Orchestrating Distributed Energy Resources to Provide Power System Ancillary Services

From the Perspective of Virtual Power Plants

Orchestrating Distributed Energy Resources to Provide Power System Ancillary Services

From the Perspective of Virtual Power Plants

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of the Rector Magnificus Prof.dr.ir. H. Bijl,
chair of the Board for Doctorates,
to be defended publicly on
Wednesday, 29 April 2026 at 15:00.

by

Haiwei XIE

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

Rector Magnificus,	chairperson
Dr. J.L. Cremer,	Delft University of Technology, promotor
Prof.dr. P. Palensky,	Delft University of Technology, promotor
Independent Members:	
Prof.dr.ir. B. De Schutter,	Delft University of Technology
Prof.dr.ir. G. Deconinck,	KU Leuven, Belgium
Prof.dr. F. Dörfler,	ETH Zurich, Switzerland
Prof.dr. H. Pandžić,	University of Zagreb, Croatia
Prof.dr.ir. J.A. la Poutré,	Delft University of Technology, reserve member
Other Member:	
Dr. P. Mohajerin Esfahani,	Delft University of Technology University of Toronto, Canada

Keywords: Virtual power plants, power system ancillary services, distributed optimization

ISBN 978-94-6518-291-9

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

博观而约取，

厚积而薄发。

Contents

Summary

Rapid decarbonisation and large-scale deployment of inverter-based renewables are turning power systems into low-inertia, weakly meshed grids where stability can no longer be guaranteed by a few synchronous machines. At the same time, Distributed Energy Resources (DERs) such as storage, electric vehicles, rooftop PV, and flexible loads are proliferating in distribution networks, creating substantial but fragmented flexibility. Virtual Power Plants (VPPs) aggregate these DERs into controllable portfolios, yet ancillary services are still often treated as static reserve capacities or power set-points, neglecting the dynamic behaviour required for Fast Frequency Response (FFR), inertia emulation, and other short-timescale services. This thesis, therefore, adopts the concept of Dynamic Virtual Power Plants (DVPPs), which explicitly model internal device dynamics.

The thesis develops a coordination framework and associated operation models that enable aggregated DERs to provide reliable dynamic ancillary services under scalability constraints, limited disclosure of local information, and behavioural considerations. It is organised around three themes: modelling DVPP operation for dynamic services, designing distributed and scalable solution methods, and integrating prosumer motivation into ancillary-service provision.

We first study the coordination of several DER aggregators, each managing a portfolio of DERs, to collectively provide the FFR specified by the DVPP operator. FFR is represented as a transfer function from frequency deviations to active-power response, and an $\mathcal{H}_\infty$-based performance metric is used to quantify the dynamic mismatch between the aggregate response and the target FFR behaviour. This results in a convex optimisation problem with linear matrix inequality constraints that jointly capture dynamic regulation quality and economic efficiency. To avoid a fully centralised implementation, which would require full disclosure of local information, the problem is reformulated as a state-based potential game in which each aggregator acts as an agent with its own cost function and feasible set. A tailored game-theoretic learning scheme enables the agents to update their FFR provisions using only local information and communication with neighbouring agents, and to converge to a Nash equilibrium that closely approximates the optimal solution, while reducing information disclosure and avoiding single points of failure.

Second, the DVPP composed of a large number of DERs in a distribution network is studied to provide aggregate FFR to the transmission grid. The DVPP's operation is still framed as matching a desired FFR transfer function at the grid interface with the aggregate transfer function of all participating devices. This matching is enforced in the frequency domain by aligning real and imaginary parts of the two transfer functions at selected angular frequencies. This leads to a nonlinear, nonconvex optimisation problem that links device-level design variables to DVPP performance and operating cost. Latencies are treated as decision variables, allowing the DVPP to shape both the amplitude and timing of its response and to unlock additional temporal flexibility. To preserve computation tractability at large portfolio sizes, a dual-decomposition solution method is proposed to decompose the global problem into per-device subproblems, coupled only through low-dimensional dual variables. The proposed solution method provides provable guarantees on solution quality, thereby enabling scalable and computationally efficient coordination of a large number of DERs while preserving near-optimal FFR regulation performance.

Finally, we introduce prosumer motivation into VPP operation by recognising that DER prosumers are not purely price-driven actuators but may be intrinsically motivated to support system reliability. A nudging mechanism is proposed based on customised recommendations: each prosumer receives a price signal and a recommended service level and adjusts their decision based on these signals. Prosumer responsiveness is captured by behavioural surrogate models that quantify their tendency to follow recommendations, and a Social Conformity Index is introduced to measure how closely realised responses track the nudged targets. A data-driven inverse-optimisation procedure is proposed to learn the surrogate models from observed behaviour, enabling offline calibration and online adaptation. The proposed framework enables the VPP operator to reduce incentive costs while maintaining coordination performance. Moreover, it opens new avenues for human-centric control in decentralised energy systems.

Taken together, this thesis advances the state of the art by developing the dynamic modelling and the coordination framework for VPP-based provision of ancillary services and by addressing key challenges of scalability, limited disclosure of local information, and efficiency. By combining control-theoretic design with distributed algorithms and behavioural insights, it offers a coherent toolbox for future VPP implementations in low-inertia, inverter-dominated power systems.

Samenvatting

Versnelde decarbonisatie en de grootschalige uitrol van op omvormers aangesloten hernieuwbare energiebronnen veranderen elektriciteitssystemen in lage-inertie, zwak gekoppelde netten waarin de stabiliteit niet langer door een beperkt aantal synchrone machines kan worden gewaarborgd. Tegelijkertijd nemen gedistribueerde energiebronnen (distributed energy resources, DER's) zoals opslag, elektrische voertuigen, PV op daken en flexibele belastingen sterk toe in distributienetten, waardoor aanzienlijke maar gefragmenteerde flexibiliteit ontstaat. Virtuele energiecentrales (virtual power plants, VPP's) aggregeren deze DER's tot bestuurbare portfolio's, maar ondersteunende diensten (ancillary services) worden nog vaak behandeld als statische reservevermogens of vermogenssetpoints, waarbij het dynamische gedrag dat nodig is voor snelle frequentierespons (fast frequency response, FFR), inertienabootsing en andere kortetijdschaaldiensten wordt veronachtzaamd. Dit proefschrift hanteert daarom het concept van dynamische virtuele energiecentrales (dynamic virtual power plants, DVPP's), waarin de interne dynamica van apparaten expliciet wordt gemodelleerd.

Het proefschrift ontwikkelt een coördinatiekader en bijbehorende operationele modellen die geaggregeerde DER's in staat stellen om betrouwbare dynamische ondersteunende diensten te leveren, rekening houdend met schaalbaarheidsbeperkingen, beperkte openbaarmaking van lokale informatie en gedragsmatige overwegingen. Het is opgebouwd rond drie thema's: het modelleren van DVPP-operatie voor dynamische diensten, het ontwerpen van gedistribueerde en schaalbare oplossingsmethoden, en het integreren van prosumermotivatie in de levering van ondersteunende diensten.

Allereerst onderzoeken wij de coördinatie tussen meerdere DER-aggregatoren, die elk een portfolio van DER's beheren, om gezamenlijk FFR te leveren overeenkomstig de doelstelling van de DVPP-operator. FFR wordt weergegeven als een overdrachtsfunctie van frequentieafwijkingen naar actief-vermogensrespons, en een op $\mathcal{H}_\infty$ gebaseerde prestatiemaat kwantificeert de dynamische afwijking tussen de geaggregeerde respons en de FFR-doelwaarde. Dit resulteert in een convex optimalisatieprobleem met lineaire matrixongelijkheidsbeperkingen, waarin zowel de kwaliteit van de dynamische regulering als de economische efficiëntie worden vastgelegd. Om een volledig gecentraliseerde op-

lossing te vermijden die het delen van vertrouwelijke gegevens zou vereisen, wordt het probleem geherformuleerd als een toestandsgebaseerd potentieelspel, waarbij elke aggregator optreedt als een agent met een eigen kostenfunctie en toelaatbare set. Een specifiek ontworpen speltheoretisch leerschema stelt de agents in staat hun FFR-bijdragen bij te werken op basis van uitsluitend lokale informatie en communicatie met naburige agents, en te convergeren naar een Nash-evenwicht dat de optimale oplossing nauw benadert, terwijl de openbaarmaking van lokale informatie wordt beperkt en single points of failure worden vermeden.

Vervolgens wordt de DVPP bestudeerd, die bestaat uit een groot aantal DER's in een distributienet en als geheel geaggregeerde FFR aan het transmissienet levert. De werking van de DVPP wordt nog steeds geformuleerd als het matchen van een gewenste FFR-overdrachtsfunctie op het netknooppunt met de geaggregeerde overdrachtsfunctie van alle deelnemende apparaten. Deze matching wordt in het frequentiedomein afgedwongen door de reële en imaginaire delen van beide overdrachtsfuncties op geselecteerde hoeksnelheden op elkaar af te stemmen. Dit leidt tot een niet-lineair, niet-convex optimalisatieprobleem dat ontwerpvariabelen op apparaatsniveau, inclusief responstijden (latencies), koppelt aan de DVPP-prestaties en de operationele kosten. Latencies worden behandeld als beslissingsvariabelen, waardoor de DVPP zowel de amplitude als de timing van zijn respons kan vormgeven en aanvullende temporele flexibiliteit kan ontsluiten. Om de computationele tracteerbaarheid bij grote portfolio's te waarborgen, wordt een dual-decompositiemethode voorgesteld die het globale probleem opsplitst in deelproblemen per apparaat, die uitsluitend via laag-dimensionale duale variabelen aan elkaar zijn gekoppeld. De voorgestelde oplossingsmethode biedt aantoonbare garanties voor de oplossingskwaliteit en maakt daarmee schaalbare en computationeel efficiënte coördinatie van een groot aantal DER's mogelijk, terwijl een nagenoeg optimale FFR-regelprestatie behouden blijft.

Tot slot introduceren we prosumentmotivatie in de VPP-operatie door te erkennen dat DER-prosumers niet louter prijs-gedreven actuatoren zijn, maar ook intrinsiek gemotiveerd kunnen zijn om de systeem-betrouwbaarheid te ondersteunen. We stellen een nudging-mechanisme voor op basis van op maat gemaakte aanbevelingen: elke prosumer ontvangt een prijssignaal én een aanbevolen serviceniveau en past zijn/haar besluit aan door lokale kosten af te wegen tegen afwijkingen van de aanbeveling. De responsiviteit van prosumers wordt vastgelegd met een commitmentfactor, die hun neiging om aanbevelingen te volgen kwantificeert, en we introduceren een Social Conformity Index om te meten hoe nauw gerealiseerde reacties de genudgede doelwaarden volgen. Daarnaast stellen we een datagedreven inverse-optimalisatieprocedure voor om commitmentfacto-

ren uit geobserveerd gedrag te identificeren, waardoor zowel offline kalibratie als online adaptatie mogelijk wordt. Het voorgestelde raamwerk stelt de VPP-operator in staat de incentive-kosten te verlagen terwijl de coördinatieprestatie behouden blijft. Bovendien opent het nieuwe perspectieven voor mensgerichte (human-centric) regeling in gedecentraliseerde energiesystemen.

Alles bij elkaar genomen brengt dit proefschrift de stand van de techniek verder door de dynamische modellering en het coördinatiekader voor VPP-gebaseerde levering van nevendiensten te ontwikkelen en door kernuitdagingen op het gebied van schaalbaarheid, beperkte openbaarmaking van lokale informatie en efficiëntie aan te pakken. Door regeltechnisch ontwerp te combineren met gedistribueerde algoritmen en gedragsinzichten, biedt het een samenhangende gereedschapskist voor toekomstige VPP-implementaties in laag-inertie, door omvormers gedomineerde elektriciteitssystemen.

List of Abbreviations

AC	Alternating Current
ADMM	Alternating Direction Method of Multipliers
BRL	Bounded Real Lemma
CO_2	Carbon Dioxide
DC	Direct Current
DER	Distributed Energy Resource
DVPP	Dynamic Virtual Power Plant
FFR	Fast Frequency Response
FOP	Forward Optimization Problem
IOP	Inverse Optimization Problem
LMI	Linear Matrix Inequality
PV	Photovoltaics
RLS	Recursive Least Squares
RoCoF	Rate-of-Change-of-Frequency
SCC	Short-Circuit Current
SCI	Social Conformity Index
TSO	Transmission System Operator
VPP	Virtual Power Plant

1

Introduction

In this chapter, the background and motivation for the thesis are introduced. The research scope is then clarified, followed by an outline of the overall structure of the thesis.

1

1.1. Background and Motivation

As global warming intensifies, the world faces growing pressure to cut greenhouse gas emissions and decarbonise energy systems. Different regions and countries have set various targets regarding the trends of decarbonisation. The EU targets at least a 55% emissions reduction by 2030 from 1990 levels and climate neutrality by 2050 [1]. The United States pledges a 50–52% reduction below 2005 emissions by 2030 and an economy-wide net-zero goal by 2050 [2]. China commits to peaking CO_2 emissions before 2030 and achieving carbon neutrality before 2060, supported by large-scale renewable deployment and power-system transformation [3]. Similarly, the United Kingdom has legislated a net-zero target for 2050 and aims at a 68% reduction by 2030 and 78% by 2035 relative to 1990 [4].

In this context, renewable sources such as wind and solar are increasingly replacing traditional fossil fuels, thanks to their low-carbon and sustainable nature. Although these increasing renewable resources provide substantial clean power, they also introduce new operational challenges for power systems [5]. As wind and solar generation depend on weather conditions, their output is inherently uncertain and time-varying, making it difficult to maintain the generation-load balance [6]. Moreover, the renewables are fundamentally DC sources and connect to the AC grid through inverters, which significantly alter the grid's dynamic behaviour [7]. The resulting inverter-dominated system exhibits low rotational inertia and weaker system strength due to limited overloading capability [8]. Together, these emerging features pose unprecedented stability concerns and operational risks for modern power systems [9].

Numerous blackout events have highlighted these operational risks. Most prominently, on 28 April 2025 the Iberian Peninsula experienced an almost complete system blackout: within seconds, Spain and Portugal, as well as parts of southwest France, lost supply following a cascade of events in a power-electronics-dominated grid. At the time, more than half of generation came from inverter-interfaced renewables, while only a minimum number of synchronous units with dynamic voltage-control capability were online [10]. Technical investigations report low-frequency inter-area oscillations between Iberia and the rest of Continental Europe [11]. In response, corrective switching actions were taken to control these oscillations, but they unintentionally increased transmission-level voltages and further reduced system strength. A combination of overvoltages, weak-grid conditions, and heterogeneous protection settings then caused a large number of (mostly converter interfaced) generating units to trip in rapid succession, eventually isolating Iberia from the rest of Continental Europe [12]. Once islanded, the Iberian system

experienced a fast frequency decline and under-frequency load shedding. Although the instantaneous share of variable renewables on that day was not exceptional for springtime operation, the coexistence of low online inertia, limited dynamic voltage control, stressed interconnections, and oscillatory behaviour created a system that was dynamically fragile to coupled frequency-voltage disturbances.

Another notable blackout occurred in Great Britain on 9 August 2019. The event was triggered by a lightning-induced fault and was then compounded by the near-simultaneous loss of a large synchronous generator and a wind plant, together with the tripping of distributed generation. System frequency dropped to about 48.8 Hz, triggering automatic under-frequency load shedding that disconnected roughly 1.1 million customers and resulting in over £10.5 million losses [13, 14]. This shows that in a low-inertia system with uncertain distributed-generation behaviour, frequency can drop faster and further than the procured reserves and existing under-frequency load shedding schemes are designed to cope with.

In South Australia, on 28 September 2016, a state-wide blackout occurred under conditions of high wind penetration and reduced online synchronous generation. A series of storm-induced transmission faults caused repeated voltage depressions. Multiple wind farms then disconnected after exhausting their low-voltage ride-through limits, leading to a rapid loss of generation and a high Rate-of-Change-of-Frequency (RoCoF). These effects ultimately caused the collapse of the islanded grid [15]. Beyond frequency stability, the event also highlighted emerging challenges in voltage support, fault current provision, and protection coordination in inverter-dominated, weak grids. Current-limited wind inverters provided only limited three-phase short-circuit current and dynamic reactive support during faults, so voltage depressions were deeper and longer than expected, and protection schemes designed for higher synchronous fault levels contributed to cascading disconnections of generation and transmission elements [15].

These severe events collectively underscore the necessity of well-designed ancillary services to help low-inertia power systems maintain secure operation. They also motivate a conceptual shift: ancillary services can no longer be regarded as static power set-points or reserve margins, but as diverse, dynamic behaviours of frequency, voltage, and fault response that must be explicitly defined, engineered, procured, and coordinated at the system level [16].

At the same time, the resources capable of delivering ancillary services are also transforming. Rather than relying solely on a small number of large synchronous plants, modern power systems increasingly draw on Distributed Energy Resources (DERs) such as

battery energy storage systems, rooftop photovoltaics, small wind, electric vehicles, heat pumps, and flexible industrial and commercial loads to shoulder a growing share of these service obligations [17, 18]. Individually, these units are small, geometrically distributed in the system, and often operated to optimise customer objectives. However, when appropriately coordinated, these numerous small units can collectively reproduce sizable flexibility and responses, which are traditionally expected to be provided by conventional generators [19].

Virtual Power Plants (VPPs), emerge as a promising paradigm for coordinating DERs and delivering their aggregate ancillary services to the power system. They collect heterogeneous resources and present them to the system operator as one or several equivalent controllable portfolios that interact with the grid much like conventional generating units [20]. By internalising device-level constraints, customer preferences, and network limitations, VPPs unlock flexibility from numerous small-scale DERs that would otherwise be invisible or infeasible to dispatch individually. For the transmission system operator, this provides a simplified interface: instead of scheduling a large population of DERs, only a limited number of aggregated portfolios with larger effective size need to be considered [21]. This reduction in dimensionality eases operational planning, market clearing, and real-time control, while the detailed coordination of individual assets remains within each VPP. In this way, VPPs enable system operators to procure reliable ancillary services from DERs without being exposed to the full complexity and diversity of individual devices.

In many countries, recent policy and market reforms explicitly recognise and support VPPs in their role as aggregators of DERs [22, 23, 24]. In Europe, the Clean Energy Package and subsequent national implementations recognise independent aggregators as market participants; in the United States, FERC Order 2222 mandates the participation of aggregated DERs in wholesale markets; and in China, large-scale VPP demonstrations are being developed within the 'New Power System' agenda to unlock flexibility from industrial parks, buildings, and storage. Taken together, these developments indicate that VPPs are increasingly regarded as critical, proactive participants in system operation.

Coordination is certainly at the heart of operating a VPP, and appears as the first theme of this thesis. Intuitively, as in an orchestra, different sections—strings, brass, woodwinds, and percussion—each with their own timbres, tempi, dynamics, and attacks, must still be aligned into one coherent ensemble. Likewise, each DER has its own dynamics, constraints, location, and owner objectives, yet the power system only "hears" the aggregate response of the portfolio [25]. From a modelling perspective, effective VPP operation therefore requires formulating mechanisms that translate system-level ancillary-service

requirements into compatible local actions. Most existing VPP formulations describe this coordination in terms of *static* quantities: energy schedules, reserve capacities, and set-points for active and reactive power, embedded in day-ahead or intra-day optimisation models [26, 27, 28, 29]. In contrast, the coordination of DERs for dynamic ancillary services is much less explored. Services such as inertia emulation, fast frequency response, oscillation damping, and dynamic voltage support depend on closed-loop device dynamics and their interactions over short time scales, which must be captured explicitly rather than via static set-points. This dynamic viewpoint is often termed a Dynamic Virtual Power Plant (DVPP), where the VPP explicitly models internal device dynamics to deliver dynamic ancillary services. This thesis builds on this DVPP perspective, focusing on modelling and formulating coordination problems so that the aggregate response of many DERs realises the required dynamic ancillary services in a reliable way [30, 31, 32].

A second theme of this thesis is *computation and solution methods*. In practice, a VPP may comprise a large population of DERs, each with its own state, constraints, and preferences [30, 33]. Solving a detailed, centrally coordinated optimisation or control problem at this scale, and at the required time resolutions, quickly becomes computationally prohibitive and communication-intensive [34]. In addition, many DER prosumers are unwilling to reveal their detailed cost functions or local constraints to the VPP operator due to privacy and autonomy concerns, and the underlying communication graph and local information can change frequently as devices connect, disconnect, or change operating modes. The orchestra analogy again offers insights: in a real performance, coordination in a symphony does not rely solely on the conductor's baton; musicians constantly listen, adapt, and self-synchronise within and across sections, as visual access to the conductor is limited and the musical interactions are too rich to be micro-managed note by note. Similarly, VPP operation calls for distributed solution methods and operation schemes, in which DERs adapt locally in response to compact coordination signals, yet collectively realise the desired dynamic ancillary services [35, 36].

A third theme of this thesis concerns the *self-motivation of DER prosumers*. With increasing awareness of climate change and targeted education and outreach, many households, businesses, and communities that own DERs are no longer passive recipients of control signals; they actively express interest in contributing to system flexibility and ancillary-service provision, provided that their comfort, production needs, and autonomy are respected. This idea is perhaps better understood through the orchestral analogy: great performances rely not only on the conductor's technique, but on the intrinsic artistry of the musicians, who aim for beauty and cohesion rather than merely playing the right

notes for a paycheck. In a similar spirit, VPPs should not treat prosumers as reluctant actuators, but as motivated collaborators whose local decisions shape the aggregate dynamic response. This calls for coordination and incentive mechanisms that leverage prosumers' intrinsic motivations and preferences so that contributing to dynamic ancillary services becomes aligned with what DER prosumers are willing and eager to do [20]. This thesis therefore also explores how such human-centric mechanisms can be embedded in VPP operation to harness prosumer engagement and increase operational efficiency in a systematic way.

Motivated by the circumstances described above, the overarching research objective of this thesis is as follows:

Develop a coordination framework and associated operational models, from the perspective of VPPs, for orchestrating large populations of DERs to provide reliable power system ancillary services.

The key philosophy behind this thesis lies in filling the gaps in modelling and formulation, computational scalability, and prosumer-oriented design of VPP operation, ultimately enabling VPPs to act as effective orchestrators of DERs in future low-inertia power systems.

1.2. Research Scope and Questions

These developments motivate a fundamental shift in how ancillary services are defined, engineered, and operated in future low-inertia power systems. Rather than treating ancillary services as static reserves supplied by a few large generators, this thesis views them as dynamic behaviours delivered by large populations of heterogeneous DERs coordinated through VPPs. The overall conceptual framework adopted in this thesis is summarised in Figure 1.1. As illustrated in Figure 1.1, $G(s)$ denotes the dynamic model of the power system. Within the VPP, there is a set of DERs indexed by $i \in I = \{1, 2, \dots, |I|\}$, and each DER $i \in I$ is represented by a local model $F_i(s)$ that delivers an ancillary-service signal r_i. From the system viewpoint, the individual contributions r_i are combined in the equivalent aggregation layer into a single response $\hat{r}$, which is exposed to the grid as the VPP's service provision. For each device, the realised service r_i is determined by a set of controllable parameters or set-points, collected in θ_i, and constrained by a feasible set Θ_i capturing device-level limits and behavioural requirements. The core problem addressed in this thesis lies in the VPP coordination block: designing algorithms that compute θ_i for all DERs such that (i) the aggregate response $\hat{r}$ tracks the VPP-level service target $r^\star$, (ii) all local constraints Θ_i are satisfied, and (iii) the overall cost of providing ancillary ser-

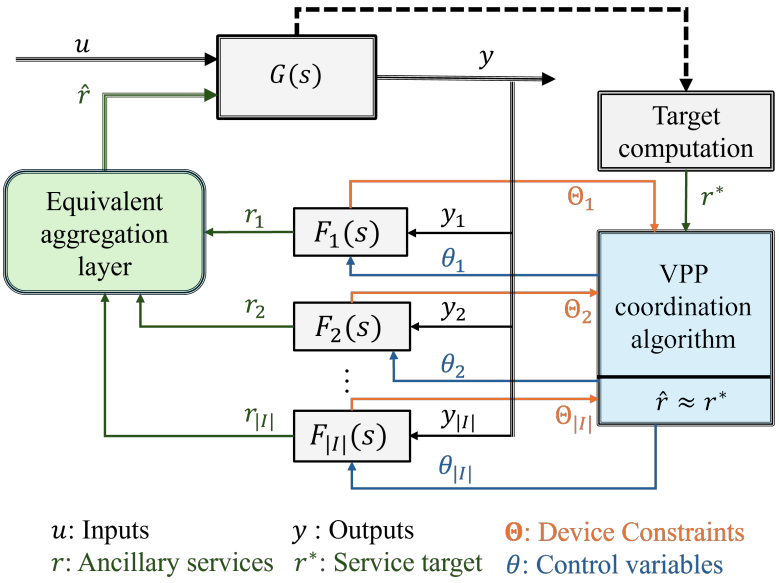

Figure 1.1: Block diagram of VPP operation for ancillary service provision.

vices is minimised while respecting the behaviour patterns of the prosumers behind the devices. Here, $r^\star$ represents the ancillary service target requested by the system operator. The framework assumes that the VPP can participate in providing this service, but it does not assume that the target can always be met exactly. Rather, the mismatch between the realised aggregate response $\hat{r}$ and the target $r^\star$ is explicitly quantified in the objective function as a measure of service quality, and the coordination problem is formulated to minimise this mismatch subject to device-level and operational constraints.

Building on this conceptual framework, the remainder of the thesis is organised around three central research questions, each aligned with one theme identified in the background and motivation: (i) closing the formulation gap for dynamic ancillary services via the DVPP perspective, (ii) ensuring computational scalability and distributed implementability of coordination schemes, and (iii) modelling DER prosumers' motivations and nudging them to participate into VPP operation.

RQ1 – Modelling DVPP operation for dynamic ancillary services. *Taking a DVPP providing fast frequency response (FFR) as a prototype, how can we mathematically formulate its regulation task for dynamic ancillary services? Based on this formulation, how should the*

1

DVPP operation profiles be chosen so that regulation quality and economic efficiency are both taken into account? How can the resulting model be structured so that it remains scalable and analytically tractable as the DVPP portfolio grows?

To answer RQ1, we first need to clarify what FFR means as a dynamic ancillary service and find a precise expression for it. FFR is an automatic active-power response triggered by locally measured frequency deviations, acting within a time frame of a few hundred milliseconds to a few seconds. It modifies active power continuously over this short interval, rather than through discrete set-point changes. Therefore, its representation should compactly capture the amplitude, timing, and shape of this dynamic behaviour. In other words, it is natural to describe FFR as a dynamic system rather than as a static power set-point. This motivates us to use parametric transfer functions to represent the dynamic systems associated with FFR. Within the block-diagram framework of Figure 1.1, for the FFR case the ancillary-service signal $r(s)$ represents the active-power response during the transient, $y(s)$ denotes the system frequency deviation, and FFR itself is described by the underlying system model $F(s) := r(s)/y(s)$. A key objective in operating a DVPP providing FFR is to ensure that the aggregate response of the DERs $\hat{r}(s)$ closely tracks the requested response $r^\star(s)$ specified by the system operator, for any frequency deviation signal $y(s)$. Equivalently, we aim to coordinate the individual DER FFR models $F_i(s)$ so that their aggregation approximates a desired FFR target $F^\star(s)$.

The core technical problem is how to quantify the mismatch between the target FFR and the DVPP's aggregate response. This is challenging and non-trivial, because the mismatch itself is a dynamic system. We must extract a metric that translates this dynamic behaviour into a scalar measure of regulation quality. It is also critical that the chosen metric is analytically and computationally tractable, so that it can be embedded into a common framework that simultaneously accounts for the overall cost of DERs' FFR provision. Representing FFR using transfer functions naturally points us to leverage system norms and frequency-domain response characteristics as candidates for such a metric.

Another core question here concerns the scalability of DVPP operation. As the number of DERs increases, and as their FFR responses differ in latency and shape, the aggregate response and the associated mismatch dynamics become higher-order and more complex. Directly deriving analytical expressions and system norms for these high-order models quickly becomes intractable. It is therefore crucial to address scalability explicitly, by structuring the formulation so that the complexity of the dynamic-matching metric and the associated decision variables grows in a controlled way with the portfolio size, and by retaining a model that remains amenable to systematic analysis and practical computation

as the DVPP scales up.

RQ2 – Distributed solution methods. *Given the formulation developed in RQ1, how can we efficiently compute high-quality operating profiles for a DVPP providing dynamic ancillary services? Is there a distributed solution method that allows DERs to participate without revealing detailed local private information and that avoids single points of failure? How can we ensure that the solution method remains tractable and efficient when the underlying optimisation problem is highly nonconvex and the DVPP scales to large portfolios?*

The practical implementation of a VPP naturally advocates a fully distributed coordination scheme. Built upon the formulation developed in RQ1, DVPP coordination must respect data privacy and decentralised information. Heterogeneous DER prosumers possess sensitive local information, such as detailed operating cost functions, availability, and device constraints, which should not be fully shared with a central coordinator [37]. The DVPP portfolio is also dynamic: devices may join, leave, or change their communication neighbours during service provision. A suitable method must therefore allow each DER to compute its decision locally, exchange only limited coordination signals, and remain robust to time-varying participation and communication graphs, without single points of failure. This motivates the use of a game-theoretic learning framework and raises the critical question of how to design it so that the resulting equilibrium coincides with, or closely approximates, the desired DVPP operating point.

From the perspective of tractability, when devices with variable response latencies are included, the operation problem formulated in RQ1 becomes a nonlinear, nonconvex programme. As the DVPP scales up, this leads to a large-scale nonconvex optimisation problem that is prohibitive for existing centralised solvers. It is therefore important to address this computational issue so that the DVPP operation model remains scalable. This motivates us to exploit the specific structure of the problem and develop solution methods that can deliver high-quality operating profiles. By leveraging decomposition methods, the large computational task can be distributed into multiple smaller subproblems and solved in parallel, thereby reducing computation time and enabling practical implementation for large DVPP portfolios.

RQ3 – Prosumer motivation and nudging the participation. *How can we systematically leverage the intrinsic willingness of DER prosumers to enhance the efficiency and reliability of VPP operation? In particular, how should prosumer behaviour be modelled, including their self-motivation, comfort and autonomy preferences, so that the incentive mechanisms*

align their local decisions with the provision of ancillary services?

