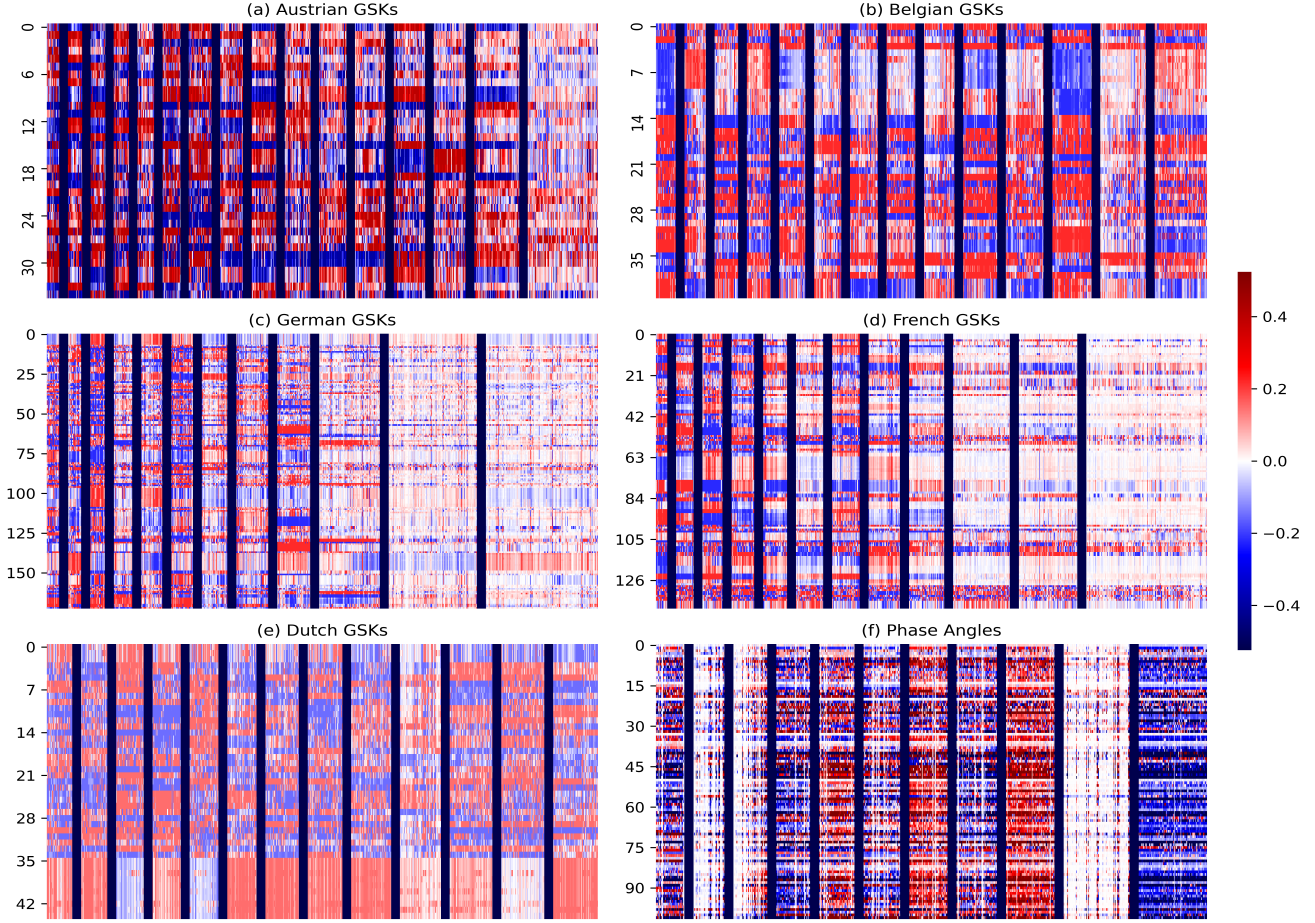


Fig. 5. Clustering results for the Generation Shift Keys and Phase Angles. For each subplot, the x axis corresponds to time indices while the y axis corresponds to entries in g_t^1 and α_t^L respectively, and the vertical dark bars are delimiters between the clusters.

imperfections in the grid structure, particularly regions where fine grid data is unavailable. While the optimisation model proposed in IV-C dramatically improves the quality of fit, due to the problem structure, multiple solutions exist, and finding the global minimum is NP hard. Further, we do not optimise over A , as this would require solving an integer optimisation problem over variables corresponding to entries in A - where each sub-problem is itself NP hard (the problem we are currently solving). This is clearly computationally infeasible. Nevertheless, our results generally give good reconstruction and generalisation, and outperform the naive approach, and we next show that we recover structure in the GSKs and PAs.