To answer RQ3, let us first consider how self-motivation manifests in observable behaviour. Intrinsic motivation becomes visible when individuals accept and carry out requests voluntarily, without expecting payment or direct financial benefit. Agreeing to an explicitly uncompensated request reveals a genuine willingness to contribute to the underlying goal rather than a response driven by monetary incentives. The same logic underlies charitable behaviour, where people donate time or money to public causes without direct material returns. In the context of VPPs, prosumers who accept flexibility requests with little or no compensation can thus be seen as signalling intrinsic motivation to support system reliability or the energy transition. An operation framework that systematically recognises and leverages such behaviour can move beyond purely price-based coordination, better align VPP operation with prosumers' internal motivations, and at the same time help manage or even reduce the VPP's overall incentive budget.

Enlightened by these observations, we advocate providing each DER prosumer with an additional piece of information: a recommendation on the amount of ancillary service they are requested to provide. Without enforcing any obligation, and while still offering an incentive price, the intrinsic motivations of DER owners can then be realised when they voluntarily adjust their decisions towards these recommendations. We term such a mechanism a *nudge*, as it acts as a minimalist intervention that predictably influences DER owners' behaviour without prohibiting any options or significantly altering their economic incentives [38].

To design such a nudging mechanism for VPP operation goals, the first critical step is to establish a behavioural model for each DER prosumer that captures how they respond jointly to prices and recommendations. Such a model should quantify the strength of intrinsic motivation, explain how prosumers adjust their service provision in the presence of a nudge, and remain consistent with their observed decisions. Data-driven and mining-based approaches are natural candidates here, since behavioural patterns can only be recognised and identified from observed responses, for example using data collected via questionnaires, pilots, or controlled tests. The next challenge is to design customised recommendation signals for heterogeneous DERs. Different prosumers may exhibit very different levels of motivation to follow service recommendations, depending on factors such as technical knowledge, confidence in participating in system services, sense of social responsibility, and local community environment [39, 40]. To achieve the VPP's operational goals, nudges must therefore be tailored so that the aggregate behaviour of all participating DERs collectively realises the desired response. How to construct such customised

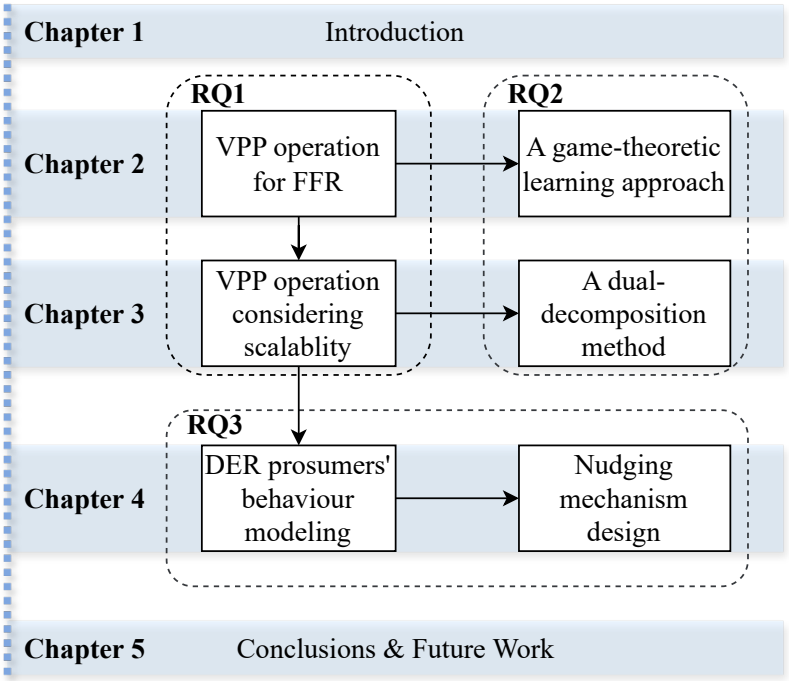

Figure 1.2: Thesis outline and relation to the research questions.

recommendations in a systematic and scalable manner remains a key problem for RQ3.

1.3. Contributions and Thesis Outline

Figure 1.2 illustrates the structure of the thesis and shows how the chapters jointly address the key research questions (RQ1, RQ2, RQ3) in pursuit of the overarching goal of efficient VPP operation. The remainder of this thesis is structured as follows:

- **Chapter 2** first answers **RQ1** by proposing a convex $\mathcal{H}_\infty$-norm-based DVPP coordination framework for FFR, which models FFR explicitly as a dynamic ancillary service, quantifies DVPP regulation performance, and computes optimal FFR provisions from multiple DER aggregators at the portfolio level. Building on this formulation, the chapter then addresses **RQ2** by developing a game-theoretic learning solution method that enables fully distributed, graph-based computation of the co-

ordinated FFR provisions without requiring centralised control or revealing the internal information of each aggregator. The proposed generic method reformulates the original centralised optimisation with a global coupling linear matrix inequality constraint into a state-based potential game by introducing local estimator states that decouple this shared semidefinite constraint. From a theoretical perspective, the resulting game is shown to admit stationary Nash equilibria that recover the centralised optimal solution and to be robust to asynchronous updates and time-varying communication graphs, thereby ensuring computational efficiency, privacy preservation, and practical implementability for DVPPs. The developments in this chapter are validated through a case study with multiple DER aggregators, in which the convergence behaviour of the learning scheme is examined under changing communication graphs and the resulting DVPP response is further assessed in the time domain against a baseline portfolio, showing improved FFR tracking performance.

- **Chapter 3** addresses the scalability issues identified in **RQ1** by presenting a frequency domain dynamic-matching formulation for the operation of the large-scale DVPP, in which large populations of heterogeneous DERs are coordinated to provide FFR to the system operator. The proposed formulation operates on frequency-domain representations of the devices' FFR impulse responses and keeps the number of coupling constraints dependent only on the chosen frequency grid, rather than on the number of DERs, thereby ensuring scalability to very large portfolios. Moreover, it naturally accommodates devices with variable response latencies by treating latency parameters as decision variables, thus explicitly exploiting the additional temporal flexibility inherent in DER FFR delivery. Building on this formulation, the chapter further contributes to **RQ2** by developing a dual-decomposition-based distributed solution method that solves the resulting large-scale nonconvex DVPP dynamic-matching problem efficiently. The proposed method decomposes the original problem into low-dimensional device-level subproblems coupled only through a small set of dual variables. It supports parallel implementation with simple projected subgradient updates and provides provable guarantees on solution quality. This enables scalable and computationally efficient coordination of large numbers of DERs while preserving near-optimal FFR regulation performance. These developments are validated on a modified IEEE 33-bus distribution system through three studies. The first examines DVPP regulation performance under low- and high-DER-penetration scenarios. The second evaluates the operational benefits of

latency-inclusive coordination. The third analyses the scalability and runtime of the distributed solution method for DVPP sizes ranging from 10 to 1000 DERs, demonstrating its computational efficiency for large-scale nonlinear optimisation problems.

- **Chapter 4** aims to answer **RQ3** by incorporating behavioural aspects of DER owners into VPP operation. To this end, it develops a behaviour-aware nudging mechanism, based on a tâtonnement process, to coordinate DER owners within a VPP through customised service recommendations and incentive payments. By tailoring recommendations to individual DER owners, the mechanism systematically translates intrinsic behavioural motivations into operational benefits for the VPP while supporting budget balance. To support this mechanism, the chapter further develops a behavioural surrogate model that captures how prosumers respond jointly to prices and recommendations, together with a data-driven inverse-optimization procedure that calibrates and updates this surrogate from historical or simulated signal-response data. In this way, the chapter enables both behaviour-aware coordination and adaptive learning of heterogeneous prosumer response patterns. The proposed mechanism is validated on a modified IEEE 39-bus transmission system with a VPP of 20 DERs through two ancillary-service use cases: synthetic inertia provision and short-circuit current support. The case studies demonstrate both the effectiveness of the nudging mechanism and the accuracy and adaptivity of the proposed data-driven behavioural-surrogate modelling approach.
- **Chapter 5** provides overall conclusions, summarising the key findings of the thesis and offering recommendations for future research and practical implementation.

These contributions advance the state of the art by providing efficient, reliable, and scalable tools for VPP operation, validated through both theoretical analysis and numerical case studies.

2

Distributed Game-Theoretic Learning for FFR

The Dynamic Virtual Power Plant (DVPP) envelops the Distributed Energy Resources (DERs) in the distribution network into a single operational entity that provides aggregate dynamic ancillary services to the transmission grid. This chapter investigates the operation of one DVPP consisting of several DER aggregators, where each aggregator is modeled as a strategic agent that responds to transmission-level Fast Frequency Response (FFR) requests. The problem is formulated from the perspective of suboptimal $\mathcal{H}_\infty$ control, leading to an efficient and tractable optimization framework. On this basis, a distributed solution method based on game-theoretic learning is developed that enables the DER aggregator agents to adaptively learn their optimal provisions. Specifically, the original problem is reformulated as a state-based potential game, in which agents interact to reach a carefully designed Nash equilibrium.

This chapter is based on **H. Xie**, *and J. L. Cremer, "Game-Theoretic Learning for Power System Dynamic Ancillary Service Provisions," IEEE Control Systems Letters, Vol. 8, June 2024.* [41].

2.1. Introduction

The large-scale integration of renewable energy resources deteriorates power system inertia and, in turn, poses severe challenges to frequency stability under major contingencies. To arrest the ensuing frequency deviations, active power support must be delivered rapidly in the immediate post-disturbance period. Dynamic Virtual Power Plants (DVPPs) have emerged as promising providers of such Fast Frequency Response (FFR) dynamic ancillary services [20, 42]: a DVPP aggregates (typically inverter-based) DERs and coordinates them to deliver a rapid, aggregate post-disturbance active-power injection. In this chapter, the coordination is performed at the *aggregator* level: each DER aggregator represents and controls a local DER portfolio and is modeled as one agent. The stringent activation-time requirements imply that FFR must be pre-programmed before contingencies occur, so that predetermined post-disturbance active-power dynamics are automatically triggered by measured frequency deviations.

The FFR service requested from a given DVPP must be specified so as to guarantee acceptable frequency performance under the most critical contingencies. This requirement depends on regional grid codes, the considered contingency set, and the current system condition (e.g., inertia and damping levels). In practice, the post-disturbance system frequency response is evaluated to verify that key performance indices—such as frequency nadir and rate of change of frequency (ROCOF)—satisfy the relevant grid-code limits [43]. Once an FFR target is defined and assigned to a DVPP, its operation is typically driven by two overarching objectives. First, the aggregate frequency response of the participating DER aggregators should closely track the dynamic profile specified by the Transmission System Operator (TSO) and encoded in ancillary-service products [44, 45]. Second, the operating point should be economically efficient, commonly expressed as minimizing the sum of individual DER aggregator cost functions [32, 46].

A substantial body of work has examined how to coordinate DERs within a DVPP so that their FFR provisions align with TSO requirements [32, 43, 46, 47, 20, 42, 48]. Broadly, two modeling and coordination paradigms can be distinguished. The first, and still most common, line of research replaces the full dynamics of FFR with a small number of equivalent adjustable parameters—synthetic inertia, damping, and primary active-power gain being typical examples [32, 43, 46]. The DVPP's aggregate response is then obtained by summing these parameters over all DER aggregators, and the operation problem becomes an algebraic matching between these aggregated quantities and the TSO's specified targets. While this approach offers clear physical intuition consistent with synchronous-generator behavior and yields analytically tractable formulations, it oversimplifies the

underlying dynamics and disregards device-dependent response delays. In reality, DER aggregators exhibit heterogeneous latencies due to differences in control bandwidths and hardware implementations [47, 20]; neglecting such latency effects can therefore degrade regulation accuracy and overall system performance.

The second research stream explicitly recognizes the dynamic nature of FFR and incorporates latency directly into the model [47, 42, 48]. In this line of work, FFR is described as a dynamic input–output system with frequency deviations as input and active power as output, often represented by rational parametric transfer functions in the frequency domain [49, 50, 51]. Such representations more accurately capture DER aggregator's behavior, while at the same time facilitating grid-code verification and systematic embedding of the response into converter control architectures.

Building on these dynamic models, several studies propose coordination strategies based on divide-and-conquer principles, whereby the overall control task is decomposed into device-level subproblems [20, 42]. However, these schemes often compromise economic efficiency: they typically assume that all contracted DERs participate in regulation and rely heavily on very fast resources (e.g., supercapacitors) to provide flexibility, which may increase costs and impose significant operational stress on the DVPP [20, 42]. Within this second stream, [47] adopts a complementary time-domain viewpoint. The authors derive explicit expressions for the system frequency trajectory and adjust DER behaviors based on associated performance metrics. This time-domain formulation is conceptually intuitive and highly interpretable, as it links regulation decisions directly to observable frequency responses. Nonetheless, it requires laborious analytical derivation of closed-form frequency trajectories and depends on approximate models of system frequency dynamics, which can limit both accuracy and scalability when network conditions, control structures, or portfolios of participating devices become more complex.

To address these challenges, we approach DVPP operation from a suboptimal $\mathcal{H}_\infty$-control perspective. The $\mathcal{H}_\infty$ norm is employed as a metric to assess the discrepancy between the total FFR provided by the aggregators and the prescribed DVPP FFR target. On this basis, we formulate a convex optimization problem that jointly maximizes overall economic efficiency and regulates DVPP's dynamic performance.

When coordinating aggregators for ancillary-service provision, it is also important to limit the disclosure of local information, such as operating costs and device limits, in order to address data confidentiality concerns [37]. In addition, the population of participating DERs is inherently uncertain: local operating conditions can change frequently, and units may connect to or withdraw from the DVPP over time. Distributed optimization tech-

niques offer a natural way to circumvent the need for sharing all local information with a central coordinator. Primal–dual methods and the Alternating Direction Method of Multipliers (ADMM) have been extensively studied to decompose coupling constraints and have seen successful applications in energy-management problems [52, 37]. However, these schemes typically still rely on a centralized *collect-and-broadcast* step to update dual variables, which prevents fully distributed, graph-based communication. As a result, they are vulnerable to information loss and communication uncertainties [53], a critical drawback in DVPP settings.

To mitigate these limitations, prior work has explored game-theoretic control techniques for distributed coordination under coupling constraints [53, 54]. Nonetheless, existing designs are not directly applicable to our problem, where the coupling constraint takes the form of a semidefinite cone arising from the underlying Linear Matrix Inequality (LMI) conditions. In this vein, we extend the game-theoretic decoupling approach of [55] to a semidefinite-program setting tailored to the proposed $\mathcal{H}_\infty$-based DVPP operation problem.

The main contributions of this work are summarized below.

- We propose an efficient and computationally tractable convex optimization framework for the optimal DVPP operation problem, leveraging $\mathcal{H}_\infty$ control theory to rigorously couple DVPP regulation targets with DER aggregators coordination and economic efficiency.

- We develop a fully distributed, game-theoretic solution methodology that enables aggregators to autonomously learn their optimal FFR provisions. The method does not require aggregators to disclose sensitive local cost or constraint information to a central entity or to other aggregators, and relies only on limited peer-to-peer communication. Moreover, the proposed scheme is robust to dynamically varying communication graphs.

- We carry out a comprehensive case study to demonstrate the effectiveness of the proposed framework, its convergence properties, and in particular its robustness against realistic communication uncertainties.

This integrated approach provides DVPP operators with an efficient and reliable methodology for harnessing inverter-based DERs to deliver dynamic ancillary services. The remainder of this chapter is organized as follows. Section 2.2 formulates the problem under

study. Section 2.3 introduces the game-theoretic control framework and state-based potential games. A distributed learning algorithm is proposed in section 2.4 to solve the formulated problem. Section 2.5 presents the case study, and section 2.6 concludes the chapter.

2.2. Problem Formulation

We consider a set of DER aggregators that provide FFR to the system operator, indexed by $I = \{1, \dots, |I|\}$. For each aggregator $i \in I$, the FFR provision is modeled as the sum of K components with different response latencies:

$$T_i(s) = \sum_{k=1}^{K} T_i^k(s) = \sum_{k=1}^{K} \frac{-D_i^k - H_i^k s}{\tau_k s + 1}, \tag{2.1}$$

where τ_k denotes the effective response latency of component k, that is, the characteristic time scale at which the corresponding FFR contribution becomes available. This latency represents the aggregate effect of measurement, communication, control, and actuation delays. Since different DER technologies can respond at different speeds, we assume that these heterogeneous responses are grouped into K predefined latency classes $\{\tau_1, \dots, \tau_K\}$. Accordingly, each aggregator can distribute its FFR provision across these classes according to the capabilities of its underlying resources. The coefficients H_i^k and D_i^k denote the controllable inertia-like and damping-like parameters associated with aggregator i's contribution in latency class k. The decision variables of aggregator i are collected in

$$v_i \doteq \{H_i^1, D_i^1, \dots, H_i^K, D_i^K\}.$$

The FFR demand from the system operator is assumed to be computed according to standard grid codes, taking into account the current operating condition, disturbance size, and closed-loop stability requirements. This demand is communicated to the DVPP as a target transfer function

$$T(s) = \frac{-D - Hs}{\tau s + 1},$$

where τ is the desired aggregate latency, and H and D are the target (equivalent) inertia and damping coefficients, respectively, all treated as given constants.

Remark. The structure in (2.1) has been shown to be achievable by DER aggregator, see [42, 21]. We adopt this form as it is consistent with advanced inverter-based control paradigms [56], so inverter-based DERs can be naturally accommodated in the framework.

In this context, H_i^k is commonly interpreted as (virtual) inertia, whereas D_i^k plays the role of damping.

Given the above setting, we define the residual between the DVPP FFR target and the aggregate response from all aggregators as

$$\begin{aligned}\mathcal{R}(s) &= T(s) - \sum_{i\in I}\sum_{k=1}^{\mathrm{K}} T_i^k(s) \\ &= \frac{-D - Hs}{\tau s + 1} - |I| \sum_{k=1}^{\mathrm{K}} \frac{-\bar{D}^k - \bar{H}^k s}{\tau_k s + 1} \\ &= \frac{b_{\mathrm{K}+1} s^{\mathrm{K}+1} + b_{\mathrm{K}} s^{\mathrm{K}} + \cdots + b_1 s + b_0}{a_{\mathrm{K}+1} s^{\mathrm{K}+1} + a_{\mathrm{K}} s^{\mathrm{K}} + \cdots + a_1 s + a_0},\end{aligned}$$

where the overbar $\bar{(\cdot)}$ denotes averaging over all aggregators, e.g., $\bar{D}^k = \sum_{i\in I} D_i^k/|I|$, and analogously for $\bar{H}^k$.

Remark. The coefficient vectors $\boldsymbol{a} = [a_0, \dots, a_{\mathrm{K}+1}]$ and $\boldsymbol{b} = [b_0, \dots, b_{\mathrm{K}+1}]$ collect the coefficients of the denominator and numerator polynomials of degree (K+1), respectively. The vector $\boldsymbol{a}$ is completely determined by the fixed parameters τ and τ_k; it does not depend on the aggregators' decisions. In contrast, each b_k is a linear combination of the decision variables in v_i, so $\boldsymbol{b}$ (and hence $\mathcal{R}(s)$) depends on the chosen FFR allocations across all aggregators.

We next derive a state-space representation of the residual system $\mathcal{R}(s)$ in controllable canonical form:

$$\begin{aligned}\dot{q} &= \mathcal{A}q + \mathcal{B}w \\ y &= \mathcal{C}q + \mathcal{D}w\end{aligned}$$

where q is the state vector, w denotes the frequency deviation (input), and y is the residual (output). The system matrices are given by

$$\mathcal{A} = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ -\hat{a}_0 & -\hat{a}_1 & -\hat{a}_2 & \cdots & -\hat{a}_K \end{bmatrix}_{(\mathrm{K}+1)\times(\mathrm{K}+1)} \qquad \mathcal{B} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix}_{(\mathrm{K}+1)\times 1} \tag{2.2}$$

$$\mathcal{C} = \begin{bmatrix} \hat{b}_0 - \hat{a}_0\hat{b}_{K+1} & \hat{b}_1 - \hat{a}_1\hat{b}_{K+1} & \cdots & \hat{b}_K - \hat{a}_K\hat{b}_{K+1} \end{bmatrix}_{1\times(\mathrm{K}+1)} \qquad \mathcal{D} = \hat{b}_{K+1}$$

with normalized coefficients

$$\hat{a}_i = \frac{a_i}{a_{\mathrm{K}+1}}, \quad i = 0, \dots, \mathrm{K}, \qquad \hat{b}_i = \frac{b_i}{a_{\mathrm{K}+1}}, \quad i = 0, \dots, \mathrm{K}+1.$$

By construction, the matrices $\mathcal{C}$ and $\mathcal{D}$ depend on the polynomial coefficients $\boldsymbol{b}$, and thus on the collection of all aggregator decision vectors $v := \{v_1, \dots, v_n\}$. Since b_k are linear functions of v, $\mathcal{C}$ and $\mathcal{D}$ are also linear functions of v. We can therefore decompose them into constant and decision-dependent parts as

$$\mathcal{C}(v) = \mathcal{C}_0 + \Delta\mathcal{C}(v), \qquad \mathcal{D}(v) = \mathcal{D}_0 + \Delta\mathcal{D}(v),$$

where $\mathcal{C}_0$ and $\mathcal{D}_0$ collect all constant terms, while each entry of $\Delta\mathcal{C}(v)$ and $\Delta\mathcal{D}(v)$ is a linear function (inner product) of v with a constant coefficient vector. This explicit dependence will be exploited later when formulating the DVPP coordination problem in terms of the decision variables v.

Ideally, we would like the residual to be *small* enough, meaning the aggregators' provisions are well achieving the system-level requirements. We therefore employ $\mathcal{H}_\infty$ norm to quantify this *smallness*, as $\mathcal{H}_\infty$ norm of a system provides the maximum energy gain of the output signal y (the residual) for a given input signal w (the frequency deviation). This quantification further endows us with a measure of the system-level performance. To see this, we introduce a variable γ here to constrain the $\mathcal{H}_\infty$ norm of the residual system $\mathcal{R}(s)$, i.e., $|\mathcal{R}(s)|_\infty \leq \gamma$, and γ should be minimized as a term in the objective function. As γ goes to 0, the residual system's output will also reduce to 0, which means the responses from the aggregators are perfectly matching the target. In order to achieve economic optimality, we also introduce $\sum_{i\in I} g_i(v_i)$ in the objective function, with g_i the continuous differentiable convex function representing the cost function of aggregator i regarding its decision variables. To sum up, with a parameter w controlling the trade-off between two terms, we have:

$$\min_{v_i, \gamma, Q} \quad \sum_{i\in I} g_i(v_i) + w\gamma \tag{2.3a}$$

$$\text{s.t.} \quad v_i \in V_i, \quad \forall i \in I \tag{2.3b}$$

$$|\mathcal{R}(s)|_\infty \leq \gamma. \tag{2.3c}$$

Eq. (2.3b) encodes the local constraints limiting the decision variables of each aggregator, where V_i is a polytope representing the domain of v_i. Based on the Bounded Real Lemma

(BRL) [57], we transform the constraint (2.3c) into a LMI constraint (with $\mathcal{Q} \succeq 0$):

$$\underbrace{\begin{bmatrix} \mathcal{Q}\mathcal{A}^T + \mathcal{Q}\mathcal{A} & \mathcal{Q}\mathcal{B} & \mathbf{0} \\ \mathcal{B}^T\mathcal{Q} & -\gamma I & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -\gamma I \end{bmatrix}}_{:=L(\mathcal{Q},\gamma)} + \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{0} & \Delta\mathcal{C}(v)^T \\ \mathbf{0} & \mathbf{0} & \Delta\mathcal{D}(v)^T \\ \Delta\mathcal{C}(v) & \Delta\mathcal{D}(v) & \mathbf{0} \end{bmatrix}}_{:=\sum_{i\in I} L_i(v_i)} \preceq \underbrace{\begin{bmatrix} \mathbf{0} & \mathbf{0} & -\mathcal{C}_0^T \\ \mathbf{0} & \mathbf{0} & -\mathcal{D}_0^T \\ -\mathcal{C}_0 & -\mathcal{D}_0 & \mathbf{0} \end{bmatrix}}_{:=F_0}.$$

The LMI constraint defines a semidefinite cone, which preserves the nice property of convexity. By construction, each matrix $L_i(v_i)$ collects the part of the global LMI that depends on the local decision vector v_i; every entry of $L_i(v_i)$ is a linear function of v_i. All purely constant contributions are gathered in the fixed matrix F_0.

To unify the notation, we also model the central regulator as agent 0, with decision vector $v_0 := (\gamma, \mathcal{Q})$ and associated LMI contribution $L_0(v_0) := L(\mathcal{Q}, \gamma)$. We denote its local cost by $g_0(v_0) := w\gamma$. Each DER aggregator $i \in I$ retains its own decision vector v_i and local cost function $g_i(v_i)$. Defining the extended agent set $\tilde{I} := \{0\} \cup I$, the global objective (2.3a) can be written compactly as the sum of separable, agent-wise costs:

$$\sum_{i\in\tilde{I}} g_i(v_i).$$

In the same spirit, the BRL-based LMI constraint can be expressed as

$$\sum_{i\in\tilde{I}} L_i(v_i) \preceq F_0.$$

Hence, the problem (2.3) has a separable objective but a coupling LMI constraint that links all agents through the common residual performance requirement. This structure motivates us to further develop distributed algorithms that effectively decouple the coupling constraint (2.3c) across agents.

2.3. Game-Theoretic Control and State-based Potential Games

The existence of coupling constraint (2.3c) implies that the decision made by one agent is constrained by the other agents' decisions. This is inherently consistent with the feature of a game, where the utility of one agent is affected not only by its own decisions, but also by the other agent's decisions. Stemming from this similarity, game-theoretic control provides a distributed alternative to handle the coupled optimization problems concerning

multiple agents [58]. The procedures of game-theoretic control design is to reformulate a control problem into a game, via i) endowing each agent a utility function (which should well encode the original costs), and ii) endowing each agent a learning rule (which will be programmed as control laws). It is expected that such a successful reformulation could bring desirable collective behavior for the system designer. In our context, the desirable collective behavior is that all the aggregators achieve their optimal provision as in (2.3). Before we proceed, let us revisit some preliminaries on the state-based potential game, the properties of which attract us to formulate problem (2.3) into such a game.

2.3.1. Preliminaries: State-Based Potential Games

Definition 2.1 (State-based game [59]). A deterministic game is regarded as a *state-based game* when there exists i) agent set I; ii) state space X; iii) state-dependent action sets of the form $U_i(x), x \in X$ for each agent $i \in I$; iv) state-dependent cost functions $J_i(x,u) \in \mathbb{R}, u \in U(x) = \prod U_i(x)$ for each agent $i \in I$; and v) deterministic state transition function $\tilde{x} = f(x,u) \in X$ for $\forall x \in X$ and $u \in U(x)$. For any $x \in X$, it holds true that the state won't change if all the agents take the null action $0 \in U(x)$, i.e., $x = f(x,0)$. We denote such a state based game by the tuple $G = \{I, X, \{U_i(x)\}, \{J_i(x,u)\}, f\}$.

Definition 2.2 (State-based potential game [59]). A state-based game is a *state-based potential game* when there exists a potential function $\Phi : X \times U \to \mathbb{R}$. The potential function satisfies the following two conditions for every state $x \in X$:

C-1 For every agent $i \in I$, $u \in U(x)$, and $u_i' \in U_i(x)$, it holds that $J_i(x,u) - J_i(x,u_i',u_{-i}) = \Phi(x,u) - \Phi(x,u_i',u_{-i})$, where $-i$ denotes the set of all the agents other than agent i, and u_{-i} includes the actions of agents in $-i$.

C-2 for every action $u \in U(x)$ and ensuing state following the dynamic $\tilde{x} = f(x,u)$, we have $\Phi(x,u) = \Phi(\tilde{x},0)$.

Definition 2.3 (Stationary state Nash equilibrium [60]). A state-action pair $[x^*, u^*]$ is a *stationary state Nash equilibrium* for a state-based game if it satisfies the following two conditions:

C-1 We have $u_i^* \in \arg\min \{J_i(x^*, u_i, u_{-i}^*) | u_i \in U_i(x^*)\}$ for $\forall i \in I$,

C-2 x^* is a fixed point of the state transition function with u^*, i.e., $x^* = f(x^*, u^*)$.

2.3.2. Connection to the Game-Theoretic Control

State-based potential games possess two properties that are attractive in game-theoretic control:

2

- *Existence of Nash equilibrium.* The learning algorithms of the game ideally end up in the Nash equilibrium where no agent can improve their utility by updating their decisions anymore. We therefore wish to find a way to engineer the Nash equilibrium of the game, such that it is exactly located at the optimum we desire. The premise here is to have the guarantee that the Nash equilibrium does exist. Since the stationary state Nash equilibrium in a state-based potential game is guaranteed to exist [59], it is theoretically appealing to design a state-based potential game for the original problem, and accurately locate the equilibrium to be equivalent or close to the optimum.

- *Existence of distributed Nash equilibrium-seeking algorithms.* It is established that for a state-based potential game, there are feasible choices of distributed learning algorithms that are theoretically guaranteed to converge to equilibrium [60, 54]. This distributed fashion is preferred in the sense that the agents can make use of only local information to make the decision. The local information includes i) the agent's own information and ii) the information communicated from its neighbours. Therefore, each agent is exchanging the information with only its neighbours.

2.4. A Distributed Solution Method

In this section, we show how we fit the problem (2.3) into a state-based potential game, and provide a distributed learning algorithm such that the agents achieve the Nash equilibrium (i.e. optimum) following the learning rules.

2.4.1. Game Design

1) The agent set and the communication graph:

The agent set (including the regulator agent 0) is $\tilde{I} = \{0, 1, \dots, |I|\}$. The communication structure is described by an undirected graph $G = (\tilde{I}, E)$, where each node represents an agent and each edge $(i, j) \in E$ indicates that agents i and j can directly communicate. Accordingly, for each agent $i \in \tilde{I}$, its neighbor set is defined as

$$N_i \doteq \{ j \in \tilde{I} \mid (i, j) \in E \} \cup \{i\}.$$

That is, N_i contains all agents that can directly communicate with agent i, and also includes agent i itself. We assume that the communication graph G is connected.

2) States and state space:

We denote matrix e_i as the agent i's estimation for the coupling constraint (average share), i.e., $e_i \sim \frac{1}{|\tilde{I}|}(\sum_{j\in\tilde{I}} L_j(v_j) - F_0)$. The system states tuple $x = (v, e)$ is then composed of two parts, i) all the decision variables $v = (v_0, v_1, \cdots, v_n)$; and ii) the estimation from all the agents $e = (e_0, e_1, \cdots, e_n)$. The feasible space of the states x is

$$X = \{(v, e) \,|v_i \in V_i, e_i \prec 0\}. \tag{2.4}$$

By defining the estimation state e_i, we are taking the first step to decouple the coupling constraint, as all the agents now have an independent estimation state for this constraint.

3) Action tuples and the state-dependent action space:

The action tuple $u_i = (\hat{v}_i, \hat{e}_{i\to out})$ is defined for each agent. It is composed of two parts similar to the states, i) the changes to be made on the local decision variables $\hat{v}_i$, and ii) the set of respective estimation matrices to be delivered to the corresponding neighbors $\hat{e}_{i\to out} = \left\{ \hat{e}_{i\to j} \,|j \in N_i \right\}$. An intuitive interpretation of the *delivered estimation* $\hat{e}_{i\to j}$ is, the *amount of estimation* that agent i would spare from its own estimation to assist its neighbor agent j. Imagine now the estimation matrix e_i from agent i is largely negative definite, which means all the eigenvalues of e_i are relatively far below 0. In this case, it would be helpful if agent i can assist some other agent j by supporting them with a negative definite estimation matrix $\hat{e}_{i\to j}$. With this support, agent j can be *more bold* when it adjusts v_j to get higher utility, while still maintaining its estimation e_j being negative definite.

The action's feasible space is state-dependent, as it needs to ensure the ensuing states lie in the feasible space (2.4). Given a system state x, the action's feasible set/region is

$$U_i(x) = \left\{ (\hat{v}_i, \hat{e}_{i\to out}) \middle| \begin{array}{c} v_i + \hat{v}_i \in V_i, \\ \hat{e}_{i\to out} \preceq 0, \\ e_i + L_i(\hat{v}_i) - \hat{e}_{i\to out} \prec 0 \end{array} \right\}. \tag{2.5}$$

4) State transition function:

We denote the local dynamic for agent i as $\tilde{x} = f_i(x, u_i)$. It is as follows with $\tilde{v}_i$, $\tilde{e}_i$ and $\tilde{e}_{out}$ being the ensuing states

$$\begin{cases} \tilde{v}_i = v_i + \hat{v}_i \\ \tilde{e}_i = e_i + L_i(\hat{v}_i) - \sum\limits_{out\in N_i} \hat{e}_{i\to out} + \sum\limits_{in\in N_i} \hat{e}_{in\to i} \\ \tilde{e}_{out} = e_{out} + \hat{e}_{i\to out}, \ \forall out \in N_i. \end{cases} \tag{2.6}$$

The transition function on v_i follows that $\hat{v}_i$ is the changes in the decision variables. The transition function on e contains contributions from three parts: i) $L_i(\hat{v}_i)$ be the estimation change caused by the change in v_i; ii) $-\sum \hat{e}_{i\to out}$ is the sum of estimation delivered to all neighbors; iii) $\sum \hat{e}_{in\to i}$ is the sum of estimation received from all neighbors.

The design of the state transition function provides a further step towards the decoupling of the coupling constraint. As we can see, during the dynamic transition, we always have $\sum e_i(t) = \sum e_i(0) + \sum L_i(v_i(t) - v_i(0))$. Therefore, if we nicely select the initial estimations as $\sum e_i(0) = \sum L_i(v_i(0)) - F_0$, we will have $\sum e_i(t) = \sum L_i(v_i(t)) - F_0$ for all t in the later iterations. This means that we can now evaluate the coupling constraint with the sum of the agents' estimation states. The feasible space of the states (2.4) ensures that all the estimations from agents are negative definite, and thus the coupling constraint is strictly satisfied.

5) State-dependent cost function:

The agent's cost functions depend on the ensuing system state $\tilde{x}$, which is deterministically determined by $x \in X$ and $u \in \prod_{i\in\tilde{I}} U_i(x)$ as a result of (2.6),

$$J_i(\tilde{x} \mid x, u_i, u_{-i}) = g_i(\tilde{v}_i) - \mu \sum_{j\in N_i} \log\det(-\tilde{e}_j). \tag{2.7}$$

The additional terms added to the original decoupled local cost function are logarithmic barrier functions for positive definite matrix. These additional terms implicitly enforce the constraints that $\tilde{e}_j \prec 0$ for every $j \in N_i$, which is satisfied naturally following (2.4) and (2.5).

6) Specialization for regulator agent 0:

For completeness, we now spell out the constraint sets and utility of the regulator (agent 0), which corresponds to $L(\mathcal{Q}, \gamma) := L_0(v_0)$. The local state tuple of agent 0 is $x_0 = (v_0, e_0)$ with $v_0 = (\mathcal{Q}, \gamma)$. Its feasible state space is

$$X_0 = \left\{ (\mathcal{Q}, \gamma, e_0) \,\middle|\, \mathcal{Q} \succeq 0,\ \gamma \geq 0,\ e_0 \prec 0 \right\}.$$

The action tuple of agent 0 is

$$u_0 = (\hat{\mathcal{Q}}, \hat{\gamma}, \hat{e}_{0\to\text{out}}), \quad \hat{e}_{0\to\text{out}} = \left\{ \hat{e}_{0\to j} \,\middle|\, j \in N_0 \right\}.$$

Given a state x, its state-dependent feasible action set is

$$U_0(x) = \left\{ (\hat{Q}, \hat{\gamma}, \hat{e}_{0\to\text{out}}) \;\middle|\; \begin{array}{c} Q + \hat{Q} \succeq 0, \\ \gamma + \hat{\gamma} \geq 0, \\ \hat{e}_{0\to\text{out}} \preceq 0, \\ e_0 + L(Q, \gamma) - \hat{e}_{0\to\text{out}} \prec 0 \end{array} \right\}. \tag{2.8}$$

The local state transition for agent 0, denoted $\tilde{x} = f_0(x, u_0)$, is

$$\begin{cases} \tilde{Q} = Q + \hat{Q}, \\ \tilde{\gamma} = \gamma + \hat{\gamma}, \\ \tilde{e}_0 = e_0 + L_0(\hat{Q}, \hat{\gamma}) - \sum_{out\in N_0} \hat{e}_{0\to out} + \sum_{in\in N_0} \hat{e}_{in\to 0}, \\ \tilde{e}_{out} = e_{out} + \hat{e}_{0\to out}, \quad \forall\, out \in N_0. \end{cases} \tag{2.9}$$

Finally, the state-dependent cost function of agent 0 is

$$J_0(\tilde{x} \mid x, u_0, u_{-0}) = w\,\tilde{\gamma} - \mu \sum_{j\in N_0} \log\det(-\tilde{e}_j). \tag{2.10}$$

This is consistent with the general form (2.7) with $g_0(v_0) = w\gamma$.

Theorem 2.1 (Optimality of the Nash equilibrium)**.** Consider a game $\mathcal{G}$ defined with (2.4)-(2.7). We have that, as μ goes to 0, the stationary Nash equilibrium of $\mathcal{G}$ converges to the solution of the original optimization problem (2.3).

Proof. The proof is built on three preparatory Lemmas.

Lemma 2.1.1. The defined state-based game with the agents' utility function (2.7) is a potential game with the potential function

$$\Phi(x, u) = \sum_{i\in\tilde{I}} g_i(\tilde{v}_i) - \mu \sum_{i\in\tilde{I}} \log\det(-\tilde{e}_i).$$

Lemma 2.1.2. The state-action pair $[x^*, u^*] = [(v^*, e^*), (\hat{v}^*, \hat{e}^*)]$ is a stationary state Nash equilibrium of the game $\mathcal{G}$, if and only if:

C-1 The decision variables v^* in states is the optimum of the following optimization problem with only uncoupled local constraints:

$$\begin{aligned} \min_{v} \; & \sum_{i\in\tilde{I}} g_i(v_i) - \left|\tilde{I}\right| \mu \log\det(F_0 - \sum_{i\in\tilde{I}} L_i(v_i)) \\ \text{s.t.} \quad & v_i \in V_i \qquad \forall i \in \tilde{I}. \end{aligned} \tag{2.11}$$

C-2 The estimation variables e^* in states satisfy that for all $i \in \tilde{I}$,

$$e_i^* = \frac{1}{|\tilde{I}|}\left(\sum_{i\in\tilde{I}} L_i\left(v_i^*\right) - F_0\right).$$

C-3 For all $i \in \tilde{I}$, $\hat{v}_i{}^* = 0$ and $\sum \hat{e}^*_{i\leftarrow in} - \sum \hat{e}^*_{i\rightarrow out} = 0$.

Lemma 2.1.3. The optimum of the optimization problem (2.11) converges to the solution of the original optimization problem (2.3) as μ goes to 0.

Lemma 2.1.1 identifies the potential function of the designed game and therefore proves the existence of the stationary Nash Equilibrium. Lemma 2.1.2 identifies the equivalence between the equilibrium and v^* the optimal solution of (2.11). Lemma 2.1.3 shows the convergence of v^* to the solution of (2.3) as μ goes to 0, and therefore completes the proof of Theorem 2.1.

Lemma 2.1.1 is easy to verify. The proofs of Lemma 2.1.2 and Lemma 2.1.3 are provided in Appendices 2.7.1-2.7.2. □

2.4.2. A Distributed Learning Algorithm

Proposition 2.1 (Better response dynamics [60])**.** In a state-based potential game, if the myopic agents all make attempts to improve their own utility function based on their local information, and take turns to update their action pairs, it is guaranteed that the Nash equilibrium defined in Lemma 2.1.2 is attained when no more agents can improve their utility.

Based on Proposition 2.1, we here provide a distributed learning algorithm based on the Curnot adjustment. Each agent follows (2.12) and takes a turn to update its localized state-action pair

$$\begin{aligned} u_i(k) &= \arg\min_{u_i\in U_i(x)} J_i(\tilde{x} \mid x(k), u_i, 0), \\ x(k+1) &= f_i(x(k), u_i(k)). \end{aligned} \tag{2.12}$$

The agent assumes others are taking action 0 ($u_{-i} = 0$) when making the decisions, which means the information needed in (2.12) is all locally available by agent i. The learning process is continued until every agent confirms that it does not have the will to change its actions, as schematically illustrated in Figure 2.1. Note that the initial states are chosen in the feasible set (2.4) satisfying $\sum e_i(0) = \sum L_i(v_i(0)) - F_0$.

Remark. We can see that in Proposition 2.1, there are no strict requirements for the order of updating during the learning process. Therefore, the learning process is naturally

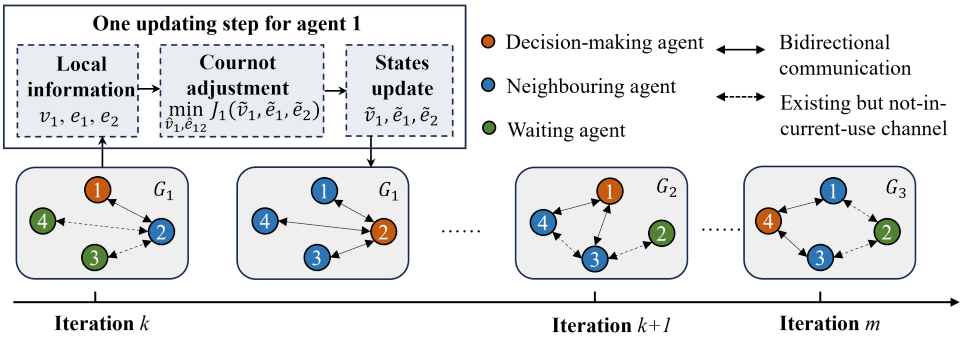

Figure 2.1: Schematic of the proposed game-theoretic learning approach. During the iterative learning process, the communication graph may vary over time; this is illustrated by the graph instances G_1, G_2, and G_3, which denote different communication topologies at different iterations.

robust to changes of agent update orders, or even occasional missing out. Furthermore, the learning rules are based on current local information, with no specifications on the communication graph G as long as it is connected. Therefore, the whole learning process is robust to communication loss or changes.

We note that, although the proposed game-theoretic learning approach shares some conceptual similarities with classical primal–dual distributed optimization methods for problems with coupling constraints, it differs markedly in several important aspects. A detailed comparison between the two is provided in Appendix 2.7.3.

2.5. Case Study

We consider a system with 3 aggregators, offering their FFR services at two standardized latency levels, $K = 2$: $\tau_1 = 0.5s$ and $\tau_2 = 5s$. The cost function related to each aggregator's service provision, and their corresponding local limits are:

1. $g_1(v_1) = 0.2H_1^1 + 2D_1^1; \quad 0 \leq H_1^1, D_1^1 \leq 5$
2. $g_2(v_2) = 0.3H_2^2 + 3D_2^1 + D_2^2; \quad 0 \leq H_2^2, D_2^1, D_2^2 \leq 5$
3. $g_3(v_3) = 1.5D_3^2; \quad 0 \leq D_3^2 \leq 10.$

The stability-economic weight w is set as 5. The setting for the system FFR target $T(s)$ is $H = 15$ p.u., $D = 20$ p.u. and a latency $\tau = 3.2$ s.

2.5.1. Convergence of the Game-Theoretic Learning Approach

Fig. 2.2 presents our results obtained under $\mu = 0.02$, and under a varying communication graph. The communication graph G encountered sudden changes twice ($G_1 \rightarrow G_2 \rightarrow G_3$), at iteration 1000 and iteration 2000. Fig. 2.2 (a)-(d) demonstrate that the approach is indeed robust to the changes happening in the communication graph during the learning. The decision variables converged to the solution of (2.11), which is close to but not exactly equal to the optimum solution of (2.3) (as shown in dashed lines). The errors between are due to the non-zero hyperparameter μ that steers the solution slightly away from the exact optimum, as a price for decoupling the constraint. μ could be adjusted lower to get results closer to the optimum, while it will deteriorate the convergence rate as the agents are *less driven* by the instructions coming from coupling constraints. Fig. 2.2 (e) illustrates the decrease of the potential during the learning process. Fig. 2.2 (f) shows the coupling constraint is strictly satisfied during the whole learning process. This implies that the approach is safe to be applied online, as when the agents are adjusting their decisions, all the constraints are guaranteed to be satisfied. The convergence of the algorithm can take a large number of iterations, showing the algorithm's limitation on the computational time. In the FFR coordination problem, though there may not be a very strict requirement on the computational time: the aggregators implement the output of the algorithm by programming their devices ahead of real-time operation; the FFR response is thereafter automatically provided in real-time when the system contingency occurs. When the number of agents is large, to address the time limitations, a solution set for multiple common FFR targets could be prepared in advance. Subsequently, when one FFR target is identified in operation, the solution pre-computed for a "close" target can be initialized for aggregators to accelerate the convergence to the optimum. Beyond the test cases considered above, in practical operation, system conditions such as DER availability, renewable generation, and ancillary service requirements may vary over time due to intermittency and forecast updates. In such cases, the coordination problem can be periodically updated using the latest system information and solved again using the proposed distributed approach. As the proposed approach is iterative, the solution obtained at the previous scheduling step can naturally serve as a warm start for the next optimization round. Since operating conditions typically evolve gradually between successive intervals, the new optimal operating point is often close to the previous one, allowing the distributed learning process to converge rapidly without restarting from scratch. This interpretation enables the framework to efficiently track time-varying operating conditions while preserving the distributed structure of the algorithm. A full implementation of such online optimization

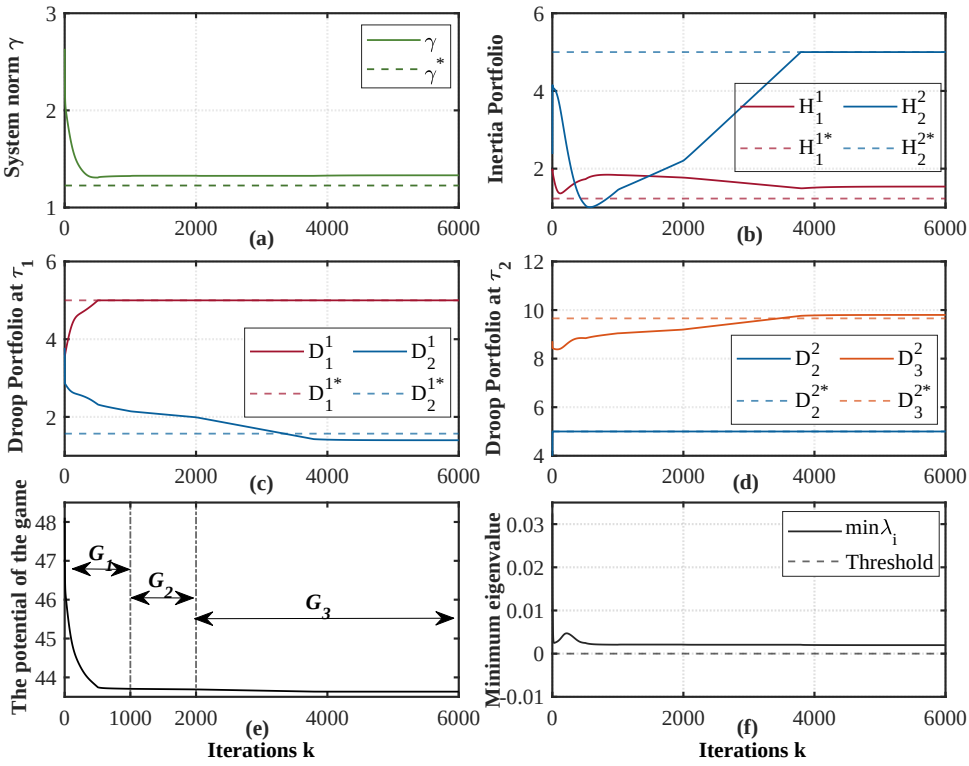

Figure 2.2: Results during the game-theoretic learning process. In (a)-(d), the color represents the decision variables coming from different agents (green-0, red-1, blue-2, orange-3).

with continuously updated measurements and forecasts is left for future work.

2.5.2. DVPP Regulation Performance

The previous subsection focused on the convergence properties of the proposed game-theoretic learning algorithm. We now complement this analysis with a time-domain validation, demonstrating that the proposed $\mathcal{H}_\infty$ formulation indeed yields a DVPP whose FFR closely tracks the desired target response under realistic frequency deviations. Specifically, we use the converged equilibrium as the DVPP FFR provision, resulting in the aggregate transfer function

$$T_{\text{prop}}(s) = \sum_i T_i(s).$$

For comparison, we construct a baseline portfolio that matches the desired aggregate inertia and damping in a purely static sense, using the cheapest possible combination of

devices. This yields a baseline aggregate transfer function $T_{\text{base}}(s)$ with $H_{11} = 5$ p.u., $H_{22} = 5$ p.u., $D_{11} = 5$ p.u., $D_{21} = 0$ p.u., $D_{22} = 5$ p.u., and $D_{32} = 10$ p.u. We excite $T_{\text{prop}}(s)$, $T_{\text{base}}(s)$, and the target $T(s)$ with the same frequency-deviation signal $\Delta f(t)$. We choose $\Delta f(t)$ as a mildly oscillatory and exponentially damped trace,

$$\Delta f(t) = -Ae^{-t/T_d} \sin(2\pi f_0 t),$$

with parameters $A = 0.5$ Hz, $T_d = 6$ s, and $f_0 = 0.25$ Hz. The corresponding target and realized FFR powers,

$$P_{\text{tar}}(t) = (T * \Delta f)(t),$$
$$P_{\text{prop}}(t) = (T_{\text{prop}} * \Delta f)(t),$$
$$P_{\text{base}}(t) = (T_{\text{base}} * \Delta f)(t),$$

are computed via time-domain simulation using MATLAB `lsim`, and the tracking errors are defined as

$$e_{\text{prop}}(t) = P_{\text{prop}}(t) - P_{\text{tar}}(t), \qquad e_{\text{base}}(t) = P_{\text{base}}(t) - P_{\text{tar}}(t).$$

Figure 2.3 illustrates the time-domain FFR tracking performance of the proposed formulation and the baseline portfolio. In the upper panel, the response of the proposed DVPP remains consistently close to the target trajectory over the entire transient. Small deviations are still visible, especially around the nadir and the early recovery phase, but their magnitude is limited. By contrast, the baseline response departs more noticeably from the target, with a clearly larger mismatch during the initial swing and a less accurate recovery. The lower panel, showing the tracking errors, quantifies this difference more clearly. The error associated with the proposed DVPP remains relatively small throughout the event, while the baseline error exhibits substantially larger peaks. In particular, the maximum absolute error of the baseline is approximately four times that of the proposed design. This confirms that the $\mathcal{H}_\infty$-based formulation significantly improves the dynamic regulation performance of the DVPP, whereas a least-cost static allocation that only matches aggregate inertia and damping is insufficient to ensure accurate FFR tracking, because it neglects the different latencies of the individual provisions and their impact on the aggregate response.

2.6. Conclusion

In this chapter, an efficient and tractable approach for the DER aggregators inside a DVPP to learn their optimal dynamic ancillary service provisions is proposed, based on $\mathcal{H}_\infty$ sys-

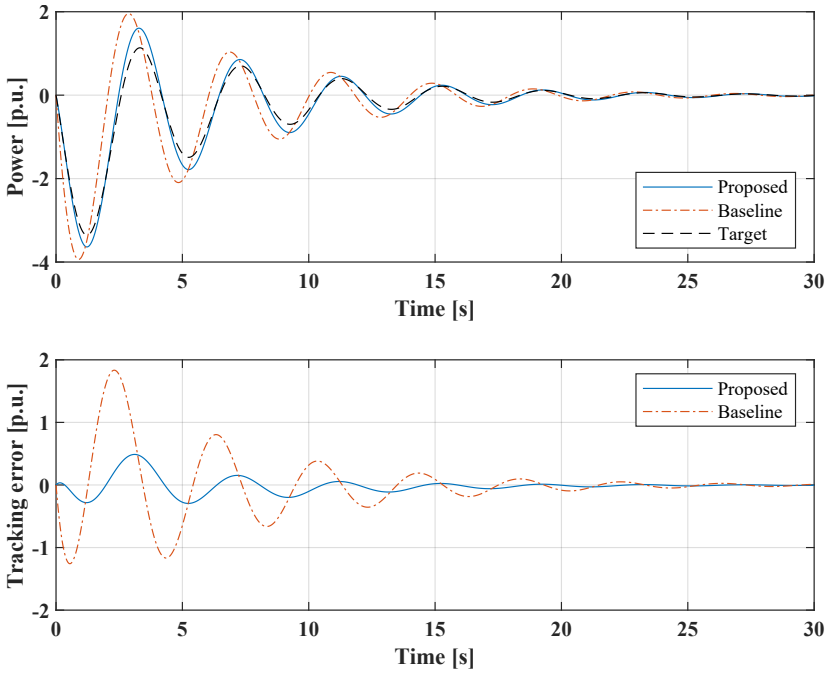

Figure 2.3: Time-domain FFR tracking performance of the proposed $\mathcal{H}_\infty$-based DVPP coordination compared with the baseline portfolio. Top: realized FFR powers $P_{\text{prop}}(t)$ and $P_{\text{base}}(t)$ versus the target trajectory $P_{\text{tar}}(t)$ under the same mildly oscillatory, exponentially damped frequency deviation $\Delta f(t)$. Bottom: corresponding tracking errors $e_{\text{prop}}(t)$ and $e_{\text{base}}(t)$, illustrating the substantially reduced peak error and overall mismatch achieved by the proposed formulation relative to the least-cost static allocation that only matches aggregate inertia and damping.

tem norm and a proposed game-theoretic learning algorithm. Although the method is brought up and introduced in a specific setting, it could be directly applied to other distributed use cases with coupling semidefinite constraints. The designed approach works as a distributed protocol. Future research can investigate this protocol for power system FFR market mechanism design, such that the agents are well incentivized and voluntarily participate in the game. It is also recommended to develop approaches that further boost the convergence time for future research.

2.7. Appendix

2.7.1. Proof of Lemma 1.2

(⇒) We first prove that if $[x^*, u^*]$ is a stationary state Nash equilibrium, then it satisfies [C-1]- [C-3]. Based on the second condition in definition 3, we have $x^* = f(x^*, u^*)$. It is only possible if [C-3] is satisfied. Then, based on the first condition in definition 3, we have $J_i(x^*, u_i^*, u_{-i}^*) = \min_{u_i \in U_i(x^*)} J_i(x^*, u_i, u_{-i}^*)$ for $\forall i \in I$. Since $J_i(x^*, u_i, u_{-i}^*)$ is a convex function on the tuple $u_i = (\hat{v}_i, \hat{e}_{i\to out}) \in U_i(x^*)$, we have the following equalities based on the sufficient and necessary optimality condition of convex optimization:

$$\operatorname{tr}\left(\frac{\partial J_i(x^*, u_i, u_{-i}^*)}{\partial \hat{e}_{i\to j}}\bigg|_{u_i=u_i^*}^{T} (\hat{e}_{i\to j} - \hat{e}_{i\to j}^*) \right) \geq 0, \quad \forall i \in I,\ \forall j \in N_i,\ \forall \hat{e}_{i\to j} \preceq 0 \tag{2.13}$$

$$\left[\frac{\partial g_i}{\partial \hat{v}_{i,d}}\bigg|_{\tilde{v}^*} + \mu \operatorname{tr}\left((\tilde{e}_i^*)^{-1} \frac{\partial L_i(v_i)}{\partial \hat{v}_{i,d}} \right) \right] (\hat{v}_{i,d} - \hat{v}_{i,d}^*) \geq 0, \forall i \in I,\ \forall \hat{v}_{i,d} \in \hat{v}_i,\ \forall \hat{v}_i \in U_i(x^*). \tag{2.14}$$

Eq. (2.13) implies that $\tilde{e}_j^* - \tilde{e}_i^* = 0, \forall i \in I, j \in N_i$. Because the communication graph is connected, we then have $\tilde{e}_j^* - \tilde{e}_i^* = 0,\ \forall i, j \in I$. Since $\sum_i \tilde{e}_i^* = \sum_{i\in I} L_i\left(\tilde{v}_i^*\right) - F_0$, we thus have

$$\tilde{e}_i^* = \frac{1}{|I|}\Big(\sum_{i\in I} L_i\left(\tilde{v}_i^*\right) - F_0\Big), \forall i \in I \qquad ([\text{C-2}] \text{ proved}). \tag{2.15}$$

Substituting (2.15) into (2.14), we obtain, for any $i \in I$, any $\hat{v}_i \in U_i(x^*)$, and any component $\hat{v}_{i,d} \in \hat{v}_i$,

$$\left[\frac{\partial g_i}{\partial \hat{v}_{i,d}}\bigg|_{\tilde{v}^*} + |I|\mu \operatorname{tr}\left(\Big(\sum L_i\left(\tilde{v}_i^*\right) - F_0\Big)^{-1} \frac{\partial L_i(v_i)}{\partial \hat{v}_{i,d}} \right) \right] \cdot (\hat{v}_{i,d} - \hat{v}_{i,d}^*) \geq 0, \tag{2.16}$$

which is exactly the sufficient optimality condition for problem [C-1].

(⇐) First, if u^* satisfies [C-3], then $x^* = f(x^*, u^*)$, x^* is a fixed point of the state transition function with u^*. In the second step, we could easily verify (2.13) from (2.15) [C-2]. Then we substitute (2.15) [C-2] into (2.16) [C-3], we get (2.14). This completes the optimality condition of

$$u_i^* \in \arg\min_{u_i \in U_i(x^*)} J_i(x^*, u_i, u_{-i}^*), \quad \forall i \in I.$$

2.7.2. Proof of Lemma 1.3

Let S^n denote the set of $n \times n$ real symmetric matrices. We use

$$S_+^n := \{X \in S^n \mid X \succeq 0\}$$

to denote the cone of symmetric positive semidefinite matrices, and

$$S^n_{++} := \{X \in S^n \mid X \succ 0\}$$

to denote the cone of symmetric positive definite matrices.

The logarithmic barrier function for the positive definite cone is defined for $X \in S^n_{++}$ as

$$B(X) := -\log\det(X) = -\ln\prod_{i=1}^{n}\lambda_i(X) = -\sum_{i=1}^{n}\ln\lambda_i(X).$$

The barrier function $B(X)$ exerts a repelling effect with respect to the boundary of the positive semidefinite cone, denoted by ∂S^n_+, where

$$\partial S^n_+ = \{X \in S^n \mid \lambda_i(X) \geq 0,\ i = 1, \dots, n,\ \text{and } \lambda_i(X) = 0 \text{ for at least one } i\}.$$

As X approaches the boundary ∂S^n_+, the barrier value $B(X)$ tends to $+\infty$.
Now consider the optimization problem

$$\begin{aligned} \min_{v}\ & \sum_{i\in I} g_i(v_i) - |I|\,\mu\log\det(F_0 - \sum_{i\in I} L_i(v_i)) \\ \text{s.t.}\quad & v_i \in V_i \qquad \forall i \in I. \end{aligned}$$

As μ gets to 0, the influence of the barrier function reduces and the solution is pushing closer to (but never can achieve) the boundary. Therefore, the problem converges to

$$\begin{aligned} \min_{v}\ & \sum_{i\in I} g_i(v_i) \\ \text{s.t.}\quad & v_i \in V_i \qquad \forall i \in I \\ & \sum_{i\in I} L_i(v_i) \prec F_0, \end{aligned}$$

which is a strict version of our original problem.

2.7.3. Comparison with Dual-Based Distributed Optimization

The proposed game-theoretic learning scheme differs from dual-based distributed optimization methods in several important aspects. To make this comparison precise, we revisit the convex formulation in (2.3). A standard dual approach introduces a global dual matrix variable $\Lambda \succeq 0$ associated with the constraint (2.3c). The dual function is

$$d(\Lambda) = \min_{v} \sum_{i\in I} g_i(v_i) + \operatorname{tr}\left(\Lambda\left(\sum_{i\in I} L_i(v_i) - F_0\right)\right) \quad \text{s.t.} \quad v_i \in V_i,\ \forall i \in I,$$

and the corresponding dual problem reads

$$\max_{\Lambda \succeq 0} \; d(\Lambda).$$

2

This dual problem decomposes across agents as

$$\max_{\Lambda \succeq 0} \; -\mathrm{tr}(\Lambda F_0) + \sum_{i \in I} \min_{v_i \in V_i} \big[g_i(v_i) + \mathrm{tr}(\Lambda L_i(v_i)) \big]. \tag{2.17}$$

Gradient-type primal-dual schemes (e.g., [61]) typically update Λ and v_i iteratively, for instance:

$$\Lambda^{k+1} = \Big[\Lambda^k + \delta_k \big(\textstyle\sum_{i \in I} L_i(v_i^k) - F_0 \big)^T \Big]_{\succeq 0}, \tag{2.18a}$$

$$v_i^{k+1} = \Big[v_i^k - \delta_k \Big(\nabla g_i(v_i^k) + \nabla_{v_i} \mathrm{tr}(\Lambda^{k+1} L_i(v_i^k)) \Big) \Big]_{V_i}, \quad \forall i \in I, \tag{2.18b}$$

with step size $\delta_k > 0$ and suitable projections. Here, Λ plays the role of a global "shadow price" penalizing each agent's contribution to the violation of the coupling constraint.