B. Structure Analysis of reconstructed GSKs and PAs

We perform clustering analysis to reveal groups of similar zonal GSKs and PAs respectively. We cluster at zonal level because the TSO in each zone chooses the GSKs independently of the others. We use the k-means algorithm [19], which in this case we found to give better results when compared to the Gaussian Mixture Models [10] solved via Expectation Maximisation. We run the algorithm 50 times and choose the result that gives the smallest within cluster sum of squares. We choose the number of clusters by using the elbow method,

but since there is no clear steep change in the derivative, we err on the parsimonious side. This is because a smaller number of clusters reduces the effect of noise and makes the recovered structure more meaningful. Because of the very high dimension, we show the results via a heat map, where each column corresponds to the values of one $g_{z,t}^1$ (or α_t^L), and we delimit the clusters with dark vertical bars. If within a cluster the entries of all $g_{z,t}^1$ (or α_t^L) are similar, then they are assigned similar colours on the heat map, and the cluster shows relatively uniformly across the x axis.

Figure 5 suggests that we recover meaningful clustering structure, with generally well balanced clusters. This suggests that the TSOs choose the GSKs in a rather simplistic and robust way, specifically from a small subset of pre-determined fixed vectors. The variations within the clusters could be explained by expertise-based value adjustment, although a part of this could be due to noisy input data and the limitations of the numerical optimisation process, which does not guarantee reaching the global optimum. Nevertheless, the recovered clusters are meaningful, and the strategies for choosing the GSKs discussed in [9] support our findings: while the exact strategies differ based on TSO, the operators choose a subset of power plants to have non-zero GSKs and then assign the values based on a simple rule. Since the power plants are

expected to switch on in order, based on the position in the generation stack, the number of considered possibilities is expected to be rather small. This, combined with the simplistic value assignment rules, results in the formation of clusters of GSKs, where within each cluster, the GSKs could be identical or very similar depending on the value assignment rule.

VII. CONCLUSION

Day Ahead Auctions in Power Markets that use the Flow Based Market Coupling methodology require linear constraints on all transmission lines to ensure that the auction leads to solutions that do not violate transmission load limits. In reality constraints corresponding to only a very small number of critical lines constraints are published. The missing data reduces market participants' risk management ability and ultimately leads to higher prices.

We propose a methodology that recovers the electricity grid constraints via a mathematical optimisation problem using only publicly available data. The model reconstructs the grid structure and the underlying time dependent signals known as the Generation Shift Keys and Phase Angles, which are then transformed to the Power Transmission Distribution Factors and Remaining Available Margins, via mappings that are determined by the grid structure. Our results show good reconstruction quality and out of sample generalisation, and we recover meaningful clustering structure for the zonal GSKs and PAs, yielding insight into the operator's behaviour and suggesting that the GSKs could be chosen from a (fixed) palette of values up to small adjustments for daily characteristics. Our results have the potential to improve the risk management by day-ahead auction participants, leading to increased transparency, reduced price inefficiency and higher social welfare.

ACKNOWLEDGMENT

We would like to thank CFM London Power and Gas desk within Macquarie Group for partially funding this research. We would like to thank the reviewers for the insightful comments that have helped us improve the paper.

REFERENCES

- [1] ACER, *Explanatory note on the day-ahead commoncapacity calculation methodology for Core CCR*, 2018, available at: www.acer.europa.eu.
- [2] ACER, *3rd ERI Progress Report*, 2015, available at: www.acer.europa.eu.
- [3] M. Aguado et al, *Flow-based market coupling in the Central Western European region - on the eve of implementation*, 2012, CIGRE .
- [4] Austrian Power Grid, *Austrian Grid Data*, 2019, available at: <https://www.apg.at/en/Stromnetz/APG-Netz>.
- [5] R. Besselink, D. Dudoignon, T. Ringelband, P. H. Schavemaker, J. Schwachheim and R. Sikora, *TSC's coordinated intraday capacity calculation concept*, 2015 12th International Conference on the European Energy Market (EEM), 2015, pp. 1-5, doi: 10.1109/EEM.2015.7216671
- [6] CWE Market Coupling, *JAO Publication Handbook*, 2017, available at <https://www.jao.eu/main>.
- [7] CWE Market Coupling, *JAO Utility Tool*, 2017, available at <http://utilitytool.jao.eu/Util>.
- [8] CWE TSOs, *Documentation of the CWE FB MC solution*, 2018, available at: <https://www.jao.eu/support/resourcecenter/overview>.
- [9] C. Dierstein, *Impact of Generation Shift Key determination on flow based market coupling*, 2017, 14th International Conference on the European Energy Market (EEM), 2017, pp. 1-7, doi: 10.1109/EEM.2017.7981901.
- [10] R.O. Duda and P.E. Hart, *Pattern classification and scene analysis*, Wiley, New York, 1973.