The proposed state-based game shares this high-level idea: penalizing the agents' impact on the coupling constraint, but implements it via *local* estimation states e_i instead of a single global dual variable:

- In dual-based schemes such as (2.18), Λ is a global variable, updated using the aggregated quantity $\sum_{i \in I} L_i(v_i^k)$ and then implicitly "seen" by all agents through the term $\mathrm{tr}(\Lambda L_i(v_i))$. This typically requires a central step (or a global consensus step) to compute and disseminate Λ^{k+1}.
- In the game-theoretic formulation, each agent i maintains its own estimation matrix e_i, representing a local estimate of its share of the coupling constraint. The state transitions (2.6) (and their specialization for agent 0) are designed so that

  $$\sum_{i \in \tilde{I}} e_i(t) = \sum_{i \in \tilde{I}} L_i(v_i(t)) - F_0$$

  holds along the trajectories, while each $e_i(t)$ remains negative definite. Thus, the global LMI is enforced via the sum of local states, and the coupling is handled through neighbour-to-neighbour exchanges of estimation matrices.

This leads to three key differences compared to classical dual-based distributed methods: (i) the semidefinite constraint is fully decoupled into local conditions $e_i \prec 0$, enforced at each iteration through the logarithmic barrier terms in (2.7); (ii) the communication

pattern is purely graph-based and fully distributed, without a central dual update; and (iii) constraint satisfaction is guaranteed throughout the entire learning process, not only at convergence, which is particularly attractive for online adjustment of FFR provisions under time-varying operating conditions and communication graphs.

3

Scaling FFR to Large DER Populations via Decomposition

Building on the previous chapter, this chapter continues to investigate the operation of a Dynamic Virtual Power Plant (DVPP) delivering Fast Frequency Response (FFR) to the system operator. It addresses the scalability of DVPP coordination as the portfolio size grows. The proposed framework has two main goals: (i) maintain high regulation quality by ensuring that the aggregate FFR tracks the DVPP target, and (ii) achieve economic efficiency by keeping the total provision cost low. To this end, the DVPP regulation is formulated as a frequency-domain matching problem, and a distributed solution based on dual decomposition is developed. The proposed approach decomposes the original large-scale nonconvex problem into multiple smaller nonconvex subproblems, substantially improving tractability and computational efficiency. Theoretical analysis and case studies show that the solution quality, in terms of optimality, remains high and can even improve as the problem size increases.

This chapter is based on 🗎 ***H. Xie***, *S. Khodakaramzadeh, P. Mohajerin Esfahani, and J. L. Cremer, "Enveloping Large Populations of Device Dynamics for Fast Frequency Response in a Dynamic Virtual Power Plant," submitted.*

3.1. Introduction

The integration of renewables poses significant challenges to power system operation [62]. This causes a pressing need for ancillary services such as inertia emulation, frequency regulation, voltage regulation and power smoothing, to maintain the secure operation in the presence of disturbances [63]. Distributed Energy Resources (DERs), while widely deployed, remain largely underutilized in this context. This underutilization stems from their geographic dispersion across distribution networks and the relatively limited flexibility of individual units, which renders them unsuitable for direct participation in transmission-level operations [64]. To harness the latent flexibility of DERs at scale, the concept of Virtual Power Plants (VPPs) has emerged as a promising solution [65]. One VPP functions as an aggregating envelope that encapsulates numerous DERs and coordinates their operations to deliver ancillary services to the grid collectively [47]. Within this framework, the VPP serves as an intermediary, ensuring that individual device constraints and objectives are respected while meeting the needs of the Transmission System Operator (TSO).

This chapter focuses on operating a VPP composed of a large number of heterogeneous DERs to respond to TSO requests for Fast Frequency Response (FFR) services. Such a VPP falls under the specific category of a Dynamic Virtual Power Plant (DVPP) [16], given that FFR represents a critical form of the dynamic ancillary services. Unlike static ancillary services such as steady-state generation–demand balancing and voltage regulation, FFR is inherently dynamic. It refers to the active-power response delivered within seconds after a disturbance to slow down and limit rapid frequency deviations. The purpose of FFR is to enhance frequency dynamics over the entire transient period [41]. To be effective, FFR must be activated automatically and almost instantaneously upon detection of frequency deviations, ensuring real-time active-power injection [44]. In practical implementations, inverter-based DERs are typically configured to provide FFR through pre-programmed control algorithms. More detailed introduction on FFR are documented in [44, 45, 66].

The key objectives in operating such a DVPP providing FFR services are twofold: i) to ensure that the aggregate frequency response of the DERs closely tracks the requested dynamic behavior specified by the TSO, and ii) to achieve the overall efficiency, typically represented by minimizing the sum of the DERs' cost functions. These objectives are preferably achieved in a distributed setting that supports data confidentiality, such that DERs are not required to reveal sensitive data, including cost functions and local device operational limits [67].

Extensive research has focused on DVPP operation, particularly on the obstacles of aggregating the DERs' FFR provisions to match TSO requests. Two main approaches have

emerged to address this problem. The mainstream approach follows the strategy of first representing FFR as a set of equivalent adjustable parameters, including synthetic inertia, damping coefficients, and primary active power responses [32, 43, 46]. Intuitively, the DVPP's aggregate FFR is represented as the summation of these equivalent parameters across DERs. The operation problem is then formulated as an algebraic matching task between these aggregate quantities and the TSO's specified requirements. However, these equivalent-parameter models collapse the underlying device dynamics into static coefficients and, by construction, ignore heterogeneous response latencies. In practice, DER activation times differ by up to several seconds due to their control bandwidths and communication delays [47, 20]. When such latency is neglected, the DVPP may be scheduled as if it can deliver the contracted FFR capacity within the required time window, while in reality a significant portion of the response arrives too late. This mismatch can lead to frequency nadirs that violate grid-code limits, undermine the certifiability and trustworthiness of DVPP-based FFR products, and force operators to apply conservative security margins or restrict DVPP participation. Thus, the simplification not merely reduces DVPP accuracy or efficiency but also creates a structural barrier to the reliable and large-scale deployment of DVPPs as FFR providers.

In contrast, the second line of investigation acknowledges the dynamic nature of FFR and explicitly incorporates latency effects [47, 41, 42, 35]. These studies advocate modeling FFR provision as a dynamic system, with frequency deviations as input and active power as output. This modeling approach not only yields a more accurate representation of DER behavior but also greatly facilitates the enforcement of grid code requirements and the practical implementation of converter control algorithms. To this end, rational parametric transfer functions in the frequency domain are applied to describe the dynamic behavior of DERs providing FFR, consistent with advanced inverter control paradigms [49, 50, 51]. Under these modeling assumptions, several studies have approached the coordination of DERs to achieve DVPP operational objectives. Notably, divide-and-conquer strategies have been explored in [20, 42], wherein the aggregate control task is decomposed among individual devices. However, these approaches fall short in achieving overall economic efficiency: they generally pre-assign all contracted DERs into the regulation task and rely heavily on high-speed resources such as supercapacitors to provide the necessary flexibility. This reliance requires over-dimensioning of these fast-storage devices and leads to an increased cost of maintaining sufficient FFR capability for the DVPP. In contrast, the work [41] proposes a sub-optimal $\mathcal{H}_\infty$ control framework, casting the DVPP operation as a convex semi-definite programming problem. While this approach demonstrates improved

dynamic performance and theoretical soundness, it faces scalability limitations. As the number of DERs increases, the associated dimension of coupling constraints grows substantially, posing challenges for computation efficiency and practical implementation at a large scale. Furthermore, all the aforementioned methods assume that DERs operate with fixed latency parameters. However, some DER technologies such as batteries and other energy storage systems offer the capability to adjust their response latency on demand within predefined limits [68]. As a result, existing approaches are unable to accommodate DERs with variable latency and consequently fail to exploit the latent flexibility embedded in heterogeneous DER portfolios.

To address the aforementioned challenges, this chapter proposes a novel approach for the distributed operation of large-scale DVPPs in delivering FFR services. The proposed framework targets two major bottlenecks in current DVPP operation: (i) the lack of a principled formulation that directly matches an FFR target with heterogeneous DER device dynamics, and (ii) the underexploited temporal flexibility arising from devices with adjustable response latencies. The main contributions of this chapter are summarized as follows:

1. **Frequency-domain dynamic matching:** The DVPP regulation problem is interpreted as a dynamic matching task between two systems: the desired FFR behavior (given) and the aggregate of all DER FFR provisions (adjustable). This matching task is formulated in the frequency domain by enforcing agreement between the corresponding transfer functions at multiple sampling frequencies, thereby capturing the essential dynamic characteristics of the FFR service.

2. **Optimization-based DVPP operation framework:** Building on the frequency domain matching formulation, a nonlinear nonconvex optimization framework that links device-level design variables to the DVPP performance and operational costs is constructed in this chapter. This formulation naturally accommodates devices with variable response latencies by treating latency parameters as decision variables, hence explicitly leveraging the additional temporal flexibility inherent in DERs.

3. **Scalable distributed algorithm:** To address the computational intractability of the resulting nonconvex program, a dual decomposition–based solution method that decomposes the large-scale problem into tractable per-device subproblems is developed. This enables distributed computation while yielding high-quality primal solutions and preserving performance guarantees as the DVPP scales.

The rest of the chapter is organized as follows. In Section 3.2, the DVPP operation problem is formulated as a dynamic matching in the frequency domain. The model is also extended to enable the treatment of latency as an optimization variable. Section 3.3 elaborates on the solution methodology, introducing the dual decomposition technique to solve the large-scale DVPP operation problem. Section 3.4 presents the case studies to demonstrate the effectiveness and scalability of the proposed approach. Finally, section 3.5 concludes the chapter and outlines directions for future research.

3.2. Problem Formulation

In this section, we formulate the problem of operating a DVPP to provide FFR, by aligning the aggregate device dynamics with the regulation target. The following subsections break down the problem into key components: regulation quality, problem structure, and hyperparameter selection.

3.2.1. Regulation Quality of One DVPP

We consider a set of inverter-based DERs that collectively provide FFR and form a DVPP. Let $I = \{1, 2, \dots, |I|\}$ index these devices. For each DER $i \in I$, we represent its FFR provision in the Laplace domain as

$$T_i(s) := \frac{\Delta P_i(s)}{\Delta f(s)} = \frac{H_i s + D_i}{\tau_i s + 1}, \tag{3.1}$$

where $s \in \mathbb{C}$ is the Laplace variable in the complex domain. Here, $\Delta P_i(s)$ denotes the Laplace transform of the active-power deviation of DER i from its baseline output P_i due to FFR provision, and $\Delta f(s)$ is the Laplace transform of the frequency deviation from the nominal value. The sign convention adopted here differs from that in Chapter 2: in the present chapter, the counteracting nature of FFR is absorbed into the definition of the deviation variables, so the transfer function is written without an explicit negative sign. Therefore, the formulation remains physically consistent with Chapter 2, where the same stabilising effect is expressed by including the negative sign explicitly in the transfer function. The variables H_i and D_i are the tunable controller parameters subject to device-specific bounds. They are commonly referred to as virtual inertia and virtual damping, respectively, as they emulate the inertia and damping effects of synchronous generators. As DERs exhibit heterogeneous device bandwidths, their response time differ when providing FFR services. The parameter τ_i in (3.1) precisely captures this characteristic, and represents the FFR delivery latency of the DER i. For many inverter-based DERs, the latency τ_i can be adjusted within an admissible interval, e.g., via the allocation of communication

bandwidth resources [69, 70, 64]. In other cases, the latency of the DER's FFR provision is fixed and unadjustable, due to the device constraints [20], communication limits [71] or user preferences. Our formulation accommodates both settings by treating τ_i either as a decision variable or as a fixed parameter.

Fixed-τ_i devices Let $I_f \subseteq I$, with $|I_f| \leq |I|$, denote the subset of DERs whose FFR provisions have fixed response time τ_i. For each device $i \in I_f$, the stacked decision variable vector contains $x_i \doteq (H_i, D_i, P_i)$, where H_i and D_i are its virtual inertia and virtual damping as introduced before, and P_i denotes the baseline active power output of DER i, i.e., its scheduled injection without any additional modulation due to FFR provision. With the dimension of decision variable vector denoted by $n_i = 3$, $x_i \in \mathbb{R}^{n_i}$ are continuously adjustable within the local constraints set $X_i \subseteq \mathbb{R}^{n_i}$. The local constraints set X_i contains

$$\begin{cases} P_i + \Delta\overline{P_i} \leq \overline{P_i}, \\ H_i \geq 0, D_i \geq 0, P_i \geq 0. \end{cases} \tag{3.2}$$

In (3.2), $\Delta\overline{P_i}$ denotes the maximum active power that is needed during the FFR provision time horizon $\mathcal{T}$, i.e.,

$$\Delta\overline{P_i} := \max_{t\in\mathcal{T}} \Delta P_i(t).$$

$\overline{P_i}$ is the maximum available power of DER i. One can estimate its maximum $\Delta\overline{P_i}$ by [46]

$$\Delta\overline{P_i} = \Delta\dot{f}_{\text{lim}} H_i + \Delta f_{\text{lim}} D_i, \tag{3.3}$$

where $\Delta\dot{f}_{\text{lim}}$ is the maximum permissible Rate of Change of Frequency (RoCoF) specified by the system operator. Similarly, Δf_{lim} is the maximum permissible frequency drop specified by the system operator, which is commonly known as frequency nadir. Combining (3.2) and (3.3), the feasible set X_i is constructed as a collection of linear constraints on the decision variables x_i.

Variable-τ_i devices Let $I_v \subseteq I$, with $|I_f| + |I_v| = |I|$, denote the subset of DER devices that have variable-τ_i in FFR provisions. For each device $i \in I_v$, the decision variable vector contains an extra variable τ_i, i.e., $x_i \in \mathbb{R}^{n_i} \doteq (H_i, D_i, P_i, \tau_i)$, and $n_i = 4$. The feasible set X_i includes (3.2) and (3.3), and an extra box constraints on τ_i, i.e.,

$$\underline{\tau_i} \leq \tau_i \leq \overline{\tau_i}, \tag{3.4}$$

where $\underline{\tau_i}$ is the minimum latency of device i, and $\overline{\tau_i}$ is the maximum latency of device i.

Aggregating the device-level contributions yields the DVPP FFR transfer function

$$T_{\text{DVPP}}(s) \doteq \sum_{i \in I} T_i(s) = \sum_{i \in I} \frac{H_i s + D_i}{\tau_i s + 1}. \tag{3.5}$$

Given a target FFR transfer function specified by the TSO as

$$\hat{T}(s) = \frac{\hat{H} s + \hat{D}}{\hat{\tau} s + 1}, \tag{3.6}$$

where the equivalent inertia $\hat{H}$, damping $\hat{D}$, and the latency $\hat{\tau}$, are treated as fixed constants. The objective of the regulation of DVPP is to align the aggregate devices' FFR with the target FFR, i.e.,

$$T_{\text{DVPP}}(s) \to \hat{T}(s).$$

Remark (Target FFR)**.** The target FFR $\hat{T}$ is assumed to be derived by the system operator under applicable grid codes, accounting for system operating conditions, disturbance magnitudes, and closed-loop stability requirements. Looking ahead, as ancillary-service markets evolve, one DVPP may obtain its target FFR directly from market-clearing outcomes.

To assess the DVPP regulation quality, we now try to quantify the regulation error between the two dynamic systems, $\sum_{i \in I} T_i(s)$ and $\hat{T}(s)$. The idea is to i) derive each system's frequency-domain response, i.e., the Fourier transform of its impulse response, ii) sample these responses over a prescribed set of angular frequencies and compare them pointwise. Ideally, if the two frequency responses coincide over the entire bandwidth of interest, the two dynamic systems can be regarded as equivalent. This indicates that the DVPP exactly achieves its FFR regulation objective.

Let us denote the sampled set of angular frequencies as $\Omega = \{\omega_k\}_{k=1}^{K}$, where K is the dimension of the set. By substituting $s = j\omega_k$ in the transfer function $T_i(s)$, we obtain the corresponding frequency response

$$T_i(j\omega_k) = \frac{j\omega_k H_i + D_i}{1 + j\omega_k \tau_i} = \underbrace{\frac{\omega_k^2 \tau_i H_i + D_i}{1 + \omega_k^2 \tau_i^2}}_{\Re(T_i(j\omega_k))} + j \underbrace{\frac{\omega_k H_i - \omega_k \tau_i D_i}{1 + \omega_k^2 \tau_i^2}}_{\Im(T_i(j\omega_k))},$$

where $j = \sqrt{-1}$ is the imaginary unit, $\Re(T_i(j\omega_k))$ denotes the real part of the complex function, and $\Im(T_i(j\omega_k))$ denotes the imaginary part. Similarly, we derive the frequency

response of the target dynamic system as the complex function

$$\hat{T}(j\omega_k) = \frac{j\omega_k\hat{H} + \hat{D}}{1 + j\omega_k\hat{\tau}} = \underbrace{\frac{\omega_k^2\hat{\tau}\hat{H} + \hat{D}}{1 + \omega_k^2\hat{\tau}^2}}_{\Re(\hat{T}(j\omega_k))} + j\underbrace{\frac{\omega_k\hat{H} - \omega_k\hat{\tau}\hat{D}}{1 + \omega_k^2\hat{\tau}^2}}_{\Im(\hat{T}(j\omega_k))}.$$

Now we can quantify the regulation error of the two FFR systems, by computing their differences in real and imaginary parts separately. Introduce the stacked decision vector $\boldsymbol{x} \doteq (x_i)_{i\in I}$. Define the real and imaginary mismatches as functions of ω_k and $\boldsymbol{x}$:

$$\begin{aligned}
\Delta\Re(\omega_k, \boldsymbol{x}) &\doteq \sum_{i\in I} \Re(T_i(j\omega_k)) - \Re(\hat{T}(j\omega_k)) \\
&= \sum_{i\in I} \frac{\omega_k^2\tau_i H_i + D_i}{1 + \omega_k^2\tau_i^2} - \frac{\omega_k^2\hat{\tau}\hat{H} + \hat{D}}{1 + \omega_k^2\hat{\tau}^2}, \\
\Delta\Im(\omega_k, \boldsymbol{x}) &\doteq \sum_{i\in I} \Im(T_i(j\omega_k)) - \Im(\hat{T}(j\omega_k)) \\
&= \sum_{i\in I} \frac{\omega_k H_i - \omega_k\tau_i D_i}{1 + \omega_k^2\tau_i^2} - \frac{\omega_k\hat{H} - \omega_k\hat{\tau}\hat{D}}{1 + \omega_k^2\hat{\tau}^2}.
\end{aligned} \tag{3.7}$$

Given a fixed $\boldsymbol{x}$, we quantify the regulation performance by per-ω_k slack variables $\boldsymbol{\varepsilon} \doteq (\varepsilon_k)_{k=1}^K$ via

$$\begin{aligned}
\min_{\boldsymbol{\varepsilon}\geq 0} \quad & \|\boldsymbol{\varepsilon}\|_2^2 \\
\text{s.t.} \quad & |\Delta\Re(\omega_k, \boldsymbol{x})| \leq \varepsilon_k \quad 1 \leq k \leq K \\
& |\Delta\Im(\omega_k, \boldsymbol{x})| \leq \varepsilon_k \quad 1 \leq k \leq K.
\end{aligned} \tag{3.8}$$

where ε_k captures the worst-component mismatch at ω_k, and $\|\boldsymbol{\varepsilon}\|_2$ represents the vector's Euclidean norm, $|\cdot|$ denotes the absolute value of a scalar. As the objective $\|\boldsymbol{\varepsilon}\|_2^2$ approaches zero, each ε_k vanishes, which forces both the real and imaginary mismatches to zero at every sampled frequency. Hence, the two frequency responses coincide on the grid Ω. When Ω is chosen to be sufficiently dense over the bandwidth of interest, this provides an accurate certification that the DVPP dynamics align with the target FFR. Conversely, larger values of ε_k indicate greater spectral discrepancy and thus poorer regulation performance. Therefore, the proposed frequency domain error metric offers an efficient assessment of the DVPP regulation quality.

Remark (Zero- and high-frequency limits)**.** If $\Omega = \{0\}$, then $T_i(j0) = D_i$ and $\hat{T}(j0) = \hat{D}$. The matching reduces to the static damping condition $\sum_{i\in I} D_i = \hat{D}$. Conversely, in the high-frequency limit $\omega \to \infty$, $T_i(j\omega) \to H_i/\tau_i$, and $\hat{T}(j\omega) \to \hat{H}/\hat{\tau}$, yielding $\sum_{i\in I} H_i/\tau_i =$

$\hat{H}/\hat{\tau}$; under homogeneous fixed latencies $\tau_i \equiv \hat{\tau}$, this simplifies to $\sum_{i\in I} H_i = \hat{H}$. Hence, the static matching formulations in the literature can be regarded as the special cases of our broader frequency-domain framework.

3.2.2. The DVPP Operation Problem

Now consider the settings above and given a target FFR transfer function (3.6), the DVPP operation problem is to adjust the decision variables of all devices $\boldsymbol{x}$, such that i) the DVPP achieves a good regulation quality. ii) the overall costs attached to the DERs' variables are minimized. We formulate the problem as follows:

$$\min_{\boldsymbol{x},\boldsymbol{\varepsilon}\geq 0} \quad \|\boldsymbol{\varepsilon}\|_2^2 + w\sum_{i\in I} f_i(x_i) \tag{3.9a}$$

$$\text{s.t.} \quad x_i \in X_i \quad \forall i \in I \tag{3.9b}$$

$$|\Delta\Re(\omega_k, \boldsymbol{x})| \leq \varepsilon_k \quad 1 \leq k \leq K \tag{3.9c}$$

$$|\Delta\Im(\omega_k, \boldsymbol{x})| \leq \varepsilon_k \quad 1 \leq k \leq K \tag{3.9d}$$

$$P_{\min} \leq \sum_{i\in I} \pi_P(x_i) \leq P_{\max}. \tag{3.9e}$$

Here, $f_i : \mathbb{R}^{n_i} \to \mathbb{R}$ represents a convex device-level cost, the weight $w > 0$ is a predefined parameter tuning the trade-off between regulation quality and economic efficiency. Constraints (3.9b) enforce the local feasibility of each DER device, with $X_i \subseteq \mathbb{R}^{n_i}$ being a convex and compact set defined by (3.2)-(3.4). The real/imaginary frequency-domain mismatches at each ω_k are bounded in (3.9c) and (3.9d). Constraint (3.9e) bounds the aggregate baseline active power of the DVPP. Recall that x_i stacks the decision variables of device i, and the projector $\pi_P : \mathbb{R}^{n_i} \to \mathbb{R}$ selects the active power variable, i.e., $\pi_P(x_i) = P_i$. The limits $P_{\min}$ and $P_{\max}$ are specified by the system operator based on network-operational constraints, e.g., transmission/ feeder congestion, transformer thermal ratings, voltage/ security margins, or market dispatch setpoints. Together with the device-level headroom conditions (3.2) and (3.3), this coupling constraint ensures that the scheduled baselines P_i are simultaneously feasible at the device and system levels while reserving sufficient headroom for FFR activation. The proposed formulation provides a modular and extensible framework that can incorporate a range of DVPP services: additional operational requirements (e.g., ramp-rate limits, energy budgets, or market setpoint tracking) can be incorporated without requiring significant changes to the core problem.

We would like to highlight several structural features of the formulated problem (3.9):

Scaling: The problem size grows with the number of DER devices $|I|$. The decision vector dimension is $\sum_{i\in I} n_i$, and the constraints consist of all local sets X_i. However, the number of coupling constraints in (3.9c)–(3.9e) is independent of $|I|$; instead, it depends on the number of frequency samples $|\Omega| = K$. Specifically, there are $4K$ FFR regulation constraints (due to the real and imaginary parts of the absolute value) and 2 aggregate power constraints. This structure effectively overcomes the scalability limitations of previous formulations, such as that in [41], where the dimension of the coupling constraints grows explosively with the number of DER devices $|I|$.

Coupling: The objective $\sum_{i\in I} f_i(x_i)$ and the local feasibility sets are fully separable across devices (in x_i). The coupling across the DER devices arises only through the coupling constraints (3.9c)- (3.9e) and ε_k. As made explicit in (3.7), each coupling constraint is *additively separable across devices*: it is a sum of functions that depend on a single block x_i, namely $\Re\{T_i(j\omega_k)\}$, $\Im\{T_i(j\omega_k)\}$, and $\pi_P(x_i)$. Consequently, the constraint Jacobian has a *block-angular* pattern: block-diagonal local parts plus a small number ($4K+2$) of global "sum" rows. While this structure prevents a straightforward primal decomposition into independent subproblems, it is well-suited for dual decomposition, which forms the foundation of our proposed solution method.

Non-Convexity: For each DER $i \in I$ and each $\omega_k \in \Omega$, let us define the per-device mappings

$$\alpha_i^k(H_i,\tau_i) := \frac{\omega_k^2\tau_i H_i}{1+\omega_k^2\tau_i^2}, \quad \beta_i^k(D_i,\tau_i) := \frac{D_i}{1+\omega_k^2\tau_i^2},$$
$$\gamma_i^k(H_i,\tau_i) := \frac{\omega_k H_i}{1+\omega_k^2\tau_i^2}, \quad \delta_i^k(D_i,\tau_i) := -\frac{\omega_k\tau_i D_i}{1+\omega_k^2\tau_i^2}.$$

With these definitions, and easing the notations as $\hat{R}^k \doteq \Re(\hat{T}(j\omega_k))$, $\hat{J}^k \doteq \Im(\hat{T}(j\omega_k))$, (3.9c) and (3.9d) can be equivalently written as *sum-separable* rows:

$$\begin{aligned}\left|\sum_{i\in I}(\alpha_i^k+\beta_i^k)-\hat{R}^k\right| &\le \varepsilon_k, \forall 1\le k\le K\\ \left|\sum_{i\in I}(\gamma_i^k+\delta_i^k)-\hat{J}^k\right| &\le \varepsilon_k, \forall 1\le k\le K.\end{aligned} \tag{3.10}$$

We can see from above that, for the fixed-τ_i devices $i \in I_f$, the mappings α_i^k, β_i^k, γ_i^k, δ_i^k are all affine in (H_i, D_i). In the extreme case of $I_f = I$, (3.10) are linear in x_i, thus the overall problem (3.9) is convex. However, for the variable-τ_i devices $i \in I_v$, the mappings

Algorithm 1: Iterative frequency selection

Inputs : $K_{\max} \in \mathbb{N}$, tolerance $\varepsilon_{\text{tol}} > 0$, candidate grid Ξ, initial seed ω_1

Outputs: Selected angular-frequency set Ω

$\Omega \leftarrow \{\omega_1\}, \quad t \leftarrow 1$

repeat

  Solve (3.9) on Ω to obtain $\boldsymbol{x}^{(t)}$, $T^{(t)}_{\text{DVPP}}$

  $\omega^\star \in \arg\max_{\omega \in \Xi \setminus \Omega} \Delta T(\omega; \boldsymbol{x}^{(t)}); \quad \Delta T^\star \leftarrow \Delta T(\omega^\star; \boldsymbol{x}^{(t)})$

  if $\Delta T^\star \le \varepsilon_{\text{tol}}$ ***or*** $|\Omega| = K_{\max}$ **then**

    break

  $\Omega \leftarrow \Omega \cup \{\omega^\star\}; \quad t \leftarrow t + 1$

until $\Delta T^\star \le \varepsilon_{\text{tol}}$ ***or*** $|\Omega| = K_{\max}$

return Ω

above become *rational* in (H_i, D_i, τ_i) with a common quadratic denominator $1 + \omega_k^2 \tau_i^2$. This renders (3.10) and the problem (3.9) *nonconvex*.

To summarize, we confront a large-scale, smooth but nonconvex program whose objective is device-separable and whose inter-device interactions arise only through additively separable coupling constraints (3.9c)–(3.9e). In general, attaining global optima for nonconvex programs is intractable (typically NP-hard), and centralized nonlinear solvers scale poorly as the variables' dimension $\sum_{i \in I} n_i$ increases. These difficulties motivate us to propose the solution method (Sec. 3.3) that exploits the problem structure to handle the coupling constraints.

3.2.3. Selection of Angular Frequencies Ω

The set of angular frequencies $\Omega = \{\omega_k\}_{k=1}^K$ can either be chosen based on the frequencies of interest (e.g., centered around critical resonance frequencies) or empirically determined. To further optimize the selection of frequencies, we adopt an iterative scheme as illustrated in Algorithm 1. Let Ξ be a candidate set. Given $\boldsymbol{x}$, define the pointwise mismatch

$$\Delta T(\omega, \boldsymbol{x}) \doteq \sqrt{\Delta \Re(\omega, \boldsymbol{x})^2 + \Delta \Im(\omega, \boldsymbol{x})^2}.$$

This iterative rule concentrates angular frequency samples where the current DVPP response deviates most from the target, yielding a compact Ω that can subsequently be reused as a reference for similar operating conditions or new instances.

3.3. Solution Method

In this section, we present a dual-decomposition–based solution method for (3.9). The developed method decomposes the large-scale nonconvex program into a number of smaller nonconvex subproblems, yielding tractable and scalable computation with high-quality primal solutions, in terms of optimality. We also derive a computable performance bound (a duality-gap certificate) for the obtained solutions, which quantifies suboptimality and demonstrates the scalability of the proposed method.

3.3.1. Dual–Decomposition Solution

We solve (3.9) by dualizing the coupling constraints (3.9c)–(3.9e). Recall there are $4K + 2$ coupling constraints, we attach each constraint with one multiplier:

$$\begin{aligned}
\lambda_k^{\mathrm{R},+} &\leftrightarrow \sum_{i\in I}(\alpha_i^k + \beta_i^k) - \varepsilon_k \le \hat{R}^k, 1 \le k \le K \\
\lambda_k^{\mathrm{R},-} &\leftrightarrow -\sum_{i\in I}(\alpha_i^k + \beta_i^k) - \varepsilon_k \le -\hat{R}^k, 1 \le k \le K \\
\lambda_k^{\mathrm{J},+} &\leftrightarrow \sum_{i\in I}(\gamma_i^k + \delta_i^k) - \varepsilon_k \le \hat{J}^k, 1 \le k \le K \\
\lambda_k^{\mathrm{J},-} &\leftrightarrow -\sum_{i\in I}(\gamma_i^k + \delta_i^k) - \varepsilon_k \le -\hat{J}^k, 1 \le k \le K \\
\mu^{+} &\leftrightarrow \sum_{i\in I}\pi_P(x_i) \le P_{\max} \\
\mu^{-} &\leftrightarrow -\sum_{i\in I}\pi_P(x_i) \le -P_{\min}.
\end{aligned} \tag{3.11}$$

To ease notations, bold symbols denote vectors stacked over the frequency grid Ω, e.g., $\boldsymbol{\lambda}^{\mathrm{R},+} \doteq (\lambda_1^{\mathrm{R},+}, \cdots, \lambda_K^{\mathrm{R},+})^\top$, $\boldsymbol{\alpha_i} \doteq (\alpha_i{}^1, \cdots, \alpha_i{}^K)^\top$. We write $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ for the full collections of multipliers. Let $\boldsymbol{r}(\boldsymbol{x}, \boldsymbol{\varepsilon}) : \mathbb{R}^{\sum_{i\in I} n_i + K} \to \mathbb{R}^{4K+2}$ denote the vector obtained by stacking the left–hand sides of (3.11), and let $\boldsymbol{v} \in \mathbb{R}^{4K+2}$ collect the corresponding right–hand-side constants. The dual function $d : \mathbb{R}^{4K+2} \to \mathbb{R}$ of problem (3.9) is defined as

$$d(\boldsymbol{\lambda}, \boldsymbol{\mu}) \doteq \min_{\boldsymbol{x}\in X, \boldsymbol{\varepsilon}\ge 0} \quad \mathcal{L}(\boldsymbol{x}, \boldsymbol{\varepsilon}, \boldsymbol{\lambda}, \boldsymbol{\mu}),$$

where $\mathcal{L} : \mathbb{R}^{\sum_{i\in I} n_i+5K+2} \to \mathbb{R}$ is the Lagrangian defined below:

$$
\begin{aligned}
&\mathcal{L}(\boldsymbol{x}, \boldsymbol{\varepsilon}; \boldsymbol{\lambda}, \boldsymbol{\mu}) = \|\boldsymbol{\varepsilon}\|_2^2 + w \sum_{i\in I} f_i(x_i) + (\boldsymbol{\lambda}^{\mathrm{R},+})^\top \Big[\sum_{i\in I} (\boldsymbol{\alpha}_i + \boldsymbol{\beta}_i) - \boldsymbol{\varepsilon} - \hat{\boldsymbol{R}} \Big] \\
&+ (\boldsymbol{\lambda}^{\mathrm{R},-})^\top \Big[\hat{\boldsymbol{R}} - \sum_{i\in I} (\boldsymbol{\alpha}_i + \boldsymbol{\beta}_i) - \boldsymbol{\varepsilon} \Big] + (\boldsymbol{\lambda}^{\mathrm{I},+})^\top \Big[\sum_{i\in I} (\boldsymbol{\gamma}_i + \boldsymbol{\delta}_i) - \boldsymbol{\varepsilon} - \hat{\boldsymbol{I}} \Big] \\
&+ (\boldsymbol{\lambda}^{\mathrm{I},-})^\top \Big[\hat{\boldsymbol{I}} - \sum_{i\in I} (\boldsymbol{\gamma}_i + \boldsymbol{\delta}_i) - \boldsymbol{\varepsilon} \Big] + \mu^+ \Big[\sum_{i\in I} \pi_P(x_i) - P_{\max} \Big] + \mu^- \Big[P_{\min} - \sum_{i\in I} \pi_P(x_i) \Big] \\
&= \underbrace{(\boldsymbol{\lambda}^{\mathrm{R},-} - \boldsymbol{\lambda}^{\mathrm{R},+})^\top \hat{\boldsymbol{R}} + (\boldsymbol{\lambda}^{\mathrm{I},-} - \boldsymbol{\lambda}^{\mathrm{I},+})^\top \hat{\boldsymbol{I}} - \mu^+ P_{\max} + \mu^- P_{\min}}_{l(\boldsymbol{\lambda},\boldsymbol{\mu})} \\
&+ \underbrace{\|\boldsymbol{\varepsilon}\|_2^2 - (\boldsymbol{\lambda}^{\mathrm{R},+} + \boldsymbol{\lambda}^{\mathrm{R},-} + \boldsymbol{\lambda}^{\mathrm{I},+} + \boldsymbol{\lambda}^{\mathrm{I},-})^\top \boldsymbol{\varepsilon}}_{h(\boldsymbol{\varepsilon};\boldsymbol{\lambda})} \\
&+ \sum_{i\in I} \underbrace{w f_i(x_i) + (\boldsymbol{\lambda}^{\mathrm{R},+} - \boldsymbol{\lambda}^{\mathrm{R},-})^\top (\boldsymbol{\alpha}_i + \boldsymbol{\beta}_i) + (\boldsymbol{\lambda}^{\mathrm{I},+} - \boldsymbol{\lambda}^{\mathrm{I},-})^\top (\boldsymbol{\gamma}_i + \boldsymbol{\delta}_i) + (\mu^+ - \mu^-)\pi_P(x_i)}_{g_i(x_i;\boldsymbol{\lambda},\boldsymbol{\mu})}.
\end{aligned}
\tag{3.12}
$$

In the last equality of (3.12), the blue group represents $h(\boldsymbol{\varepsilon}, \boldsymbol{\lambda})$ the terms with regard to the regulation errors, the red group represents $g_i(x_i, \boldsymbol{\lambda}, \boldsymbol{\mu})$ the terms with regard to each DER i, and the (black) remainder comprises only affine terms in the multipliers. Associated to the dual function d, we define the dual problem of (3.9) as

$$
\begin{aligned}
\max_{\boldsymbol{\lambda},\boldsymbol{\mu}} \quad & \left\{ l(\boldsymbol{\lambda}, \boldsymbol{\mu}) + \min_{\boldsymbol{\varepsilon}\geq 0} h(\boldsymbol{\varepsilon}, \boldsymbol{\lambda}) + \sum_{i\in I} \min_{x_i\in X_i} g_i(x_i, \boldsymbol{\lambda}, \boldsymbol{\mu}) \right\} \\
\text{s.t.} \quad & \boldsymbol{\lambda}, \boldsymbol{\mu} \geq 0.
\end{aligned}
\tag{3.13}
$$

We call the minimizations in the dual problem (3.13) as inner problems. We solve (3.13) with an iterative primal–dual scheme. In maximizing the dual, we also recover the inner minimizers $(\boldsymbol{x}^*, \boldsymbol{\varepsilon}^*)$ evaluated at the optimal multipliers $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$. These inner minimizers constitute the DVPP operating decisions of primary interest. The solution method is summarized in Algorithm. 2, with its key steps outlined below:

Dual subgradient ascent (outer loop) At iteration $t+1$, the multipliers are updated by projected subgradient ascent

$$
\begin{bmatrix} \boldsymbol{\lambda}^{t+1} \\ \boldsymbol{\mu}^{t+1} \end{bmatrix} = \left[\begin{bmatrix} \lambda^t \\ \mu^t \end{bmatrix} + \theta_t \begin{bmatrix} \frac{\partial \mathcal{L}(\boldsymbol{x}^t, \boldsymbol{\varepsilon}^t, \boldsymbol{\lambda}, \boldsymbol{\mu})}{\partial \boldsymbol{\lambda}} \\ \frac{\partial \mathcal{L}(\boldsymbol{x}^t, \boldsymbol{\varepsilon}^t, \boldsymbol{\lambda}, \boldsymbol{\mu})}{\partial \boldsymbol{\mu}} \end{bmatrix} \right]_{\geq 0}
\tag{3.14}
$$

where $(\boldsymbol{x}^t,\ \boldsymbol{\varepsilon}^t,\ \boldsymbol{\lambda}^t,\ \boldsymbol{\mu}^t)$ denote the variable values at iteration t, and $[\cdot]_{\geq 0}$ is the elementwise projection onto the nonnegative orthant, i.e., $[z]_{\geq 0} := (\max\{z_i, 0\})_i$ applied component-

wise. The step size follows a diminishing rule, $\theta_t = \theta_0/\sqrt{t}$, which promotes stable updates and convergence of the primal–dual iterates.

Inner minimizers At iteration $t+1$, we update the inner minimizers $(\boldsymbol{x}^{t+1}, \boldsymbol{\varepsilon}^{t+1})$ by solving the inner problems with the updated multipliers $(\boldsymbol{\lambda}^{t+1}, \boldsymbol{\mu}^{t+1})$, separately. These subproblems are decoupled and can be solved in parallel, enabling a distributed implementation. For the update of the regulation error, minimizing the inner problem over $\boldsymbol{\varepsilon}$ yields the closed form

$$\begin{aligned}\boldsymbol{\varepsilon}^{t+1} &= \arg\min_{\boldsymbol{\varepsilon}} h(\boldsymbol{\varepsilon}, \boldsymbol{\lambda}^{t+1}) \\ &= \frac{1}{2}(\boldsymbol{\lambda}^{t+1,\mathrm{R},+} + \boldsymbol{\lambda}^{t+1,\mathrm{R},-} + \boldsymbol{\lambda}^{t+1,\mathrm{I},+} + \boldsymbol{\lambda}^{t+1,\mathrm{I},-}).\end{aligned} \tag{3.15}$$

For the update of x_i^{t+1}, similarly follow minimizing the inner problem of device i

$$x_i^{t+1} = \arg\min_{x_i \in X_i} g_i(x_i, \boldsymbol{\lambda}^{t+1}, \boldsymbol{\mu}^{t+1}). \tag{3.16}$$

As discussed in (3.10), for variable-τ_i devices the mappings $\boldsymbol{\alpha}_i, \boldsymbol{\beta}_i, \boldsymbol{\gamma}_i, \boldsymbol{\delta}_i$ are nonlinear in x_i, so the subproblem (3.16) remains nonconvex. Nevertheless, (3.16) is per-device and thus of much smaller dimension than the full program (3.9), making it substantially easier to solve in practice. We propose two approaches to solve the subproblem (3.16).

(i) Traversal τ_i-grid search. With the box constraint $\underline{\tau_i} \le \tau_i \le \overline{\tau_i}$, choose a resolution $\Delta\tau > 0$ and form

$$\mathcal{T}_i := \left\{ \underline{\tau_i} + m\,\Delta\tau \;\middle|\; m = 0, \dots, \left\lfloor (\overline{\tau_i} - \underline{\tau_i})/\Delta\tau \right\rfloor \right\}.$$

For each fixed $\tau \in \mathcal{T}_i$, the subproblem (3.16) reduces to a small convex quadratic program in (H_i, D_i, P_i); solve it and record the objective $g_i(\tau)$. Select

$$\tau_i^\dagger \in \arg\min_{\tau \in \mathcal{T}_i} g_i(\tau), \qquad x_i^\dagger = (H_i^\dagger, D_i^\dagger, P_i^\dagger, \tau_i^\dagger).$$

As $\Delta\tau \to 0$, this grid-search approximation converges to the optimal value.

(ii) Projected first-order approach. Update the device solution to (3.16) via projected subgradient steps

$$x_i^{(j+1)} = \left[x_i^{(j)} - \theta_j \frac{\partial g_i\big(x_i^{(j)}, \boldsymbol{\lambda}, \boldsymbol{\mu}\big)}{\partial x_i} \right]_{X_i}, \tag{3.17}$$

where $\theta_j > 0$ is a stepsize and $[\cdot]_{X_i}$ denotes Euclidean projection onto the compact set X_i. The iteration stops upon convergence (e.g., a small relative change in

$x_i^{(j)}$). Because (3.16) is nonconvex, this approach is not guaranteed to find the global optimum and may converges to a stationary/local minimum. Compared with the τ_i-grid traversal (which can approach the global minimum as the grid is refined), the projected first-order method is generally more time-efficient in practice.

Stopping and feasibility certificate We terminate the iterations when (i) the relative multiplier change is small,

$$\frac{\left\|[\boldsymbol{\lambda}^{t+1};\, \boldsymbol{\mu}^{t+1}] - [\boldsymbol{\lambda}^{t};\, \boldsymbol{\mu}^{t}]\right\|_2}{\max\{1, \|[\boldsymbol{\lambda}^{t};\, \boldsymbol{\mu}^{t}]\|_2\}} \;\le\; \epsilon_{\text{tol}}^{1}, \tag{3.18}$$

and (ii) the coupling–constraint violation is below a feasibility tolerance,

$$\left\| \left[\, \boldsymbol{r}(\boldsymbol{x}^{t+1}, \boldsymbol{\varepsilon}^{t+1}) - \boldsymbol{v} \,\right]_{\ge 0} \right\|_\infty \;\le\; \epsilon_{\text{tol}}^{2}. \tag{3.19}$$

Although the stopping test includes a feasibility check, there is generally no guarantee that the last primal iterate $(\boldsymbol{x}^\star, \boldsymbol{\varepsilon}^\star)$ is feasible for the coupling constraints (3.9c) and (3.9d), due to nonconvexity. Related work [72] enforces resource contraction to ensure feasibility for the coupling constraints in MILPs. In our setting, we observe that $(\boldsymbol{x}^\star, \boldsymbol{\varepsilon}^\star)$ typically satisfies the aggregate power bounds (3.9e), while the FFR constraints (3.9c) and (3.9d) are more challenging. To ensure feasibility and a fair objective evaluation, we recompute the slacks at $\boldsymbol{x}^\star$ as

$$\tilde{\varepsilon}_k^\star = \max\Big\{ \big|\Delta\Re(\omega_k, \boldsymbol{x}^\star)\big|, \big|\Delta\Im(\omega_k, \boldsymbol{x}^\star)\big| \Big\}, \;\; \forall\, \omega_k \in \Omega. \tag{3.20}$$

We then report the objective and operational profiles using the reconstructed pair $(\boldsymbol{x}^\star, \tilde{\boldsymbol{\varepsilon}}^\star)$.

3.3.2. Scalability and Performance

Let $J_{\mathcal{P}}^\star$ denote the optimal value of (3.9), and let $J_{\mathcal{P}}(\boldsymbol{x}, \boldsymbol{\varepsilon})$ be the value attained by a candidate $(\boldsymbol{x}, \boldsymbol{\varepsilon})$. Likewise, we define $J_{\mathcal{D}}^\star$ and $J_{\mathcal{D}}(\boldsymbol{\lambda}, \boldsymbol{\mu})$ for the dual problem (3.13). Because (3.9) is nonconvex, in general $J_{\mathcal{D}}^\star \le J_{\mathcal{P}}^\star$ with a nonzero duality gap, and the primal pair $(\boldsymbol{x}^\star, \boldsymbol{\varepsilon}^\star)$ recovered from the dual might not be globally optimal. We next show that $(\boldsymbol{x}^\star, \boldsymbol{\varepsilon}^\star)$ enjoys a quantitative performance bound, and more importantly, the *relative* gap vanishes as the problem size grows. The forthcoming results build on the decomposition analysis in [72, Thm. 2.3], originally developed for mixed-integer linear programs.

Proposition 3.1 (Bound on duality gap)**.** Assume the slacks are uniformly bounded componentwise, i.e., $\|\boldsymbol{\varepsilon}\|_\infty \le \bar{\varepsilon}$. Then

$$J_{\mathcal{P}}^\star - J_{\mathcal{D}}^\star \;\le\; (4K+3) \cdot \max\Big\{ w \max_{i \in I} \gamma_i,\; K\, \bar{\varepsilon}^{\,2} \Big\}, \tag{3.21}$$

Algorithm 2: Dual–decomposition solver for (3.9)

Inputs : Feasible $\{x_i^0 \in X_i\}_{i\in I}$; $\boldsymbol{\varepsilon}^0 = \mathbf{0}$; $\boldsymbol{\lambda}^0 = \mathbf{0}$; $\boldsymbol{\mu}^0 = \mathbf{0}$; $t \leftarrow 0$

Outputs: $(\boldsymbol{x}^\star, \tilde{\boldsymbol{\varepsilon}}^\star)$ and $(\boldsymbol{\lambda}^\star, \boldsymbol{\mu}^\star)$

repeat
  Dual ascent: update $(\boldsymbol{\lambda}^{t+1}, \boldsymbol{\mu}^{t+1})$ via (3.14)
  Slack update: update $\boldsymbol{\varepsilon}^{t+1}$ with (3.15)
  foreach $i \in I$ **do**
    Device update: update x_i^{t+1} via (3.16)
  $t \leftarrow t + 1$
until *stopping tests* (3.18) *and* (3.19) *are satisfied*
Reconstruction: compute $(\boldsymbol{x}^\star, \tilde{\boldsymbol{\varepsilon}}^\star)$ from $(\boldsymbol{x}^t, \boldsymbol{\varepsilon}^t)$ using (3.20)
return $(\boldsymbol{x}^\star, \tilde{\boldsymbol{\varepsilon}}^\star)$, *and* $(\boldsymbol{\lambda}^\star, \boldsymbol{\mu}^\star) \leftarrow (\boldsymbol{\lambda}^t, \boldsymbol{\mu}^t)$

where

$$\gamma_i \doteq \max_{x_i \in X_i} f_i(x_i) - \min_{x_i \in X_i} f_i(x_i). \tag{3.22}$$

As a consequence of Proposition 3.1, let $|I|$ increase while K remains constant and the sets $\{X_i\}_{i\in I}$ are uniformly bounded. Suppose $J_{\mathcal{P}}^\star$ grows linearly with $|I|$. Then

$$\frac{J_{\mathcal{P}}^\star - J_{\mathcal{D}}^\star}{J_{\mathcal{P}}^\star} \to 0 \qquad \text{as } |I| \to \infty, \tag{3.23}$$

which indicates that the duality gap vanishes in relative terms as the DVPP size increases, and the recovered solution becomes asymptotically near-optimal. In other words, as the problem size of (3.9) grows, it tends to closely approximate a convex program.

The proof of Proposition 3.1 with a more general problem setting is presented in [73, Prop. 5.7.4, p. 223].

3.4. Case Study

In this section, we present three studies on a modified distribution test system. Study I (Sec. 3.4.1) evaluates how the proposed latency-aware dynamic matching achieves accurate tracking of the FFR target and efficient DVPP regulation, thereby demonstrating contribution 1. Study II (Sec. 3.4.2) quantifies the additional flexibility and operational benefits unlocked by allowing device latencies to vary, thereby demonstrating contribution 2. Study III (Sec. 3.4.3) assesses the efficacy, scalability, and runtime performance of

Table 3.1: Main Parameters of the Test System

System Parameters	Low DER penetration	High DER penetration
$\|I\|$	10	100
Total demand	3 MW	
Disturbance	0.32 MW	
Grid inertia	80 kW·s/Hz	48 kW·s/Hz
Grid damping	50 kW/Hz	18 kW/Hz
SG droop	1 Hz/MW	1.33 Hz/MW
SG delay	10 s	8 s
$\overline{P_i}$	[0; 500] kW	[0; 50] kW
$\Delta \dot{f}_{\text{lim}}$	1 Hz/s	
Δf_{lim}	0.8 Hz	
τ_i	[1; 5] s	
$f_i(x_i)$	$a_i P_i^2 + b_i H_i^2 + c_i D_i^2$	
a_i	[0; 1] k€/(MW)^2	
b_i	[1; 5] k€/(MW · s/Hz)^2	
c_i	[1; 5] k€/(MW/Hz)^2	
w	0.2 (MW/Hz)^2/k€	
Ω	$\{0, 0.25, 0.5, 0.75\}$ rad/s	

- SG: Existing Synchronous Generators in the grid.
- Values in the brackets are sampled from a uniform distribution.
- The default bounds of variable-latency DERs are $\underline{\tau}_i = 1$ and $\overline{\tau}_i = 5$, except in Study II, where this interval is varied in a sensitivity analysis.

the proposed distributed solution method to validate contribution 3. The parameters of the modified IEEE 33-bus distribution system are summarized below in Table. 3.1.

To demonstrate effectiveness, Study I evaluates two operating scenarios: low DER penetration with 10 DERs versus high DER penetration with 100 DERs. Study II focuses on the high-penetration case (100 DERs). Study III scales the DVPP from 10 to 1000 DERs

to stress-test the method. In the low DER penetration scenario, the FFR target is

$$\hat{T}(s) = \frac{0.6s + 0.5}{2.3s + 1}[\text{MW/Hz}].$$

In the high DER penetration scenario, the FFR target is

$$\hat{T}(s) = \frac{0.6s + 0.5}{1.3s + 1}[\text{MW/Hz}].$$

In both scenarios, the aggregate active power constraint follows $P_{\min} = 0.85\,\text{MW}$, and $P_{\max} = 1.95\,\text{MW}$. We compare the performance of our proposed solution method with the nonlinear solver `fmincon` in MATLAB R2025a [74]. All the computation was performed on a desktop PC with RAM of 64GB and a 3.70 GHz processor. The dynamic simulations in study I and II were performed in MATLAB Simulink.

3.4.1. Study I: Effectiveness of the DVPP Regulation

Figures 3.1 shows the post-disturbance frequency under two operating scenarios. We tested five cases per scenario:

- *Baseline*: static matching of inertia and damping;
- *Case 1*: dynamic matching with all τ_i fixed, solved by `fmincon`;
- *Case 2*: dynamic matching with all τ_i fixed, solved by the proposed method;
- *Case 3*: dynamic matching with all τ_i variable, solved by `fmincon`;
- *Case 4*: dynamic matching with all τ_i variable, solved by the proposed method.

As shown in the left panel (a) and (c) from both scenarios, the dynamic matching markedly improves the frequency nadir and satisfies the 49.2 Hz safety threshold, whereas static matching does not. Beyond nadir values, an additional observation is that dynamic matching also yields a visibly faster and smoother recovery trajectory, indicating that the aggregate response is shaped appropriately over the full transient rather than only at the most critical instant. This improvement is strongest under high DER penetration, where the baseline violates the threshold by approximately 0.201 Hz while all dynamic-matching cases restore a secure margin. Right panel (b) and (d) show the same mechanism at the power level: with dynamic matching, the DVPP's emulated FFR active power closely tracks the reference implied by the target transfer function throughout the event horizon, whereas the baseline exhibits a clear trajectory-shape mismatch (under-delivery in the

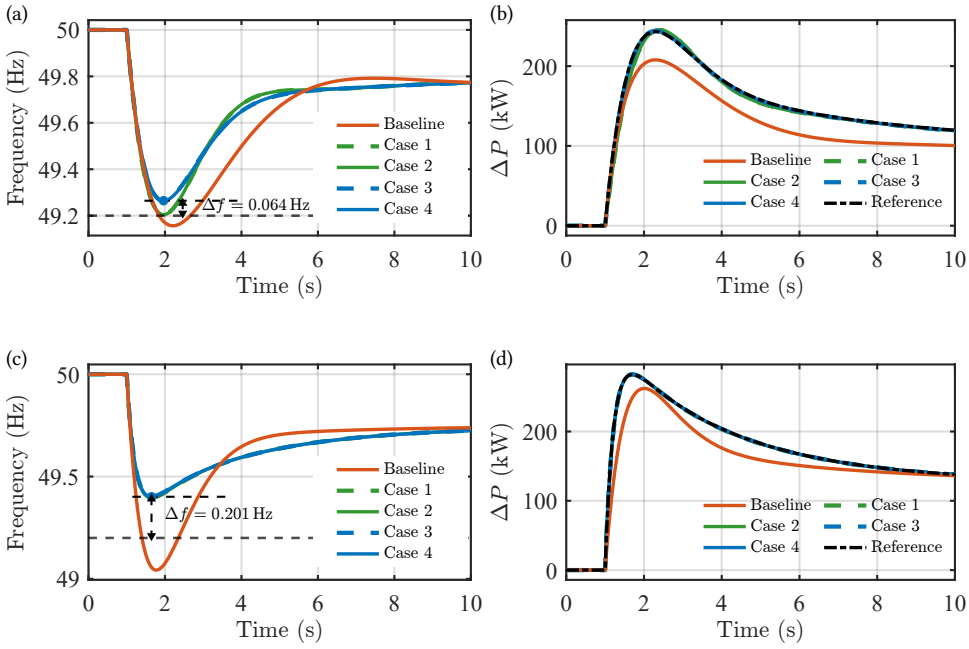

Figure 3.1: Comparison of DVPP regulation under low and high DER penetration: (a) low-penetration post-disturbance frequency; (b) low-penetration FFR power; (c) high-penetration post-disturbance frequency; (d) high-penetration FFR power.

early transient and/or overly rapid decay thereafter). The trajectories in Case 1 vs. Case 2 and Case 3 vs. Case 4 are nearly indistinguishable, indicating that the proposed solver achieves essentially the same solution quality as `fmincon` while enabling scalable computation. Finally, allowing variable τ_i does not sacrifice tracking quality in either scenario; its advantage is more pronounced in the low-DER case, where limited portfolio diversity reduces mutual compensation among devices. Under high penetration, aggregation already mitigates latency heterogeneity even with fixed τ_i, and the primary benefit of optimizing τ_i shifts from feasibility/tracking to reduced operating cost, as detailed in Study II.

We further examine how the choice of angular frequencies Ω affects DVPP regulation. Figures 3.2 compare using only the DC point $\Omega = \{0\}$ (i.e., $K = 1$) with a richer grid $|\Omega| = 10$ selected by the proposed iterative scheme. In figure 3.2(a) and figure 3.2(c), enlarging Ω substantially improves the alignment of both the Fourier-domain magnitude and the time-domain emulated FFR active power with the reference trajectory implied by

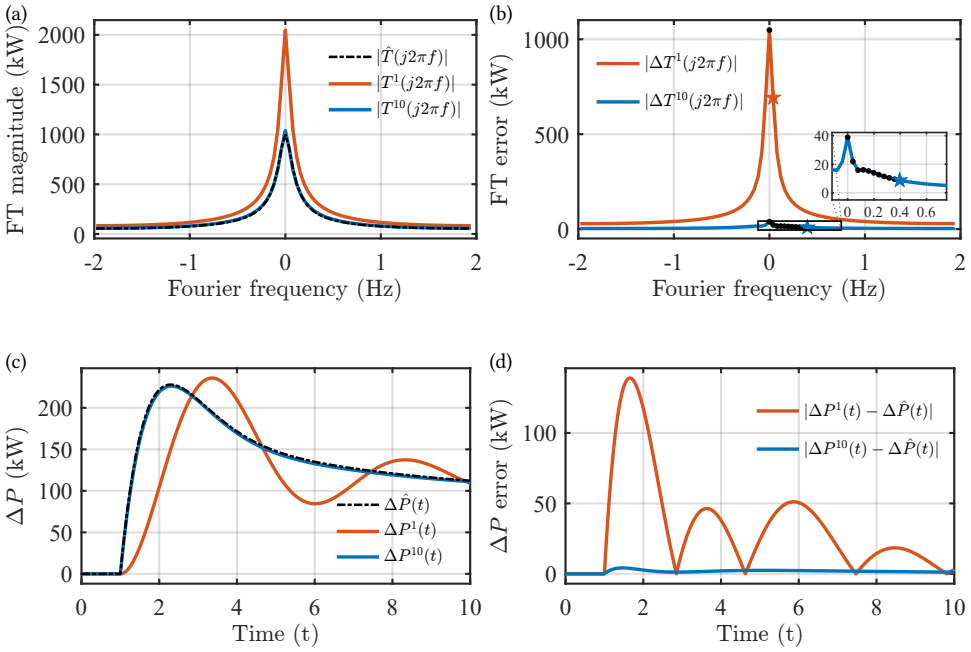

Figure 3.2: Comparison of frequency-domain and time-domain tracking under $|\Omega| = 1$ and $|\Omega| = 10$ (iterative): (a) Fourier-spectrum magnitude; (b) spectrum error $\Delta T(\omega)$; (c) time-domain FFR power; and (d) absolute tracking error. In (a) and (b), the horizontal axes denote the Fourier frequency f (Hz), displayed in the standard zero-centered two-sided spectrum after the conventional FFT shift.

the target FFR transfer function. By contrast, using a very small Ω (e.g., $|\Omega| = 1$) yields a visible shape mismatch, which is most pronounced during the early transient where accurate delivery is critical for arresting the frequency drop. Figure 3.2(b) and figure 3.2(d) plot the corresponding tracking errors, confirming that the improvement is not merely qualitative: the error energy is markedly reduced when $|\Omega| = 10$, indicating that the selected frequencies better constrain the dynamic profile rather than only the steady-state level. Figure 3.2(b) further illustrates how the iterative scheme works. At each iteration, we evaluate the magnitude mismatch $\Delta T(\omega)$ over a candidate set Ξ, choose the next frequency at the largest residual mismatch, $\omega^\star = \arg\max_{\omega\in\Xi\setminus\Omega} |\Delta T(\omega)|$ (denoted by star in the figure 3.2(b)), and append $\omega^\star$ to Ω (denoted by dots). This "greedy" selection concentrates new constraints where the current aggregate response deviates most, thereby accelerating mismatch reduction with a small number of frequencies.

Table 3.2: The norm of matching error $\boldsymbol{\varepsilon}$ under different Ω selection schemes

K	Iterative		Random					
	$\|\boldsymbol{\varepsilon}\|_1$	$\|\boldsymbol{\varepsilon}\|_2$	$\|\boldsymbol{\varepsilon}\|_1$			$\|\boldsymbol{\varepsilon}\|_2$		
			Min.	Avg.	Max.	Min.	Avg.	Max.
2	0.497	0.248	0.296	1.548	2.467	0.170	0.717	1.179
4	0.697	0.291	0.390	0.730	1.342	0.226	0.354	0.612
6	0.453	0.179	0.233	0.545	1.145	0.136	0.260	0.526
8	0.178	0.065	0.269	0.594	1.010	0.146	0.284	0.466
10	0.062	0.027	0.233	0.515	0.804	0.099	0.250	0.396

Table 3.2 compares the matching error $\boldsymbol{\varepsilon}$ for frequencies chosen by the *iterative* scheme (seeded at $\omega_1 = 0$) versus *random* selection. For the random baseline, we draw 10 independent instances per $|\Omega|$ and report min/avg/max. As $|\Omega|$ increases, both schemes generally reduce $\|\boldsymbol{\varepsilon}\|_1$ and $\|\boldsymbol{\varepsilon}\|_2$, reflecting progressively tighter dynamic matching. Notably, the iterative scheme achieves substantially smaller errors at the same K: for $K = 10$, it reduces $\|\boldsymbol{\varepsilon}\|_1$ to 0.062 and $\|\boldsymbol{\varepsilon}\|_2$ to 0.027, whereas the random averages remain at 0.515 and 0.250, respectively. Moreover, the random selection exhibits a wide spread (min/avg/max), especially at small K, indicating that naive frequency choices can be brittle and occasionally lead to large mismatches; the iterative rule mitigates this sensitivity by consistently targeting the most informative frequencies.