- [11] W.M. Edjroubi C.M. Atke, *SciGRID Open Source Transmission Network Model, User Guide v0.2*, 2015, available at: <https://www.power.scigrid.de/pages/downloads.html>.
- [12] Davis Edwards, *Energy Trading and Investing: Trading, Risk Management and Structuring Deals in the Energy Market*, 2010, New York : McGraw-Hill.
- [13] Elia, *Belgian Grid Data*, 2020, available at: <https://www.elia.be/en/grid-data>.
- [14] ENTSOE, *Grid Map*, available at: <https://www.entsoe.eu/data/map/>.
- [15] ENTSOE, *Transparency Data*, available at: <https://transparency.entsoe.eu/dashboard/show>.
- [16] B. Felten, P. Osinski , T. Felling, C. Weber, *The flow-based market coupling domain - Why we can't get it right*, Utilities Policy, 2020.
- [17] D. P. Kingma and J. Ba, *ADAM: A Method for Stochastic Optimization*, 2017.
- [18] L. Wyroll, K. Kollenda, C. Müller A. Schnettler, *Impact of Flow-Based Market Coupling Parameters on European Electricity Markets*, 2018, 53rd International Universities Power Engineering Conference (UPEC).
- [19] J. MacQueen, *Some methods for classification and analysis of multivariate observations*, Proc. Fifth Berkeley Symp. on Math. Statist. and Prob., Vol. 1 (Univ. of Calif. Press, 1967), 281–297.
- [20] NEMO Committee, *Euphemia Public Description Single Price Coupling Algorithm*, 2019.
- [21] S. J. Reddi, S. Kale, S. Kumar, *On the convergence of ADAM and beyond*, 2018.
- [22] SciGRID, *The Open Source Model of European Energy Networks*, available at: <https://www.scigrid.de/>.
- [23] D. Schönheit and R. Sikora, *A statistical approach to generation shift keys*, 2018 15th International Conference on the European Energy Market (EEM), doi: 10.1109/EEM.2018.8469900.
- [24] D. Schönheit, *An improved statistical approach to generation shift keys: Lessons learned from an analysis of the Austrian control zone* 2019 Z Energiewirtschaft 43, 193–212, doi: 10.1007/s12398-019-00261-w.
- [25] K. Van den Bergh, E. Delarue, W. D'Haeseleer, *DC power flow in unit commitment models*, 2014, Energy and Environment 240.
- [26] K. Van den Bergh and E. Delarue, *An improved method to calculate injection shift keys*, 2016, Electric Power Systems Research Volume 134, Pages 197-204, doi: <https://doi.org/10.1016/j.epsr.2016.01.020>.
- [27] K. Van den Bergh, Boury, E. Delarue, *The Flow-Based Market Coupling in Central Western Europe: Concepts and definitions*, 2016, Electricity Journal, 29(1), 24–29. <https://doi.org/10.1016/j.tej.2015.12.004>.
- [28] S. Voswinkel, B. Felten, T. Felling, C. Weber, *Flow-based market coupling: What drives welfare in Europe's electricity market design?*, 2019, online at: <http://hdl.handle.net/10419/201591>.
- [29] D. Schönheit, M. Kenis, L. Lorenz, D. Möst, E. Delarue, and K. Bruninx, *Toward understanding flow-based market coupling: An open-access model*, Advances in Applied Energy 2, 2021.
- [30] C. Byers and G. Hug, *Modeling flow-based market coupling: Base case, redispatch, and unit commitment matter*, 2020 17th International Conference on the European Energy Market (EEM), 2020, pp. 1-6, doi: 10.1109/EEM49802.2020.9221922.
- [31] D. Schönheit, R. Weinhold, C. Dierstein, *The impact of different strategies for generation shift keys (GSKs) on the flow-based market coupling domain: A model-based analysis of Central Western Europe*, Applied Energy, Vol 258, 2020, 114067, ISSN 0306-2619.



Ioan Alexandru Puiu holds a B.Eng. degree from the University of Manchester and an M.Sc. in Applied Mathematics from the University of Oxford. He is currently pursuing a D.Phil. degree in Applied Mathematics at the University of Oxford. His research interests include machine learning, data science, optimisation, game theory and energy markets.



Raphael Andreas Hauser is an Associate Professor at the Oxford Mathematical Institute, a Turing Fellow at the Alan Turing Institute in London, and Tanaka Fellow in Applied Mathematics at Pembroke College Oxford. He is a former recipient of the SIAM Activity Group on Optimization Prize. His research areas encompass numerical optimization, operations research, mathematical modelling and data science.