3.4.2. Study II: Benefits of Latency-Inclusive Operation

Figure 3.3 reports the DVPP overall operating cost (Eq. (3.9a)) as the fraction of DERs with *variable* latency increases. We consider two cases:

- Narrow adjustment interval: $2\,\text{s} \le \tau_i \le 4\,\text{s}$;
- Wide adjustment interval: $1\,\text{s} \le \tau_i \le 5\,\text{s}$.

Relative to the fixed-latency case (proportion of adjustable devices equals to 0), enabling latency tuning reduces the overall costs as more flexibilities are unlocked. The reduction strengthens as the proportion of adjustable devices increases. Moreover, the wide adjustment interval case consistently outperforms the narrow interval case, reflecting the added temporal flexibility and more effective frequency-domain matching.

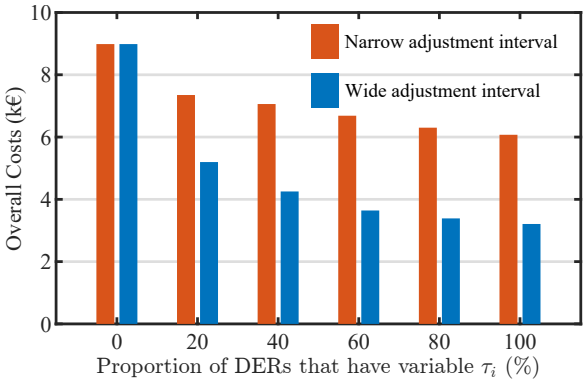

Figure 3.3: DVPP overall costs under different share of variable-τ_i DERs for two adjustment intervals.

3.4.3. Study III: Performance of the Solution Method

In this study, we focus on the settings where all DERs have variable latencies, thereby rendering the problem strongly nonconvex. For each problem size, 10 random problem instances are generated independently, and both the proposed method and the baseline `fmincon` solver are applied to each instance. The results for optimality gap and runtime reported in Fig. 3.5 are based on these 10 instances.

Convergence and feasibility. Figure 3.4(a) displays the evolution of the primal objective over the iterations for a randomly chosen instance, showing stable descent and convergence under a diminishing stepsize. Figure 3.4(b) monitors the aggregate-power coupling during the iterations for the same instance, which falls into the feasible interval despite dualizing the constraint. In practice, the linear structure of the power window makes it readily controlled by the multipliers, yielding primal feasibility at convergence.

Optimality certificate and gap. Given any primal-feasible pair $(\boldsymbol{x}, \boldsymbol{\varepsilon})$ and dual-feasible multipliers $(\boldsymbol{\lambda}, \boldsymbol{\mu})$, the weak duality relation holds:

$$J_D(\boldsymbol{\lambda}, \boldsymbol{\mu}) \;\le\; J_P^\star \;\le\; J_P(\boldsymbol{x}, \boldsymbol{\varepsilon}).$$

Hence, the optimal dual value $J_D^\star(\boldsymbol{\lambda}^\star, \boldsymbol{\mu}^\star)$ provides a computable tight lower bound for $J_P^\star$. We quantify the suboptimality of a candidate solution via the relative duality gap

$$\frac{J_P(\boldsymbol{x}, \boldsymbol{\varepsilon}) - J_D^\star}{J_D^\star}.$$

Figure 3.5(a) reports the resulting gaps for the proposed method and the baseline `fmincon` solver. The proposed distributed algorithm consistently achieves small duality gaps and

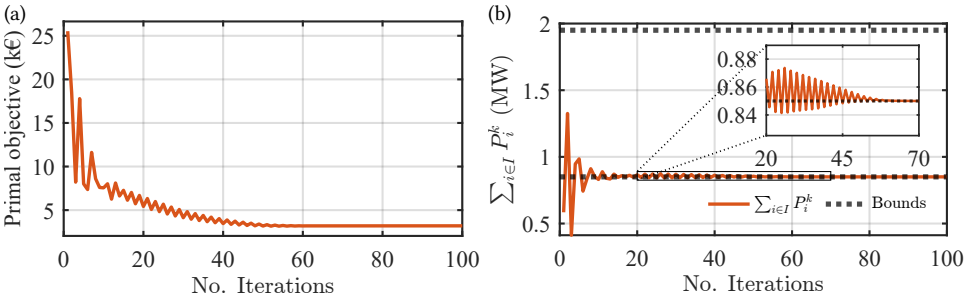

Figure 3.4: The convergence of the solution method for an instance: (a) primal objective; (b) aggregate active power within $[P_{\min}, P_{\max}]$.

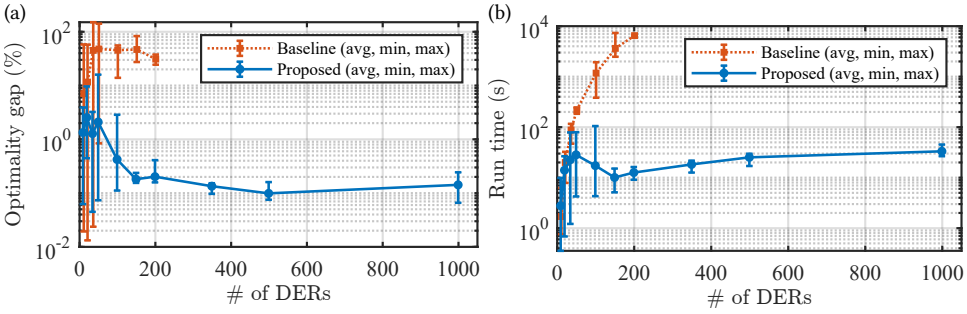

Figure 3.5: Performance of the proposed solution method: (a) optimality gap of the solutions using the proposed method and the baseline; (b) solve time using the proposed method and the baseline.

maintains solution quality as the number of DERs increases, whereas `fmincon` often fails to produce competitive solutions beyond a few tens of DERs and eventually becomes inapplicable. Moreover, the observed decrease in the gap as the system size increases supports Proposition 3.1, confirming that the relative optimality of the proposed method improves with scale.

Runtime scaling. Figure 3.5(b) shows the runtime as a function of problem size. The runtime of the baseline increases rapidly with $|I|$ and soon becomes prohibitive, whereas that of the proposed method grows only modestly. This improvement stems from decomposing the original large-scale nonconvex problem into tractable low-dimensional subproblems. All runs are performed on a single CPU core, suggesting that the computation time could be reduced further through parallel execution of the independent inner solves.

Overall, the results demonstrate that the proposed method scales to $\mathcal{O}(10^3)$ DERs while retaining a duality-gap certificate of solution quality.

3.4.4. Discussions

The proposed framework relies on several modelling and market assumptions. First, if DERs can be reliably partitioned into a few perfectly homogeneous groups, then a cluster-level coordination model may be sufficient and can reduce computational effort. However, in realistic DVPPs, DERs often differ in availability, costs, dynamic characteristics, and technical constraints, such that cluster-level averaging may introduce aggregation errors and require a nontrivial disaggregation step to recover feasible device-level allocations. The scalability claim of Chapter 3 should therefore be understood in this sense: the proposed decomposition-based dynamic-matching framework remains computationally tractable even when individual DER heterogeneity must be preserved.

Second, the framework assumes that the TSO specifies a well-defined FFR product in the form of a low-order target transfer function, and that the DVPP operator has access to suitable reduced-order models of the participating devices. This calls for future ancillary-service market designs that clear bids and specify dynamic services in such standardized forms, together with device-certification procedures that provide reliable reduced-order models. Nevertheless, the proposed formulation itself does not fundamentally require DER dynamics to follow a specific parametric structure. In principle, any stable transfer function can be incorporated by evaluating its frequency response at the selected sampling frequencies. This naturally includes higher-order or band-pass dynamics that may arise from converter control loops or filtering stages.

Network constraints are also treated in a simplified manner and should be incorporated more explicitly before field deployment. In the current work, we assume that the location of DERs does not influence the active-power response measured at the point of common coupling (substation to the transmission grid). This assumption is commonly adopted in system-level studies of frequency control and adequacy, where internal network constraints are neglected and each region is represented as a single node [47, 42]. Under this assumption, network topology and DER locations do not affect the aggregate FFR seen from the transmission side. A natural next step is to embed the proposed coordination scheme within more detailed distribution-network models and to account for the influence of line capacities, power losses, and related constraints.

Future work will also extend the framework to additional grid services (e.g., voltage regulation and synthetic inertia), richer network constraints, and experimental validation

in realistic DVPP demonstrators.

3.5. Conclusions

This chapter studied the operation of a DVPP that envelopes DERs to provide FFR services. The DVPP coordination is formulated as a frequency-domain dynamic matching problem between a target FFR transfer function and the aggregate of device-level FFR transfer functions, enabling accurate regulation and enhanced frequency stability. The formulation explicitly treats device latency as a decision variable, thereby exploiting additional temporal flexibility and operational benefits. To solve the resulting large-scale nonconvex program, a scalable dual-decomposition method is developed. The proposed method exploits additive structure, decomposes the original problem into per-device subproblems, and recovers high-quality primal solutions from the dual. Both theoretical results and case studies demonstrate strong performance, superior scalability, and runtime efficiency. Although designed for this specific setting, the solution method readily generalizes to other large-scale nonconvex programs. Overall, the proposed framework provides a scalable and efficient DVPP operation scheme that turns distributed renewables into dependable stability resources and reduces reliance on conventional synchronous units in high-renewable power systems.

4

Customised Nudging for Ancillary Service Provisions

Many DER prosumers are increasingly willing to contribute to ancillary services, yet this self-motivation is largely untapped. This chapter proposes a nudging mechanism that converts such willingness into operational benefits for a VPP. Beyond conventional price-based schemes, the mechanism provides each DER prosumer with customized recommendations on the suggested service level, which, together with financial incentives, iteratively steers decisions towards VPP operation targets. The intrinsic motivation of each DER is captured by a behavioural surrogate model, identified via an inverse-optimization-based, data-driven procedure that supports both offline initialization and online adaptation. We establish convergence guarantees showing that the mechanism attains the desired operating point while preserving budget balance. Case studies on synthetic inertia and short-circuit current support illustrate the effectiveness and versatility of the proposed approach for different types of ancillary services.

This chapter is based on 🗎 ***H. Xie**, A. Silani, and J. L. Cremer, "Nudging Distributed Energy Resources To Provide Power System Ancillary Services Using Customized Recommendations," submitted.*

4.1. Introduction

Power system ancillary services are essential functions that maintain system reliability, stability, and power quality, beyond the basic generation and delivery tasks [75]. These services include frequency control [76] (e.g., synthetic inertia, fast frequency response), voltage and reactive power support [77, 78], and Short-Circuit Current (SCC) contributions [79]. The rapid growth of inverter-based and Distributed Energy Resources (DERs) has increased demand for these services and has shifted responsibility from traditional synchronous generators to DERs themselves [80]. The critical importance of such services was highlighted by the April 2025 Iberian Peninsula blackout, which exposed gaps in fast, decentralized grid support [11]. In this context, Virtual power plants (VPPs) emerge as a scalable and promising means to aggregate DER flexibility and deliver these grid services as a whole [81, 16]. This chapter focuses on the operation of such a VPP deployed to provide ancillary services to the grid.

The core problem of operating such a VPP is to align the economic interests of DER prosumers with the technical objectives of the VPP operator: on the one hand, the VPP operator wants to deliver the exact requested aggregated flexibility to the system operator to make profits [82]. On the other hand, the main interests of DERs lie in maximizing their own utilization or profits. These two objectives are partially in conflict. For example, a VPP operator may be requested by the system operator to supply sufficient SCCs to support fault detection, protection device coordination, and voltage stability. However, delivering SCCs typically requires the inverter-based DERs to reserve a portion of their apparent power capacity for reactive power injection. This reserved capacity may limit their ability to export active power, potentially conflicting with the DER prosumers' economic interests. The existing literature approaches this problem in two directions to align the objectives of VPP and the DERs. The first line of research enforces rigid control protocols to the DERs, e.g., central optimal control [83], decentralized/ distributed optimization [82, 84], hierarchical control approach [85], multi-agent coordination [41]. Though these approaches enable VPP operators to achieve their operational targets, DER prosumers' autonomy is often overlooked in these mandatory schemes. The second line of research adopts incentive-based schemes, wherein the VPP operator steers DER behaviour by establishing economic signals, market mechanisms, or compensation rules, rather than issuing direct control commands. Methods including dynamic tariffs [86], auctions and bidding [87], marginal costs compensations [81], contractual and regulatory incentives [88] are extensively investigated to align DER prosumers' self-interested decisions with system-level objectives. Within this stream, *transactive control* is a decentral-

ized, market-based paradigm where DER agents exchange bid/utility information and an operator (or market mechanism) discovers an equilibrium price to coordinate supply, demand, and network constraints [89]. Implementations range from one-shot market clearing to iterative price–quantity exchanges, and primarily rely on DER price-responsiveness as the behavioural channel [89]. However, incentive-based schemes (including transactive control) typically model prosumers as predominantly price-responsive agents and therefore do not explicitly activate non-monetary motivations, leaving voluntary participation potential underutilized and potentially increasing required incentive payments. This motivates a critical yet largely unexplored question:

Can we systematically leverage the intrinsic willingness of DER prosumers to enhance the VPP operation efficiency?

This question gains increasing importance, as evidence shows that intrinsic motivations are becoming key drivers for DER prosumers to participate in ancillary services. For instance, a 2015 German survey found that approximately 69% of PV owners were willing to invest in PV–storage systems and to provide grid balancing services using their storage assets [90]. Similarly, a UK-based study [91] reported a high latent willingness and substantial untapped potential for residential demand-side participation. Notably, both surveys highlight that the motivations behind DER prosumers' willingness extend far beyond financial compensation, to assurances that their efforts support environmental goals (e.g., the integration of renewable energy and the avoidance of costly grid reinforcements). These findings indicate that many DER prosumers are driven by non-monetary motivations, including a commitment to supporting the broader energy transition. Recent research has begun to explore how non-monetary considerations, such as fairness [92], trust [93], and data privacy [94], influence the operation and coordination of power systems and VPPs. These factors, while not directly tied to compensation, are increasingly recognized as critical for sustained user engagement and effective system operation. Fairness and privacy-aware designs, for example, foster user acceptance and long-term participation. Moreover, it is demonstrated in [93] that building DER prosumers' trust in aggregator signals enables more accurate steering of their active power responses, thereby improving coordination efficiency. Motivated by these insights, this chapter aims to systematically consider DER prosumers' intrinsic motivations into VPP coordination frameworks.

To this end, a nudging mechanism is proposed in which the VPP operator provides each DER prosumer with an additional piece of information: customized recommendation on the amount of services that they are suggested to provide. The intrinsic motivations

of DERs are then realized once the DERs adapt their decisions towards aligning with the recommendations. The proposed mechanism is termed as a nudge, as it functions as a minimalist intervention that predictably influences DER prosumers' behaviour without prohibiting any options or significantly altering their economic incentives [38].

Implementing such a mechanism requires modelling prosumer responses to the joint incentive-recommendation signals. Therefore, a behavioural surrogate model of prosumer decision-making is constructed. In particular, we capture heterogeneity across DERs via a prosumer-specific commitment factor, which quantifies the tendency to follow service recommendations. Larger commitment factors correspond to more willing and proactive participation in fulfilling system-level service expectations, whereas smaller values reflect behaviour that is closer to purely price-responsive decision-making. Variations in commitment may arise from multiple sources, including technical knowledge, confidence in system participation, perceived social responsibility, and local community environment, among others [39, 40]. Our proposed mechanism explicitly respects this diversity and leverages it when tailoring recommendations to individual prosumers.

The main contributions of this chapter are as follows:

- A behaviour-aware signal-space nudging mechanism for VPP operation is proposed. The mechanism augments the broadcast incentive price with customized service recommendations. By leveraging prosumers' intrinsic willingness to align with non-binding guidance, the mechanism reduces the reliance on high monetary incentives and enables budget-compliant coordination.

- A data-driven identification method is developed to learn prosumer-specific behavioural surrogate models, in particular the commitment factor that governs recommendation following behaviour. Using inverse optimization techniques for both offline and online identification, the method makes the nudging mechanism deployable under behavioural heterogeneity and uncertainty.

- The Social Conformity Indicator (SCI) is introduced as an interpretable diagnostic metric to monitor recommendation adherence and prosumer engagement, providing operational visibility into the effectiveness of the nudging signals.

Case studies on a modified IEEE 39-bus system demonstrate the effectiveness and generality of the proposed framework in two representative ancillary-service settings: synthetic inertia provision and short-circuit current support. The results evaluate convergence and operational cost, budget compliance, and behavioural alignment, and further

assess robustness under behavioural model mismatch as well as the benefit of online commitment adaptation. Overall, the studies indicate that the proposed mechanism can reliably steer autonomous DERs toward VPP targets while converting intrinsic motivations into measurable coordination gains.

The remainder of this chapter is organized as follows. Section 4.2 formulates the VPP ancillary-service provision problem and introduces two representative use cases. Section 4.3 presents the behavioural surrogate model and the proposed nudging mechanism with customized recommendations. Section 4.4 develops a data-driven identification method for prosumers' surrogate models. Section 4.5 reports the case study results, and Section 4.6 concludes this chapter and outlines future research directions.

4.2. Optimal Ancillary Service Provisions of VPP

In this section, the general formulation of the VPP ancillary service provision problem is first established. Subsequently, two different but representative use cases are introduced to illustrate the problem: synthetic inertia and SCC.

4.2.1. Problem Formulation

Let us consider a VPP composed of a set $I = \{1, 2, \cdots, |I|\}$ of DERs, located at different buses in the power network. The role of the VPP is to coordinate these heterogeneous DERs to collectively provide ancillary services to the grid operator. Such services may include synthetic inertia, frequency response, or SCC support, depending on the requirements of the system operation.

Let us denote the target amount of aggregated service that the VPP is required to deliver as a known constant $x^* \in \mathbb{R}$. This regulation target may be derived from contractual obligations with the system operator or from market-clearing results in an ancillary service market. To fulfill this requirement, the VPP operator must determine the service contribution $x_i \in \mathbb{R}$ from each DER $i \in I$, while maintaining both service quality and economic performance. We formulate this coordination problem as follows:

$$\min_{x_i, i \in I} \quad \sum_{i \in I} f_i(x_i) + w \left| \sum_{i \in I} F_i x_i - x^* \right| \tag{4.1a}$$

$$\text{s.t.} \quad x_i \in X_i, \ \forall i \in I. \tag{4.1b}$$

In this formulation, $f_i : \mathbb{R} \to \mathbb{R}$ is the cost function that quantifies the economic expense or operational effort for DER i to provide a given level of service. F_i maps the local decision

x_i to its actual contribution in the aggregated response, possibly incorporating network effects such as electrical distance or control dynamics. w is a weighting parameter that balances cost efficiency against tracking accuracy. The selection of w is dependent on the operational priority/ performance requirements of VPP operator. X_i represents the physical or regulatory constraints that are local in DER i.

The first objective term represents the overall operation costs of the VPP. The use of the ℓ_1-norm in the second objective term reflects the goal of minimizing the absolute mismatch between the total provided service and the regulation target. This choice offers robustness to small deviations and aligns well with market-based settlement mechanisms that often penalize over- or under-delivery [82].

Remark. A key structural feature of problem (4.1) is that the objective depends on the sum of mapped local decisions, $\sum_{i\in I} F_i x_i$, forming a typical *sharing problem* in distributed optimization. This structure also resembles a block-wise variant of the *Lasso problem*, where each block corresponds to a DER's decision vector. Such a formulation is well-suited for scalable and distributed solution methods [61].

The generality of this formulation allows it to be instantiated for different ancillary services by appropriately specifying the mapping F_i, feasible sets X_i, and interpretation of x_i. In the next two subsections, we illustrate this framework through two representative use cases: synthetic inertia provision and SCC support.

4.2.2. Use Case 1: Synthetic Inertia Provision

This subsection instantiates the general coordination framework (4.1) for the case of synthetic inertia provision. For each DER $i \in I$, the decision variable is defined as $x_i \in \mathbb{R}$, denoting the synthetic inertia emulated by the DER. The VPP's goal is to regulate the aggregate response such that the total delivered synthetic inertia meet a prescribed target:

$$\sum_{i\in I} x_i \to x^*.$$

The active power that is needed to emulate the synthetic inertia of DER i during the disturbance is

$$\Delta P_i(t) = -2H_i \Delta \dot{f}(t), \quad \forall t \in \mathcal{T}_0,$$

where $\Delta P_i(t)$ dynamically varies with time t, and $\Delta \dot{f}(t)$ is the rate of change of frequency (RoCoF) over a time window $\mathcal{T}_0$ following the disturbance. Focusing on under-frequency

events, the total power output of the DER must not exceed its available power. This imposes the following constraint:

$$P_i + \Delta P_i(t) \leq P_i + \max_{t \in \mathcal{T}_0} \left| 2H_i \Delta \dot{f}(t) \right| \leq \overline{P_i}, \tag{4.2}$$

where P_i denotes the active power during normal operation, $\overline{P_i}$ is the maximum available power of DER i. It is typical to bound the maximum frequency deviation rate using the maximum permissible RoCoF specified by the system operator $\Delta \dot{f}_{\text{lim}}$, with which we have $\Delta \dot{f}_{\text{lim}} \geq \max_{t \in \mathcal{T}_0} |\Delta \dot{f}(t)|$. Therefore, (4.2) is conservatively converted as

$$P_i + 2\Delta \dot{f}_{\text{lim}} H_i \leq \overline{P_i}. \tag{4.3}$$

Together with the non-negativity of H_i, inequality (4.3) defines the local feasible set X_i for each DER.

4.2.3. Use Case 2: Short-Circuit Current Support

This section instantiates the general coordination framework (4.1) for the specific case of SCC support. Suppose the VPP operates within an m-bus transmission grid, with DERs distributed across a subset of buses. Let $\mathcal{B} \subseteq \{1, \dots, m\}$ denote the set of buses with connected DERs, and let $\psi(i) \in \mathcal{B}$ be the bus to which DER $i \in I$ is connected. Thus, we have $\cup_{i \in I} \psi(i) = \mathcal{B}$. During a short-circuit fault, inverter-based DERs providing SCC support are modeled as variable constant current sources, in accordance with the IEC 60909-0-2016 standard. From a system-level perspective, we aggregate the current injections across all buses during the fault into a vector $J := [J_1, J_2, \dots, J_m]^T \in \mathbb{R}^m$, and define $J' := [J'_1, J'_2, \dots, J'_m]^T \in \mathbb{R}^m$ as the baseline current injection without VPP participation. For each bus $b \in \mathcal{B}$, the net injection becomes:

$$J_b = J'_b + \sum_{\psi(i)=b} x_i, \quad \forall b \in \mathcal{B}, \tag{4.4}$$

where $x_i \in \mathbb{R}$ denotes the current injected from DER i during the fault, serving its decision variable in the coordination problem.

Applying Ohm's law to the equivalent network, the bus voltages satisfy $V = ZJ$, where $V = [V_1, V_2, \cdots V_m] \in \mathbb{R}^m$ denotes the vector of bus voltages, and $Z \in \mathbb{R}^{m \times m}$ is the network impedance matrix. The matrix Z is derived by combining the nodal admittance matrix and the subtransient reactances of the synchronous generators. Let f denote the faulted bus, and let the current injected into the bus as the positive direction. Since $V_f = 0$ during

the short-circuited fault, the f-th row of the voltage equation yields

$$J_{sc} := -J_f = \frac{\sum_{b\neq f} Z_{fb} J_b}{Z_{ff}} = \underbrace{\frac{\sum_{b\neq f} Z_{fb} J'_b}{Z_{ff}}}_{:=J'_{sc}} + \frac{\sum_{b\in\mathcal{B}} Z_{fb} \sum_{\substack{i\in I \\ \psi(i)=b}} x_i}{Z_{ff}}, \tag{4.5}$$

where $J_{sc} \in \mathbb{R}$ is the total SCC drawn from the fault bus, and $J'_{sc} \in \mathbb{R}$ is the SCC without VPP participation. The *incremental* SCC contribution from the VPP is therefore:

$$\Delta J_{sc} = \sum_{b\in\mathcal{B}} \frac{Z_{fb}}{Z_{ff}} \sum_{\psi(i)=b} x_i = \sum_{i\in I} \underbrace{\frac{Z_{f\psi(i)}}{Z_{ff}}}_{:=F_i} x_i. \tag{4.6}$$

Ideally, the VPP seeks to regulate the aggregate SCC provision such that $\Delta J_{sc} \to x^*$, aligning with its operational target. Equation (4.6) highlights that each DER's contribution to SCC at the fault bus is weighted by a location-dependent factor F_i, which is determined by the impedance between the DER's bus and the faulted bus.

4.3. Behaviour-Aware Nudging Mechanism

A key challenge in operating a VPP with autonomous prosumers is that the operator cannot directly set individual DER setpoints and must instead influence decisions through broadcast signals, most commonly a monetary incentive price. Relying on price alone can be inefficient: meeting a target may require high incentives, creating budget pressure for the operation of a VPP. We therefore adopt the concept of *nudging* from behavioural economics, which steers decisions by shaping the information environment while preserving freedom of choice. Concretely, we complement the incentive price with a non-binding, prosumer-specific customized recommendation indicating a suggested service provision level.

This section proceeds in two steps. First, we introduce a behavioural surrogate model that captures how one prosumer responds jointly to the nudging signals. Second, building on this surrogate model, we present the proposed nudging mechanism, which iteratively updates the price and the recommendations within a tâtonnement process to steer the aggregate VPP provisions toward the desired target while respecting the operator's budget.

4.3.1. Behavioural Surrogate Model Under Price Incentives and Recommendations

Let $\lambda \in \mathbb{R}$ denote the unit incentive price, and let $r \in \mathbb{R}$ denote the recommended provision level sent by the VPP. Given (λ, r), prosumer i selects its provision x_i by optimizing its private objectives subject to local constraints. In practice, prosumer responses to the same signal pair (λ, r) are heterogeneous and noisy due to differences in private costs $f_i(\cdot)$, feasible regions X_i, and recommendation-following behaviour. To support scalable mechanism design, we adopt a tractable *behavioural surrogate model* that captures prosumers responses to broadcast signals. Specifically, we model prosumer i's decision as the solution to

$$x_i(\lambda, r) = \arg\min_{x \in X_i} \; f_i(x) - \lambda x + \frac{\sigma_i}{2}|x - r|^2, \tag{4.7}$$

where $f_i : \mathbb{R} \to \mathbb{R}$ is prosumer i's private (convex) cost for providing the ancillary service and X_i is its local feasible set. The term $-\lambda x$ captures the monetary incentive, while $\sigma_i|x - r|^2/2$ is a recommendation-following regulariser that penalizes deviations from the recommended level r. This regulariser is a reduced-form representation of intrinsic motivations that make a prosumer more inclined to align with recommendations beyond purely price-driven optimisation. As σ_i increases, the optimal response $x_i(\lambda, r)$ places greater weight on adhering to r. We therefore refer to $\sigma_i > 0$ as the *commitment factor*.

We adopt the surrogate model (4.7) for three reasons. First, it preserves prosumer autonomy: each prosumer remains free to optimise its own objective over its feasible region X_i, and the operator imposes no hard constraints or direct control commands. Second, it explicitly accounts for financial incentives through λ, ensuring compatibility with standard compensation-based VPP operation. Third, it provides a low-parameter yet heterogeneous representation of recommendation-following behaviour through a single prosumer-specific parameter σ_i. As the commitment factor σ_i increases, deviations from the recommendation r are penalised more strongly, and the optimal response $x_i(\lambda, r)$ aligns more closely with r; when σ_i is small, the prosumer behaves closer to being purely price-responsive. Heterogeneous values of σ_i capture that prosumers may differ in their intrinsic willingness to contribute, due to behavioural preferences, perceived social responsibility, convenience or comfort considerations, trust in the operator, and other non-monetary factors. Moreover, the quadratic term $|x - r|^2$ symmetrically penalises both under- and over-provision relative to the recommendation, reflecting that, from a system-operation's perspective, insufficient service degrades reliability while excessive provision can be inefficient or violate operational limits.

The surrogate response model (4.7) serves as the behavioural foundation for the mechanism developed in subsequent subsections. It provides a systematic way to translate the operator's signals (λ, r_i) into prosumer actions, and it yields a low-dimensional behavioural parameter σ_i that can be identified from observed responses in a data-driven manner (Section 4.4). While (4.7) is used as a behavioural *baseline* for mechanism design and inference, real prosumers may follow alternative (possibly heuristic) decision-making rules that deviate from this model; to reflect this reality, the case studies additionally consider such alternative decision rules and evaluate the robustness of the proposed mechanism.

4.3.2. Design Objectives and Tâtonnement Architecture

Fig. 4.1 illustrates the proposed nudging architecture, implemented as an iterative tâtonnement coordination layer between the VPP operator and the DER prosumers. At each iteration k, the operator broadcasts an incentive price λ^k and customized recommendations $\{r_i^k\}_{i\in I}$; each prosumer then updates its provision decision and reports it back to the operator. The recommendations are customized to account for heterogeneity in feasible regions X_i, cost structures f_i, grid impact factors F_i, and behavioural responsiveness which is captured by the commitment factor, identified from data in Section 4.4. The tâtonnement iteration proceeds as follows:

1. **Operator update.** Given the previously observed decisions $\{x_i^{k-1}\}$, the operator updates the incentive price λ^k and computes new individualized recommendations $\{r_i^k\}$ according to the mechanism to be designed, aiming to steer the aggregate VPP provision toward the desired target while respecting the budget.

2. **Prosumer update.** Each prosumer i updates its decision in response to the broadcast signals (λ^k, r_i^k), and communicate the resulting decision x_i^k back to the operator.

A key feature of the tâtonnement process is that it supports prosumer autonomy while limiting the disclosure of local information: each prosumer keeps X_i and $f_i(\cdot)$ local and does not disclose them to the operator. This avoids the information barrier that the operator cannot compute the optimum of (4.1) *a priori* without collecting private costs and constraints. Instead, the mechanism only requires identifying a low-dimensional behavioural parameter: the commitment factor σ_i in (4.7), and reaches the optimal allocation *implicitly* through iterative price-and-recommendation updates and prosumers' self-optimizing responses. Importantly, the tâtonnement iterations operate on a coordination time-scale

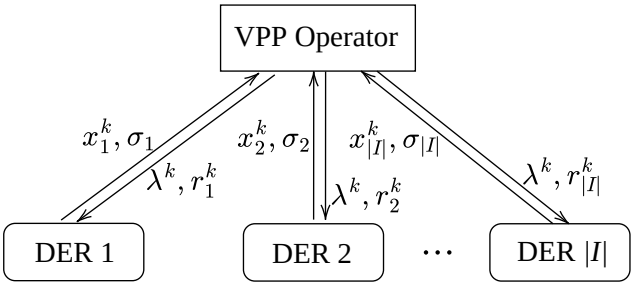

Figure 4.1: The tâtonnement process of the nudging mechanism.

(e.g., minutes to hours) and are *not* required during contingency dynamics. Once parameters are deployed, each DER implements the corresponding local ancillary-service controller mode autonomously, triggered by measured electrical quantities (e.g., frequency or terminal voltage) and responding on millisecond time-scales.

Accordingly, the nudging mechanism is designed to achieve two objectives: (i) **Convergence:** the iterates $\{x_i^k\}$ converge to the optimal solution of (4.1); (ii) **Budget compliance:** the long-run total incentive payment remains within the operator's available budget Λ, i.e.,

$$\limsup_{k\to\infty} \lambda^k \sum_{i\in I} x_i^k \le \Lambda. \tag{4.8}$$

4.3.3. Nudging Mechanism: Price-Recommendation Update Rules

The following update rules are proposed:

$$\lambda^k = \begin{cases} \frac{\Lambda}{\sum_{i\in I} x_i^{k-1}}, & \text{if } \sum_{i\in I} x_i^{k-1} \ge \epsilon, \\ \frac{\Lambda}{\epsilon}, & \text{if } \sum_{i\in I} x_i^{k-1} < \epsilon, \end{cases} \tag{4.9a}$$

$$r_i^k = z_i^k - \frac{\lambda^k + y_i^k}{\sigma_i}, \tag{4.9b}$$

where ϵ is a predefined minimal provision level, and (4.9b) are built upon the updates of two auxiliary variables z and y as

$$\begin{cases} z^k = \arg\min_{z'} w \left| \sum_{i\in I} F_i z_i' - x^* \right| + \sum_{i\in I} \frac{\sigma_i}{2} \left| z_i' - x_i^{k-1} - \frac{1}{\sigma_i} y_i^{k-1} \right|^2 \\ y_i^k = y_i^{k-1} + \sigma_i (x_i^{k-1} - z_i^k). \end{cases} \tag{4.10}$$

We establish the following convergence proof of the proposed nudging mechanism, based on the assumption that the DER prosumers in the VPP are agents whose decision-making processes are *exactly* described by the behavioural surrogate model (4.7). In particular, during the iterations, each prosumer i responds to (λ^k, r^k) by

$$x_i^k \in \arg\min_{x\in X_i} \; f_i(x) - \lambda^k x + \frac{\sigma_i}{2}|x - r_i^k|^2, \tag{4.11}$$

where $f_i(\cdot)$ is convex and the local feasibility region X_i is a convex and compact polyhedral set. This assumption corresponds to the idealized case where the identified (learned) commitment factors yield a surrogate model that perfectly matches each prosumer's true input–output response to (λ, r).

4

Proposition 4.1. Under the above ideal assumption, any sequence of DER decisions $\{x_i^k\}_{i\in I}$ (4.11) under the proposed mechanism (4.9a), (4.9b), (4.10), initialized as $(\mathrm{col}(x^0), z^0, \mathrm{col}(y^0)) \in \prod_{i\in I} X_i \times \mathbb{R}^{|I|} \times \mathbb{R}^{|I|}$, converges to an optimal solution of (4.1). Moreover, the budget compliance condition (4.8) is satisfied *at convergence*. In particular, if the mechanism terminates at an index K such that $x^K = x^{K-1}$, then $\lambda^K \sum_{i\in I} x_i^K \le \Lambda$.

Proof. **Convergence to an optimal solution.**

Let $X := \prod_{i\in I} X_i$, $x := \mathrm{col}(x_1, \dots, x_{|I|})$, $z := \mathrm{col}(z_1, \dots, z_{|I|})$, and define

$$f(x) := \sum_{i\in I} f_i(x_i) + I_X(x), \qquad g(z) := w\Big|\sum_{i\in I} F_i z_i - x^\star\Big|.$$

Introduce the splitting constraint $\Psi x - \Psi z = 0$ with $\Psi := \mathrm{diag}(\sqrt{\sigma_1}, \dots, \sqrt{\sigma_{|I|}})$, and consider the equivalent convex problem

$$\min_{x,z} \; f(x) + g(z) \quad \text{s.t.} \quad \Psi x - \Psi z = 0, \tag{4.12}$$

which is equivalent to (4.1) since Ψ is nonsingular ($\sigma_i > 0$).

Using (4.9b) and completing squares, the prosumer response (4.7) can be written in stacked form as

$$x^k = \arg\min_{x\in X} \; f(x) + \frac{1}{2}\left|\Psi x - \Psi z^k + u^k\right|^2,$$

where $u^k := \mathrm{col}\left(y_1^k/\sqrt{\sigma_1}, \dots, y_{|I|}^k/\sqrt{\sigma_{|I|}}\right)$. Similarly, the auxiliary updates (4.10) are equivalent to

$$z^{k+1} = \arg\min_{z} \; g(z) + \frac{1}{2}\left|\Psi x^k - \Psi z + u^k\right|^2, \qquad u^{k+1} = u^k + \Psi x^k - \Psi z^{k+1}.$$

Thus, the mechanism implements the standard augmented-Lagrangian splitting iterations for the convex problem (4.12) (up to a harmless index shift). Under assumptions, f and g

are closed, proper, and convex, and a saddle point exists; therefore, the classical convergence result for these iterations applies [61], yielding primal feasibility $\Psi(x^k - z^k) \to 0$ and objective convergence $f(x^k) + g(z^k) \to p^\star$. Since Ψ is nonsingular, $\Psi(x^k - z^k) \to 0$ implies $x^k - z^k \to 0$, and hence x^k converges to an optimal solution of (4.1).

Budget compliance at convergence.

Let $S^k := \sum_{i \in I} x_i^k$. From (4.9a), we know that

$$\lambda^k = \frac{\Lambda}{\max\{S^{k-1}, \epsilon\}},$$

this implies $\lambda^k S^{k-1} \leq \Lambda$ and $0 \leq \lambda^k \leq \Lambda/\epsilon$. Hence,

$$\lambda^k S^k \leq \Lambda + \frac{\Lambda}{\epsilon} |S^k - S^{k-1}|.$$

Since $x^k \to x^\star$, we have $|S^k - S^{k-1}| \to 0$, and taking $\limsup_{k \to \infty}$ yields (4.8). If the mechanism terminates at K with $x^K = x^{K-1}$ (so $S^K = S^{K-1}$), then

$$\lambda^K S^K = \lambda^K S^{K-1} \leq \Lambda.$$

□

4.3.4. Social Conformity Indicator for Mechanism Diagnosis

To facilitate diagnosing the behavioural response of DERs under the nudging mechanism, we introduce the *Social Conformity Indicator* (SCI):

Definition 4.1 (Social Conformity Indicator)**.** The *Social Conformity Indicator* for a DER $i \in I$ is defined as

$$\mathrm{SCI}_i := e^{-\alpha |x_i - r_i|}, \tag{4.13}$$

where x_i is its actual provision, r_i is the recommended provision from the VPP operator, and $\alpha > 0$ is a sensitivity parameter that penalizes deviation from the recommendation.

The SCI indicates the extent to which a DER's realized response aligns with the operator's recommendation and thus reflects the coordination intent of the mechanism. It directly quantifies behavioural deviation through $|x_i - r_i|$: a high SCI corresponds to close adherence and strong behavioural conformity, whereas a low SCI reflects large deviations, potentially signaling weak engagement or unreliable response. The SCI therefore serves as a simple yet principled diagnostic tool for monitoring the functioning of the nudging mechanism. In particular, a persistently low *average* SCI across prosumers suggests that

recommendations are not being effectively followed, for example due to low responsiveness, behavioural noise, or model mismatch; in such cases, the operator may need to adjust the recommendation design, increase incentives, or recalibrate the learned surrogate models.

4.4. Data-Driven Behaviour Surrogate Modeling of Prosumer Responses

In Section 4.3, we introduced the behaviour surrogate model (4.7) and a nudging mechanism built upon it. A key parameter is the prosumer-specific commitment factor σ_i, which quantifies recommendation-following behaviour but is not directly observable in practice and must be inferred from prosumers' observed responses.

Motivated by this, we develop a data-driven approach to estimate σ_i by fitting the surrogate model to historical and/or online signal–response data. The resulting estimation is posed as an *inverse optimization* problem [95], which infers objective parameters from observed (approximately) optimal decisions. For notational simplicity, we omit the subscript i throughout this section, as identification is performed independently for each prosumer.

4.4.1. Workflow and Inverse Optimization Formulation

The identification workflow is illustrated in Fig. 4.2. Following the standard inverse-optimization paradigm, we first specify a parametric *forward optimization problem* (FOP) that represents a prosumer's decision rule under the joint nudging signal, and then infer the unknown behavioural parameter(s) from observed decisions.

Forward problem (prosumer response model). Let $v := (\lambda, r)$ denote the observable input signal pair, where λ is the unit incentive price and r is the recommendation. For a given prosumer, consider the parametric forward model

$$\mathrm{FOP}(v;\vartheta): \quad x^\star(v;\vartheta) \in X^{\mathrm{opt}}(v;\vartheta) := \arg\min_{x\in X} \; G(x,v;\vartheta), \tag{4.14}$$

where X is the known feasible set and $G(x,v;\vartheta) := f(x) - \lambda x + \frac{\sigma}{2}|x-r|^2$. Here ϑ denotes the (unknown) behavioural parameter vector, with the *commitment factor* σ being the main quantity of interest. Under the convexity assumptions used in Section 4.3, $\mathrm{FOP}(v;\vartheta)$ is strongly convex and admits a unique optimizer $x^\star(v;\vartheta)$.

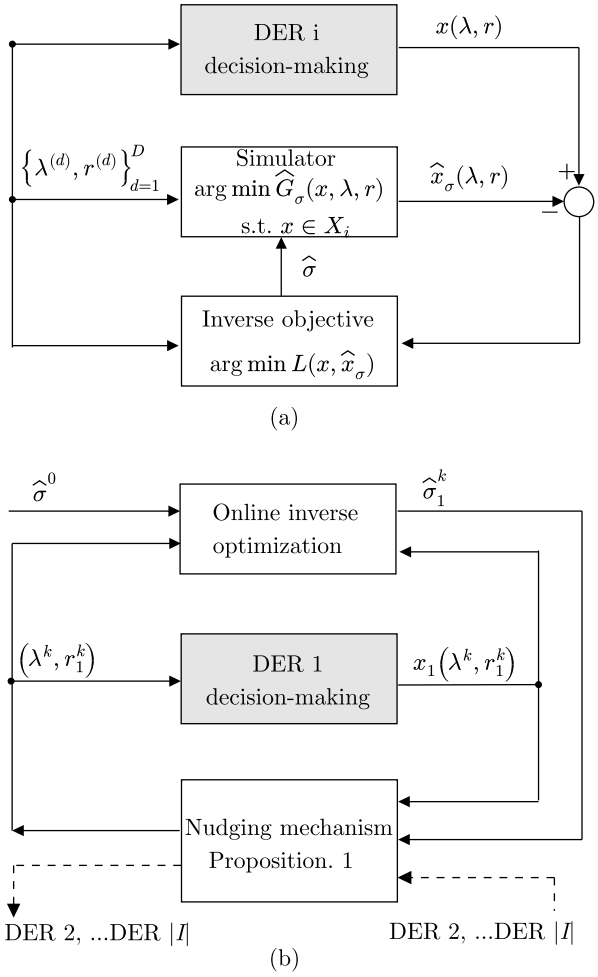

Figure 4.2: The workflow of the proposed identification toolbox.

Decision data (offline stage). In the offline stage (Fig. 4.2(a)), we query the prosumer with a finite set of trial inputs $\{v^{(d)} = (\lambda^{(d)}, r^{(d)})\}_{d=1}^{D}$ and observe the corresponding decisions $\{\check{x}^{(d)}\}_{d=1}^{D}$. Due to noise, unmodeled factors, and possible model mismatch, $\check{x}^{(d)}$ need not be exactly optimal for any ϑ. Hence, we adopt a data-driven inverse optimization formulation that estimates parameters by minimizing a loss measuring the violation of inverse-feasibility.

Data-driven inverse problem (offline identification). Given the dataset $\mathcal{D} := \{(v^{(d)}, \check{x}^{(d)})\}_{d=1}^{D}$, we estimate ϑ by solving

$$\text{IOP}-\text{DD}(\mathcal{D}) : \quad \min_{\vartheta\in\Theta} \frac{1}{D}\sum_{d=1}^{D} \ell\big(\check{x}^{(d)}, X^{\text{opt}}(v^{(d)};\vartheta)\big), \tag{4.15}$$

where Θ is the admissible parameter set and $\ell(\cdot,\cdot)$ is a model–data fit loss. In this chapter we use a squared minimum-distance (predictability) loss:

$$\ell\big(\check{x}, X^{\text{opt}}(v;\vartheta)\big) := \min_{x\in X^{\text{opt}}(u;\vartheta)} |x-\check{x}|^2. \tag{4.16}$$

Since $X^{\text{opt}}(v;\vartheta)$ is a singleton under strong convexity, (4.16) reduces to $|x^\star(v;\vartheta)-\check{x}|^2$.

Online stage (adaptive identification). Fig. 4.2(b) depicts the online stage, where new signal–response pairs $(u^k, \check{x}^k)$ are collected during tâtonnement. The estimate $\hat{\vartheta}^k$ is updated sequentially and fed back into the nudging mechanism to compute (λ^{k+1}, r^{k+1}), enabling adaptation without requiring a dedicated offline dataset for every prosumer.

4.4.2. Data-Driven Estimation of Commitment Factors

We now instantiate (4.15) under a tractable parametric form. Assume the private cost admits a quadratic surrogate $f(x) = ax^2 + bx$ with $a \geq 0$, $b \in \mathbb{R}$. For a given input v, the forward optimizer of FOP$(v;\vartheta)$ with $\vartheta = (a, b, \sigma)$ has the closed form

$$x^\star(v;\vartheta) = \left[\frac{\lambda + \sigma r - b}{2a+\sigma}\right]_X, \tag{4.17}$$

where $[\cdot]_X$ denotes projection onto X.

Linear re-parameterization. Define the regression vector $\phi(v) := \begin{bmatrix}1 & \lambda & r\end{bmatrix}^\top$ and re-parameterize

$$\theta := \begin{bmatrix}\theta_1\\ \theta_2\\ \theta_3\end{bmatrix} = \frac{1}{2a+\sigma}\begin{bmatrix}-b\\ 1\\ \sigma\end{bmatrix}, \qquad \Rightarrow \qquad x^\star(v;\vartheta) = \big[\phi(v)^\top\theta\big]_X, \tag{4.18}$$

so that the commitment factor can be recovered as $\sigma = \theta_3/\theta_2$ (with $\theta_2 > 0$).

Algorithm 3: Data-Driven Commitment Identification Method Based On Inverse Optimization

Offline Stage

Inputs: Dataset $\{(\lambda^{(d)}, r^{(d)})\}_{d=1}^{D}$ and $\{\breve{x}^{(d)}\}_{d=1}^{D}$;

Outputs: Initial estimate $\hat{\theta}^0$, commitment factor $\hat{\sigma}^0$;

Procedure `OfflineIdentification`:

  for $d = 1$ **to** D **do**

    Construct regression vector: $\phi^{(d)} \leftarrow \begin{bmatrix} 1 & \lambda^{(d)} & r^{(d)} \end{bmatrix}^\top$;

  Solve (4.19);

  Recover: $\hat{\sigma}^0 = \hat{\theta}_3^0 / \hat{\theta}_2^0$;

Online Stage

Inputs: Online data stream $\{(\lambda^k, r^k, \breve{x}^k)\}_{k=1}^{\infty}$, initial estimate $\hat{\theta}^0$, forgetting factor $\gamma \in (0, 1]$, initial covariance scale $\alpha > 0$;

Outputs: Updated parameter $\hat{\theta}^k$ and $\hat{\sigma}^k$;

Procedure `OnlineIdentification`:

  Initialize $P^0 \leftarrow \alpha I$, $\theta^0 \leftarrow \hat{\theta}^0$;

  for $k = 1, 2, \ldots$ **do**

    Observe new data $(\lambda^k, r^k, \breve{x}^k)$;

    Construct: $\phi_k \leftarrow \begin{bmatrix} 1 & \lambda^k & r^k \end{bmatrix}^\top$;

    Update gain vector: $K_k \leftarrow \frac{P_{k-1}\phi_k}{\gamma + \phi_k^\top P_{k-1}\phi_k}$;

    Update predicted decision: $\hat{x}^k \leftarrow [\phi_k^\top \theta^{k-1}]_X$;

    Update parameter estimate: $\theta^k \leftarrow \theta^{k-1} + K_k(\breve{x}^k - \hat{x}^k)$;

    Update inverse correlation matrix: $P_k \leftarrow \frac{P_{k-1} - K_k \phi_k^\top P_{k-1}}{\gamma}$;

    Recover: $\hat{\sigma}^k \leftarrow \theta_3^k / \theta_2^k$;

Offline inverse optimization. Given $\mathcal{D} = \{(v^{(d)}, \breve{x}^{(d)})\}_{d=1}^{D}$, we estimate θ by minimizing the minimum-distance loss (4.16), which here becomes the predictability loss

$$\hat{\theta} = \arg\min_{\theta \in \Theta} \frac{1}{D} \sum_{d=1}^{D} \left| \breve{x}^{(d)} - \left[\phi(v^{(d)})^\top \theta \right]_X \right|^2, \tag{4.19}$$

and recover $\hat{\sigma} = \hat{\theta}_3 / \hat{\theta}_2$. As the projection $[\cdot]_X$ induces a piecewise-affine forward map, (4.19) is generally non-convex. We therefore solve it using a local nonlinear optimizer

with a multi-start strategy to mitigate sensitivity to initialization and improve the obtained solution.

From a statistical viewpoint, (4.19) can be interpreted as an empirical risk minimization problem with a squared minimum-distance loss. For convex forward optimization problems, such minimum-distance formulations admit statistical consistency guarantees under standard regularity assumptions and i.i.d. sampling: the empirical objective converges (in probability) to the population objective (risk consistency), and under additional identifiability conditions one can further obtain parameter consistency; see [96] and the summary in [95, Thm. 2]. Importantly, these guarantees are conditional on correct model specification, i.e., that the behaviour surrogate model (4.7) can perfectly represent the prosumers' decision rule. This is generally not the case in practice due to unobserved factors, noise, and behavioural deviations. Accordingly, we do *not* claim to recover accurate commitment factors; rather, we target to identify the best-fitting surrogate parameter within the model class (4.7), which is sufficient to explain observed behaviour and is useful for the nudging signal design.

Online update. In the online stage, θ is updated sequentially as new observations $(v^k, \check{x}^k)$ arrive, using an Recursive Least Square (RLS)-style recursion with forgetting factor $\gamma \in (0, 1]$ [97]; the resulting $\hat{\sigma}^k$ is then passed to the nudging mechanism to update the next recommendation. The forgetting factor determines the weight given to past observations. A smaller γ prioritizes recent data, making the estimator more responsive to dynamic changes in commitment levels. By continuously refining $\hat{\theta}$, our proposed identification method ensures that the estimated commitment factor $\hat{\sigma}$ remains adaptive and consistent with observed behaviour. This dynamic tracking of heterogeneous DER commitment levels also improves the effectiveness of our proposed recommendation mechanism. Our proposed identification method is reported in Algorithm 3. Note that the prediction step explicitly includes the projection (step 18), so the update accounts for saturation at the feasible bounds.

Remark. In Algorithm 3, the initial estimate $\hat{\sigma}^0$ obtained from the offline stage is used to initialize the online RLS algorithm. In practical settings, DERs can often be grouped according to similarities such as geographical location, community profiles, or device types. DERs within the same group are likely to exhibit similar levels of commitment. As a result, they can share a common initial estimate $\hat{\sigma}^0$, thereby alleviating the need to collect a dedicated offline dataset for each DER.

4.5. Case Study

4.5.1. Simulation Settings and Scenarios

We use a modified IEEE 39-bus transmission system [98]. The system base is set to 100 MVA. A VPP comprising 20 DERs is integrated into the network. These DERs are located at four buses (bus 31, bus 32, bus 34, and bus 35), and have rated capacities uniformly distributed between 50 MW and 250 MW. Collectively, the DERs account for approximately 50% of the total system generation capacity (6.28 GW). Two ancillary service use cases are constructed: SCC support and synthetic inertia provision. The SCC case is used to evaluate the effectiveness of the nudging mechanism given the behaviour surrogate models are learned from data (Sec. 4.5.2). The synthetic inertia use case is used to evaluate the identification quality and online adaptivity of the behaviour surrogate model (Sec. 4.5.3). To stress-test the behaviour-aware mechanism under realistic behavioural heterogeneity, we consider the following prosumers' response modes (all respecting local constraints X_i and private costs f_i):

- **Deadband mode (tolerance to recommendations).** The prosumer exhibits weak reaction to the recommendation within a tolerance band of width $\delta = 0.1$ p.u. in the case studies, represented by a deadband term:

$$x_i^{\mathrm{DB}}(\lambda, r) = \arg\min_{x \in X_i} \; f_i(x) - \lambda x + \frac{\sigma_i}{2}\left([|x - r| - \delta]_+\right)^2,$$

  where $[t]_+ := \max\{t, 0\}$.

- **Asymmetric mode (directional response).** The prosumer reacts differently to under- vs. over-recommendation deviations by using asymmetric quadratic penalties:

$$x_i^{\mathrm{AS}}(\lambda, r) = \arg\min_{x \in X_i} \; f_i(x) - \lambda x + \frac{\sigma_i}{2}(x - r)_+^2 + \frac{\kappa\sigma_i}{2}(r - x)_+^2,$$

  where $\kappa = 0.5$ controls the directional asymmetry.

- **Bounded-rational mode (noisy/smoothed response).** The prosumer follows the baseline optimizer up to bounded rationality, modelled as a noisy response around $x_i^{\mathrm{DER}}(\lambda, r)$, which denotes the baseline surrogate response given by (4.7):

$$x_i^{\mathrm{BR}}(\lambda, r) = \left[x_i^{\mathrm{DER}}(\lambda, r) + \xi_i\right]_{X_i}, \qquad \xi_i \sim \mathcal{N}(0, \tau^2),$$

  with Gaussian perturbations and standard deviation $\tau = 0.02$.

In all the studies, DERs 1-5, 6-15, and 16-20 are assigned to the deadband, asymmetric, and bounded-rational modes, respectively. The detailed settings for each use case are described below.

SCC Support

We consider an SCC demand of $x^* = 15$ p.u. at bus 16. The location-dependent factor F_i, which quantifies the contribution of each DER's injected current to the SCC observed at bus 16, is derived from the admittance and impedance data of the IEEE 39-bus network [98]. Each DER's provision cost is modeled by a quadratic function $f_i(x_i) = a_i x_i^2 + b_i x_i$, where a_i is sampled uniformly from the range $[10, 30]$ €/p.u.^{2}, and b_i from $[0, 15]$ €/p.u. The current injections from each DER during the short-circuit faults are constrained by an upper limit of 1.2 times its rated current, consistent with typical inverter current limiter settings. The cost-tracking weight parameter is set to $w = 25$ €/p.u. DER commitment levels σ_i are categorized into three groups: low ($0 < \sigma_i \leq 2$), medium ($2 < \sigma_i \leq 6$), and high ($6 < \sigma_i \leq 10$), all in €/p.u.^{2}. In this case study, the nudging mechanism is implemented using prosumer-specific behaviour surrogate models that are identified *offline* prior to the tâtonnement iterations. Specifically, for each DER prosumer we collect 10 input–output samples of the form (λ, r, x) and fit the corresponding surrogate response model as described in Section 4.4; the resulting learned model parameters are then used by the operator to generate the individualized recommendations during the nudging process.

Synthetic Inertia Provision

In this case, the VPP is required to provide an aggregate synthetic inertia response of $x^* = 500$ MW·s/Hz. The cost function for each DER is again modeled as $f_i(x_i) = a_i x_i^2 + b_i x_i$, with $a_i \in [5, 15]$ €/(MW · s/Hz)^2 and $b_i \in [10, 15]$ €/(MW · s/Hz). The feasible set for each DER is determined by its available power margin, sampled uniformly from the interval $[0.2, 0.6]$ p.u. The weight parameter is set to $w = 20$ €/(MW · s/Hz). Commitment levels σ_i are categorized analogously: low ($0 < \sigma_i \leq 10$), medium ($10 < \sigma_i \leq 30$), and high ($30 < \sigma_i \leq 50$), all in €/(MW · s/Hz)^2. The trial incentive-recommendation pairs used for σ identification is sampled uniformly from $\lambda \in [0, 15]$ € and $r \in [0, 0.5]$ p.u.. The nonlinear non-convex problem formulated for the offline identification is solved with the solver *fmincon* in MATLAB.

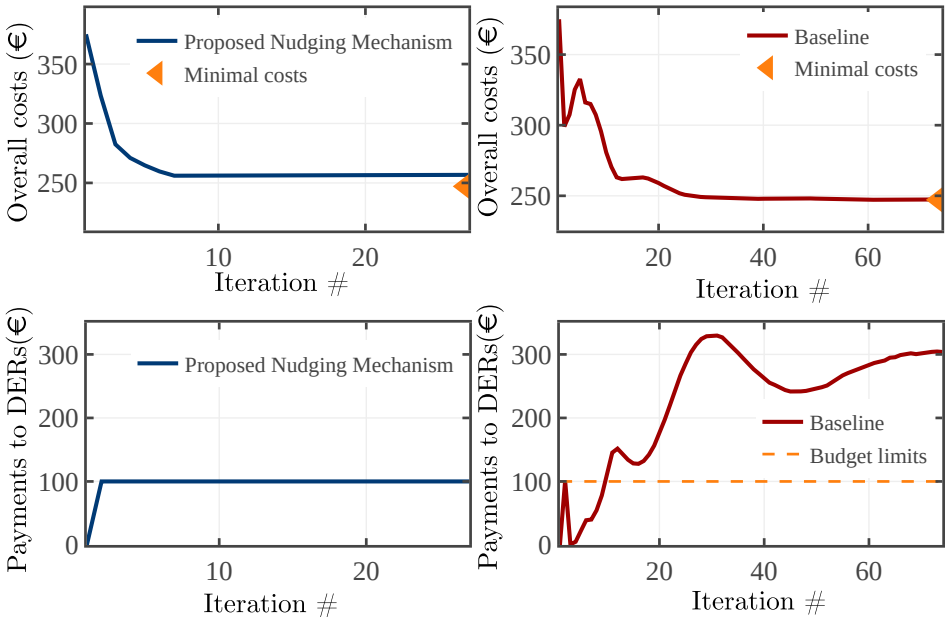

Figure 4.3: The convergence results and the payments to DERs during the tâtonnement process (the proposed nudging mechanism versus the benchmark).

4.5.2. Effectiveness of The Nudging Mechanism

The results presented in this subsection were studied on the SCC support use case. We introduce a baseline model for the same VPP operation problem, which relies fully on financial incentives. The details of the baseline are provided in appendix 4.7.1.

Convergence Performance

Fig. 4.3 presents the convergence trajectories of the proposed nudging mechanism versus the baseline. The top row illustrates the overall operational cost of the VPP over the tâtonnement iterations. The proposed nudging mechanism reaches (near-)optimal cost in noticeably fewer iterations than the baseline scheme, albeit with a small residual suboptimality in the final cost. The faster convergence of the proposed nudging mechanism comes from the fact that its tâtonnement process coincides with minimizing an augmented Lagrangian function of the original problem, which is a smooth function and thus leads to a faster convergence to the global minimum. This gap is expected: although the behavioural surrogate models are learned from offline data to approximate prosumers' responses to (λ, r), the true decision rules of DER prosumers need not be perfectly repre-

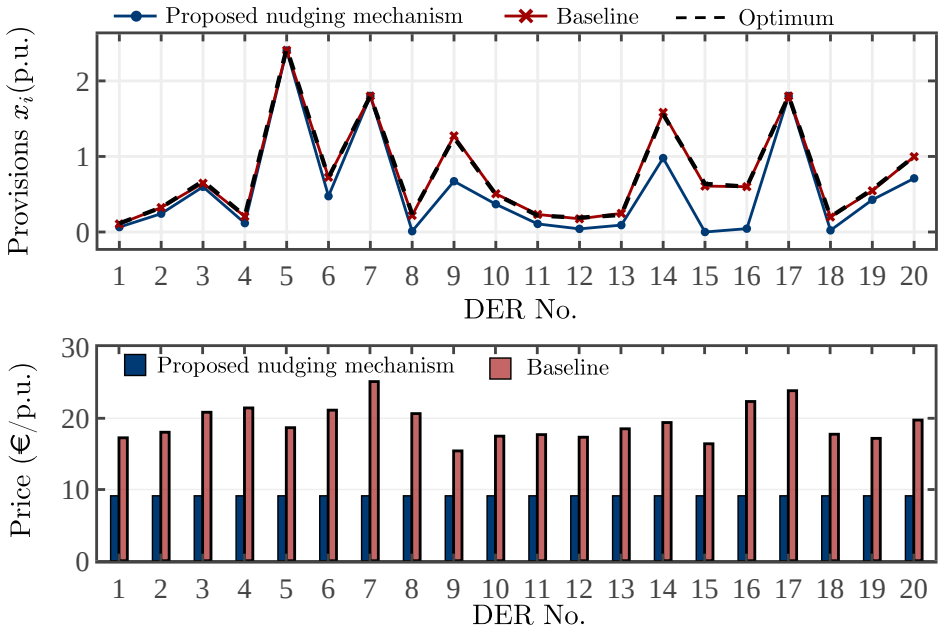

Figure 4.4: The ancillary services provisions from DERs and the corresponding clearing price to them.

sented by the surrogate class. As a result, model mismatch and inevitable behavioural noise introduce deviations between predicted and realized responses, which can prevent the iterations from attaining the exact optimum of the idealized model. The bottom row of Fig. 4.3 shows the cumulative incentive payments made to DERs. The total payments remained within the budget of VPP (100 €) across all iterations in the proposed nudging mechanism, demonstrating its budget compliance capabilities. In contrast, the baseline approach exceeded the budget threshold, as it relies fully on financial incentives to steer the behaviour of DERs. Fig. 4.4 compares the final SCC contributions and assigned prices for each DER under both methods. The top panel shows the cleared DERs' SCC support; the bottom panel shows the clearing corresponding unit prices (in €/p.u.) assigned to each DER. Both approaches steer the DERs to deliver SCC support close to the optimal target implied by (4.1). However, the proposed nudging mechanism yields a *uniform* incentive price across DERs, leading to a more even and transparent compensation rule. In contrast, the baseline produces heterogeneous prices driven primarily by marginal-cost differences. This outcome aligns with the behaviour surrogate design: customized recommendations absorb part of the coordination burden that would otherwise require stronger price dis-

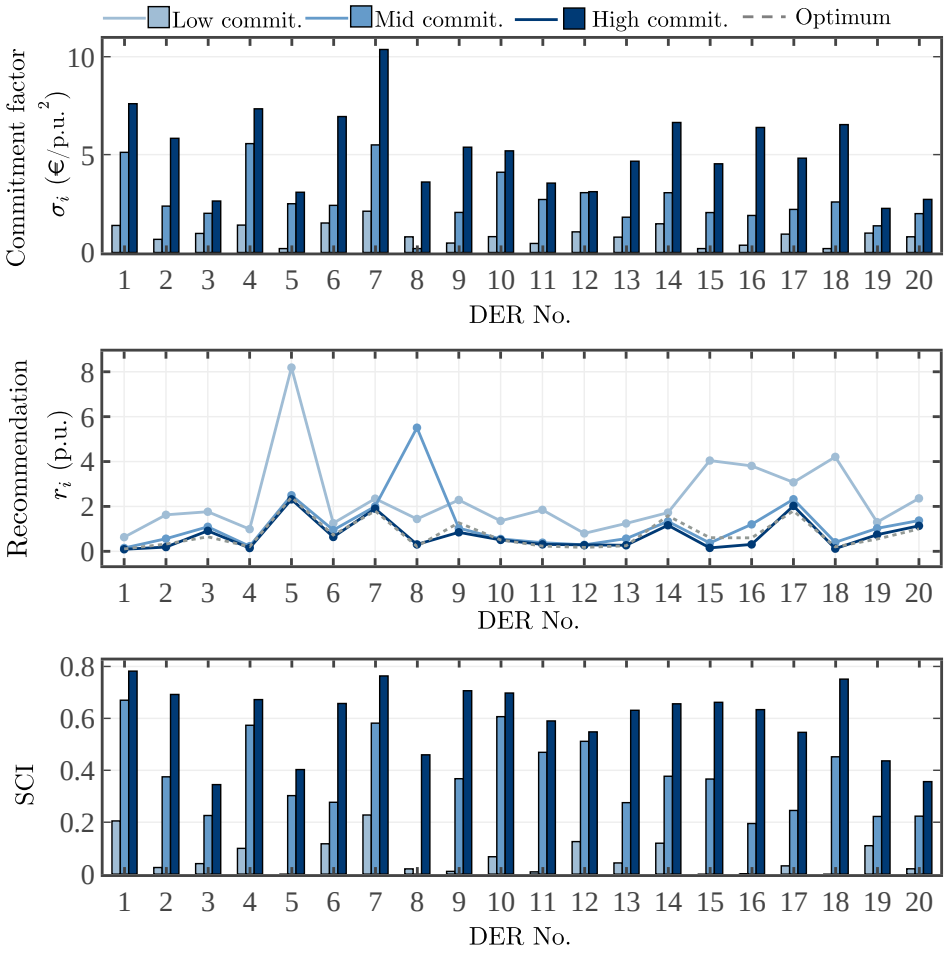

Figure 4.5: The commitment factors distribution, the customized recommendation profiles, and the corresponding SCI of DERs under three different cases.

crimination. Such a more uniform pricing structure can reduce perceived inequity among participants and may help sustain long-term engagement in the VPP.

DER Commitment Factors and Social Conformity

Fig. 4.5 reports the recommendation values and the DERs' SCI under three different levels of DER commitment factors: low ($0 < \sigma_i \leq 2$), medium ($2 < \sigma_i \leq 6$), and high ($6 < \sigma_i \leq 10$), all expressed in €/p.u.^{2}. The results exhibit a clear monotonic trend: higher commitment implies stronger conformity and less recommendation distortion. The av-

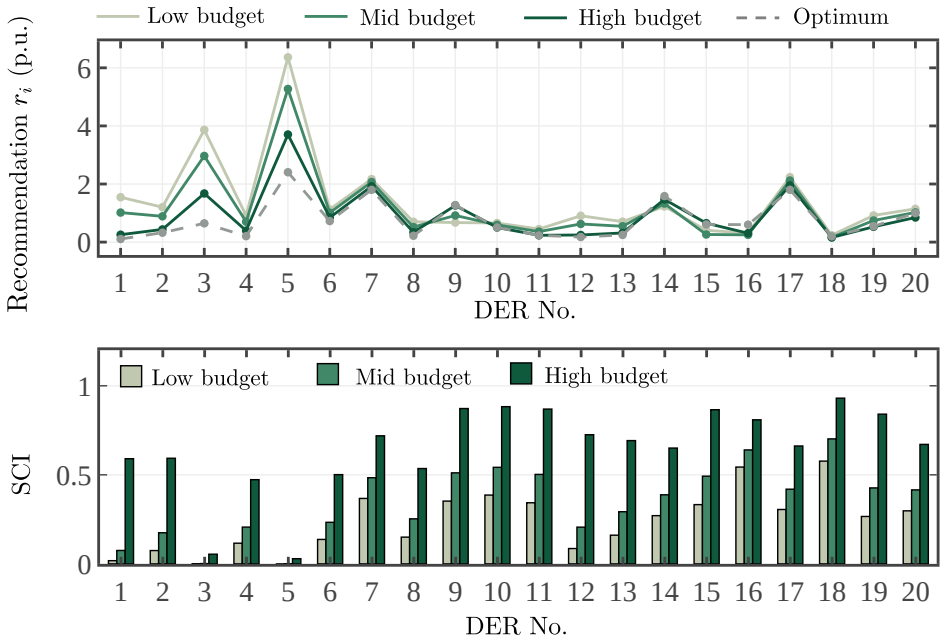

Figure 4.6: The recommendations and the SCI of DERs under different budgets.

erage SCI increases from 0.064 (low) to 0.366 (medium) and 0.599 (high). In parallel, the mean recommendation-to-optimum gap decreases from 1.554 (low) to 0.504 (medium) and 0.163 (high), indicating that the VPP operator needs less *exaggeration* in the recommendation signals to achieve effective nudging. Hence, when prosumers are more responsive to recommendations, the operator can issue recommendations closer to the true optimal allocation while still attaining reliable coordination.

Influence of Available Budget

We further examine the budget impact under three scenarios: low (50 €), medium (100 €), and high (150 €). Fig. 4.6 reports the final recommendations and the resulting SCI in each case. The mean SCI increases substantially with budget (0.239, 0.348, and 0.648 for low/medium/high budgets), indicating stronger behavioural alignment when the operator has greater financial flexibility. Meanwhile, the mean recommendation level decreases (1.388, 1.169, and 0.902), suggesting that under tight budgets the operator relies on more aggressive recommendation signals to compensate for limited monetary leverage. Therefore, although budget feasibility is guaranteed by design, overly restrictive budgets can

reduce social conformity and require stronger behavioural steering. These results also highlight SCI as a useful diagnostic indicator: it reflects how effectively the mechanism converts the available budget into behavioural alignment and can flag regimes where recommendations become excessively distorted. Finally, pushing the budget too low may weaken prosumers' willingness to participate in the coordination process; in practice, the operator should therefore set a balanced budget that controls costs while keeping DERs sufficiently motivated and engaged.

4.5.3. Effectiveness of the Data-Driven Behaviour Surrogate Modeling

This subsection evaluates the effectiveness of the proposed data-driven method for learning the behaviour surrogate models. The results in this subsection are studied on the synthetic inertia use case. The same three behavioural modes (deadband, asymmetric, bounded-rational) are assumed as in the SCC case, such that identification performance is tested under consistent behavioural heterogeneity.

Offline Identification Under Uncertainty

Table 4.1 reports the NRMSE of the *predicted responses* x produced by the identified behaviour surrogate models, comparing the proposed offline identification method with a Least-Squares (LS) baseline under three training sizes (10, 30, and 50 signal-response pairs). We evaluate identification quality via response mismatch rather than parameter error in σ_i, because the nudging mechanism uses the learned model as an *input-output map* from (λ, r) to prosumer actions. Moreover, due to projection onto X_i and saturation effects, different parameter values can generate very similar responses, so response prediction is the operationally relevant criterion.

Across all behaviour modes and different training sizes, the proposed method consistently achieves lower NRMSE than LS, indicating that it learns surrogate response models with higher predictive fidelity. Importantly, it is data-efficient: with only 10 training pairs, the average NRMSE is already modest (0.102/0.053/0.126 for DB/AS/BR), whereas LS yields substantially larger errors (0.152/0.118/0.184). Increasing the training size further improves the proposed method in a monotonic and stable manner (e.g., average NRMSE decreases from 0.102 to 0.040 to 0.029 for DB, from 0.053 to 0.030 to 0.028 for AS, and from 0.126 to 0.094 to 0.085 for BR). In contrast, LS shows limited improvement with more data and remains significantly less accurate for all modes.

This performance gap is expected. LS implicitly fits a linear regression from (λ, r)

Table 4.1: NRMSE of the offline identification results.

# Data	DER groups	Proposed method NRMSE of x			LS Method NRMSE of x		
		Min.	Avg.	Max.	Min.	Avg.	Max.
10	DB	0.051	0.102	0.181	0.127	0.152	0.189
	AS	0.033	0.053	0.091	0.068	0.118	0.161
	BR	0.098	0.126	0.183	0.132	0.184	0.222
30	DB	0.027	0.040	0.091	0.130	0.144	0.162
	AS	0.026	0.030	0.036	0.110	0.107	0.126
	BR	0.082	0.094	0.147	0.149	0.166	0.187
50	DB	0.022	0.029	0.040	0.129	0.145	0.162
	AS	0.026	0.028	0.032	0.083	0.103	0.116
	BR	0.080	0.085	0.098	0.145	0.160	0.172

to $\check{x}$, whereas the forward response induced by (4.7) is generally *piecewise affine* and becomes nonlinear when the optimizer hits bounds or other constraints become active via the projection $[\cdot]_{X_i}$. Such projection-induced nonlinearity violates the assumptions under which LS is appropriate, leading to biased fits that do not vanish with additional samples. By explicitly incorporating the constrained optimization structure (including the projection) during identification, the proposed method better captures the true input-output behaviour.

Online Adaptivity During Coordination

We next examine the benefit of *online* identification during the tâtonnement process using an RLS update. Fig. 4.7 compares three settings:

- **Case 1 (prior/nominal σ_i)**: the mechanism uses a pre-specified (prior) set of σ_i values throughout. Note that even if these values are informed by prior knowledge, they need not match the prosumers' true decision rules because σ_i is only a reduced-form parameter within the surrogate model (4.7).
- **Case 2 (offline-identified $\hat{\sigma}_i$)**: $\hat{\sigma}_i$ is identified offline from 30 signal-response pairs per DER and then kept fixed during tâtonnement.

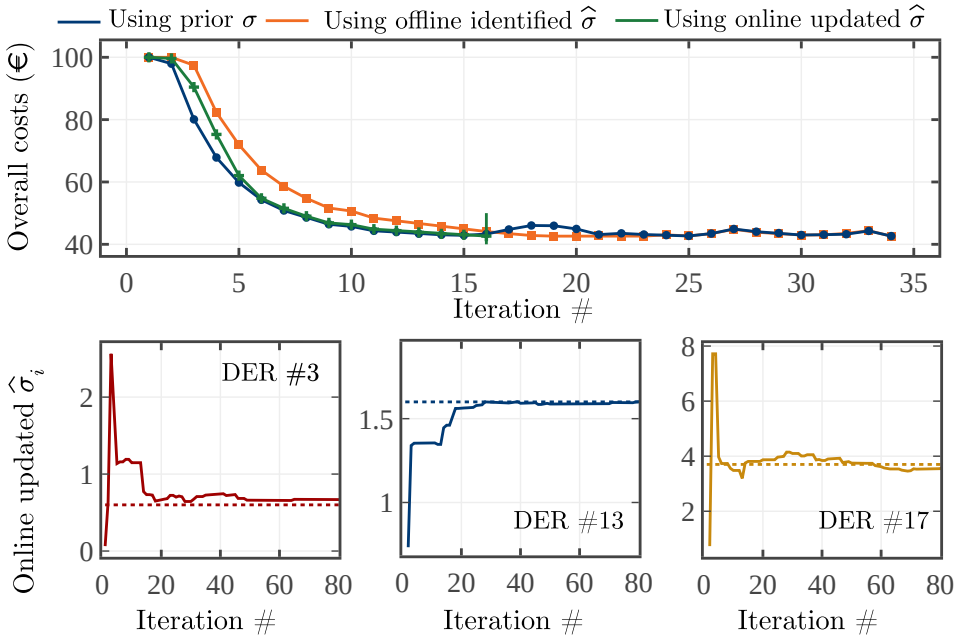

Figure 4.7: The convergence results under three cases using different ways of σ_i identification.

- **Case 3 (online-updated $\hat{\sigma}_i$)**: starting from the offline estimates in Case 2, $\hat{\sigma}_i$ is further updated online using RLS during tâtonnement.

The top panel of Fig. 4.7 plots the VPP operational cost over iterations. All three settings reach a similar steady-state cost, indicating that the nudging mechanism is robust to moderate behaviour model mismatch and imperfect parameterization. Importantly, Case 3 (online-updated $\hat{\sigma}_i$) converges *fastest*, achieving near-optimal cost in fewer iterations than both the prior/nominal initialization (Case 1) and the fixed offline model (Case 2). This faster convergence occurs because online adaptation reduces behaviour model mismatch *during* coordination, enabling the recommendations to better track the prosumers' realized responses and thus reducing transient coordination losses.

The bottom row of Fig. 4.7 shows representative RLS trajectories of $\hat{\sigma}_i$ for DERs with low, medium, and high commitment (left to right). The estimates exhibit rapid initial adjustments followed by gradual refinement toward stable values. During the early tâtonnement stage, the variability of (λ^k, r_i^k) can be limited, so full convergence of $\hat{\sigma}_i$ is not immediate; nevertheless, the partial corrections provided by RLS are already sufficient to improve recommendation quality and speed up system-level convergence. While the gain

over Case 2 is modest in this static study, the online identification becomes particularly valuable in time-varying environments where prosumer responsiveness (and thus σ_i) may drift over time.

4.5.4. Discussion

The case studies conducted on the modified IEEE 39-bus system illustrate the practical benefits and robustness of the proposed nudging mechanism and the data-driven commitment identification method. Several key observations emerge:

Nudging Mechanism Enhances Efficiency

In comparison to the baseline approach based solely on financial incentives, the proposed nudging mechanism not only guides DER behaviour toward the optimal coordination outcome, but also enables the VPP operator to consistently satisfy budget constraints. We recommend this nudging mechanism as a practical solution for VPP operators managing heterogeneous resources under both economic and operational constraints. However, it is important to emphasize that restrictive budgets can significantly reduce DERs' willingness to participate. We caution that the nudging mechanism should not be abused by VPP operators to maximize their own benefits at the expense of participant motivation, as doing so may ultimately drive DERs away and undermine long-term engagement.

Social Conformity as a Behaviour Diagnostic Tool

The SCI provides an interpretable diagnostic metric for quantifying how closely DER decisions align with the operator's coordination intent. The case studies show that SCI increases when prosumers are more responsive to recommendations and, more generally, when behavioural alignment strengthens. This makes SCI a useful indicator of DER engagement during operation and a practical signal for mechanism monitoring, enabling adaptive, feedback-based adjustments that can support sustained participation over time.

Applicability of Data-driven Behaviour Surrogate Modeling

The nudging mechanism requires a model of how DER prosumers respond to the joint incentive–recommendation signals (λ, r_i). Our data-driven behavioural surrogate modeling identifies prosumer-specific response models from observed operational data, capturing heterogeneity without relying on subjective self-declarations. This can improve transparency and trust in VPP coordination, and it can be combined with survey-based tools (e.g., questionnaires or interviews) as priors or validation checks.

Generalizability to Other Ancillary Services

Although the simulations focused on SCC support and synthetic inertia provision, the proposed mechanisms are generally applicable to a broad class of VPP-based ancillary services. The proposed behaviour-aware coordination framework is modular and can be tailored to different grid services, timescales, and resource types.

Fair compensation under heterogeneous DERs.

Fairness in the proposed framework is based on a *uniform incentive price*: DERs settled for the same service amount receive the same unit compensation. Differences in total payment arise because the settlement problem (4.1) already accounts for locational effectiveness through F_i, so more effective DERs are more likely to be selected or assigned a larger provision. However, when DER capacities differ substantially, pay-for-performance alone may not fully capture the relative effort of smaller or more constrained units. A practical extension is therefore to add an availability or reserved-capability payment on top of the performance payment. This is consistent with fairness-oriented literature on heterogeneous DER coordination [99, 100].

4.6. Conclusions

This chapter presents a behaviour-aware coordination framework for VPPs that leverages the intrinsic willingness of DER prosumers to contribute to ancillary services. The central contribution is a nudging mechanism that augments price incentives with customized recommendations and updates both signals iteratively to steer prosumer decisions toward system targets under a budget constraint. A key enabler is the proposed behavioural surrogate model, which provides a tractable input-output representation of prosumer responses to the joint signals and captures heterogeneity through a low-dimensional, learnable commitment factor. By identifying these prosumer-specific surrogate models from operational data, the VPP operator can issue more effective recommendations and reduce the need for price discrimination while maintaining coordination performance. Beyond the specific use cases explored, the proposed coordination and identification framework is modular and generalizable to other types of ancillary services. This approach provides a principled way to bridge engineering optimization with behaviour modeling, opening new avenues for scalable and human-centric control in decentralized energy systems. Future research will focus on capturing more complex behavioural dynamics and extending the approach to incorporate them.

4.7. Appendix

4.7.1. Baseline Method (Fully on Financial Incentives)

We introduce here the detailed formulation of the baseline model, which solves (4.1) using a dual-decomposition method. We first rewrite (4.1) into an equivalent consensus form:

$$\min_{x,s} \quad \sum_{i\in I} f_i(x_i) + w\left\|\sum_{i\in I} F_i s_i - x^*\right\|_1 \tag{4.20a}$$

$$\text{s.t.} \quad x_i \in X_i,\ \forall i \in I \tag{4.20b}$$

$$x_i = s_i,\ \forall i \in I. \tag{4.20c}$$

We attach dual variables u_i to each consensus equality (4.20c). The (partial) Lagrangian of (4.20) is

$$L(x, s, u) = \sum_{i\in I} f_i(x_i) + w\left\|\sum_{i\in I} F_i s_i - x^*\right\|_1 + \sum_{i\in I} u_i^{\mathsf{T}}(x_i - s_i),$$

and the dual function is $g(u) = \inf_{x,s} L(x, s, u)$. The dual problem is $\max_u g(u)$ with $u = \mathrm{col}(u_1, \dots, u_{|I|})$.

The dual decomposition method admits an *inexact* (first-order) implementation in which the primal minimizations are carried out by one projected (proximal) step per iteration. Specifically, define

$$h(s) := w\left\|\sum_{i\in I} F_i s_i - x^*\right\|_1,$$

and denote by $\Pi_{X_i}(\cdot)$ the Euclidean projection onto X_i, i.e., $\Pi_{X_i}(v) := \arg\min_{x\in X_i} \|x - v\|_2$. Moreover, denote the proximal operator of a function ϕ by

$$\mathrm{prox}_{\tau\phi}(v) := \arg\min_z \ \phi(z) + \frac{1}{2\tau}\|z - v\|_2^2.$$

Then a projected-gradient / proximal baseline is given by

$$x_i^{k+1} = \Pi_{X_i}\left(x_i^k - \alpha_x^k\left(\nabla f_i(x_i^k) + u_i^k\right)\right), \quad \forall i \in I, \tag{4.21a}$$

$$s^{k+1} = \mathrm{prox}_{\alpha_s^k h}\left(s^k + \alpha_s^k u^k\right), \tag{4.21b}$$

$$u^{k+1} = u^k + \alpha_u^k\left(x^{k+1} - s^{k+1}\right), \tag{4.21c}$$

where $\alpha_x^k, \alpha_s^k, \alpha_u^k > 0$ are stepsizes (e.g., $\alpha_\bullet^k = \alpha_\bullet/\sqrt{k+1}$), and $u^k = \mathrm{col}(u_1^k, \dots, u_{|I|}^k)$. Equation (4.21a) consists of $|I|$ independent projected-gradient steps that can be executed in parallel by the DERs. Equation (4.21b) is a proximal step associated with the nonsmooth

coupling term $h(\cdot)$, which can be computed centrally by the VPP operator. The dual update (4.21c) is a standard dual-ascent step enforcing the consensus constraints $x_i = s_i$.

The dual variables u_i^k can be interpreted as price-like signals offered by the VPP operator to DER i at iteration k of the tâtonnement process. We define $p^k := \sum_{i \in I} u_i^k$ to represent the aggregate negotiation incentive of the baseline model at iteration k.

5

Conclusions and Future Work

This chapter concludes the thesis and outlines directions for future research to further advance this work.

5.1. Conclusions

This chapter concludes the thesis by answering the three research questions stated in Chapter 1 and by highlighting how the results collectively advance the orchestration of Distributed Energy Resources (DERs) through Virtual Power Plants (VPPs) and Dynamic Virtual Power Plants (DVPPs) for the provision of ancillary services.

Answer to RQ1

RQ1 asked how the regulation task of a DVPP providing Fast Frequency Response (FFR) as a dynamic ancillary service can be formulated so that both regulation quality and economic efficiency are explicitly captured, while retaining scalability and analytical tractability as the portfolio grows. This thesis addressed this question through two complementary formulations developed in Chapters 2 and 3.

Chapter 2 introduced a convex $\mathcal{H}_\infty$-norm-based DVPP coordination framework in which FFR is modelled explicitly as a dynamic input–output behaviour, rather than as a static reserve. The aggregate FFR of multiple DER aggregators is represented by a transfer function that must track a prescribed target FFR behaviour. The mismatch between the delivered and target responses is captured by the $\mathcal{H}_\infty$ norm of a residual system, which serves as a scalar performance metric for dynamic regulation quality. Embedding this metric into a convex optimization problem with linear matrix inequality constraints enables the DVPP to jointly enforce the regulation performance management and minimise the total cost of FFR provision. In this way, the DVPP-level regulation targets, aggregator-level control parameters, and economic considerations are coherently linked within a single tractable formulation.

Chapter 3 extended the modelling perspective to a large-scale DVPP that aggregates a large population of DERs at the device level. In this setting, the DVPP operation task is again interpreted as matching a desired FFR transfer function with the aggregate transfer function of all participating DERs, but the matching is enforced directly in the frequency domain by comparing the frequency responses (i.e., the Fourier transforms of the impulse responses). By aligning real and imaginary parts of the target and aggregate responses at sampled frequencies, a dynamic-matching problem is obtained that naturally accommodates heterogeneous and variable response latencies. Treating latencies as decision variables allows the DVPP not only to allocate FFR magnitudes but also to shape the timing of the response. Although the resulting formulation is nonlinear and nonconvex, it is structured so that the complexity of the coupling constraints depends on the chosen

frequency grid rather than on the number of devices, thereby preserving scalability as the portfolio grows.

The two formulations are conceptually related, but differ in emphasis. The $\mathcal{H}_\infty$-SDP formulation in Chapter 2 can be interpreted as a worst-case frequency-domain residual minimization, since it minimizes the largest deviation between the desired and achieved responses over all frequencies. When the DER portfolio can be represented by a relatively small number of homogeneous groups, this formulation is more practical because it requires less modelling effort and less engineering tuning. By contrast, for large portfolios in which device-level heterogeneity, latency, and service-specific frequency emphasis must be retained, the dynamic-matching formulation in Chapter 3 is more appropriate, as it enforces residual matching explicitly at selected frequency points while remaining scalable across diverse device dynamic forms.

Taken together, these formulations provide a concrete answer to RQ1. They demonstrate that dynamic ancillary service provisions such as FFR can be modelled in a way that (i) captures their inherently dynamic nature, (ii) offers explicit, quantitative measures of the DVPP regulation performance, (iii) integrates economic efficiency into the same DVPP operation framework, and (iv) remains amenable to scalable analysis and computation, even when device-level latency flexibility is explicitly represented.

Answer to RQ2

RQ2 concerned the design of distributed solution methods that can efficiently compute high-quality DVPP operating profiles for dynamic ancillary services, without revealing detailed local information and without relying on single points of failure, even when the underlying optimization problems are highly nonconvex and large-scale.

Building on the convex $\mathcal{H}_\infty$-based formulation of Chapter 2, the thesis first developed a game-theoretic learning approach at the level of DER aggregators. The centralized optimization problem, which includes a global coupling linear matrix inequality constraint encoding the $\mathcal{H}_\infty$ performance requirement, is reformulated as a state-based potential game. This is achieved by introducing local estimator states that replicate the effect of the shared semidefinite constraint in a decentralised manner. The resulting game admits stationary Nash equilibria that are shown to recover, or closely approximate, the optimal solution of the original convex programme. A distributed learning algorithm is then designed in which aggregators iteratively update their FFR provisions using only local information and neighbour-to-neighbour communication. No aggregator needs to reveal its detailed cost function or operational limits to a central coordinator or to other aggrega-

tors. Theoretical analysis and numerical case studies confirm that the algorithm converges under asynchronous updates and time-varying communication graphs, thereby providing a fully distributed and practically implementable solution method for the DVPP coordination problem with limited disclosure of local information.

Chapter 3 addressed the second part of RQ2 in the context of the nonlinear, nonconvex dynamic-matching problem arising when device-level latencies are treated as decision variables. A dual-decomposition-based distributed method was proposed that exploits the separable structure of the DVPP problem. By introducing a small set of dual variables associated with the frequency-domain matching constraints, the large-scale optimization task is decomposed into low-dimensional device-level subproblems, each solved locally by an individual DER. The coupling between devices occurs only through the dual variables, which are updated via simple projected subgradient steps. This decomposition admits parallel implementation across hundreds or thousands of devices and preserves the interpretability of the dual variables as coordination signals. Theoretical analysis and numerical results demonstrate that the method scales gracefully with portfolio size, achieves close tracking of the target FFR, and yields operating points that are near-optimal relative to the ground-truth optimisers, while remaining computationally tractable.

Taken together, the contributions in Chapters 2 and 3 show that the formulations developed for RQ1 can indeed be solved by distributed methods that support data confidentiality, avoid central bottlenecks, and scale to large DVPP portfolios. Game-theoretic learning and dual decomposition emerge as complementary approaches: the former is particularly developed for convex, LMI-constrained settings, while the latter addresses large-scale nonconvex coordination problems. Both methods confirm that high-quality DVPP operation for dynamic ancillary services is achievable with decentralised computation and limited coordination signals.

Answer to RQ3

RQ3 asked how the intrinsic willingness of DER prosumers can be systematically leveraged to enhance the efficiency and reliability of VPP operation, and how prosumer behaviour should be modelled so that incentive mechanisms align local decisions with ancillary service provision.

Chapter 4 addressed this question by introducing a behavioural layer on top of the technical coordination mechanisms. Each DER owner is modelled as an agent endowed with a *commitment factor* that quantifies how strongly they tend to follow recommended service levels under given incentive prices, reflecting intrinsic motivations and hetero-

geneity across prosumers. Building on this model, the chapter proposed a nudging mechanism based on a tâtonnement-type process, in which the VPP operator iteratively adjusts a common incentive price and individual service recommendations, while prosumers update their decisions by trading off local cost against deviations from the recommended levels. The mechanism is designed so that, under convexity assumptions, the aggregate outcome converges to the optimal solution of a VPP coordination problem with an explicit budget constraint, thereby translating intrinsic behavioural motivations into operational benefits while maintaining budget balance.

To make the nudging mechanism practically deployable, a data-driven inverse optimization framework was developed to identify commitment factors from observed decisions. The approach combines an offline estimation stage, using historical or simulated data, with an online update that adapts to evolving behaviour. This allows the VPP to initialise and refine behavioural models without requiring intrusive surveys or full disclosure of preferences. In addition, a social conformity index was introduced as a diagnostic metric that quantifies how closely realised prosumer responses follow the issued recommendations, providing interpretable feedback on the effectiveness of the nudging mechanism and on the degree of alignment between technical coordination goals and actual behaviour. Case studies on synthetic inertia provision and short-circuit current support show that the proposed mechanism can steer heterogeneous prosumers towards desired operating points, maintain budget compliance, and handle diverse commitment levels more effectively than purely price-based schemes.

These results collectively answer RQ3 by demonstrating that intrinsic prosumer motivations can be modelled and harnessed systematically within VPP operation. By integrating commitment-aware behavioural models, data-driven identification, and recommendation based nudges, the thesis offers a cohesive framework in which human and technical dimensions of ancillary-service provision reinforce rather than contradict each other.

Overall Reflection

Across all three research questions, a coherent picture emerges. The thesis establishes a unified VPP/DVPP operation framework that connects system-level stability requirements, portfolio-level coordination objectives, device-level capabilities, and economic considerations within a single decision-making layer. The formulations developed for FFR demonstrate how dynamic ancillary-service requirements can be translated into optimization problems with a clear, computable regulation-quality metric and practical computational complexity. Distributed solution methods based on game-theoretic learning and

dual decomposition show that such formulations can be implemented at scale while limiting the disclosure of local information and avoiding central bottlenecks. Finally, incorporating prosumer behaviour through commitment-aware nudging mechanisms highlights that sustainable and efficient VPP operation must engage with intrinsic motivations and behavioural diversity, rather than relying solely on financial incentives and centralised control. Together, these findings advance the understanding of how VPPs and DVPPs can orchestrate large populations of DERs to provide reliable, scalable, and socially grounded ancillary services in future low-inertia power systems.

5.2. Future Work

Future research on VPPs and DVPPs can build on this thesis along five closely related directions: richer dynamic modelling of ancillary services and DER capabilities, tighter integration with grid operation and markets, deeper behavioural modelling and nudging design, more advanced distributed optimization and learning methods, and implementation-oriented validation.

A first direction is to enrich the dynamic models used to represent ancillary-service delivery and to refine the characterisation of how heterogeneous DERs provide these services, especially when multiple services must be delivered simultaneously. The linear time-invariant, frequency-domain representations adopted in this thesis provide a clean abstraction for coordination, but future work should extend them to capture nonlinearities, saturation and dead-zones, protection actions, and mode transitions of grid-following and grid-forming inverters. Time-varying behaviour and time-scale dependencies are also important, for example when state-of-charge constraints create coupling across events, when inverter control modes change with voltage or frequency conditions, or when thermal limits evolve over minutes. Moreover, uncertainty that is weather-dependent or condition-dependent (e.g., irradiance, wind, temperature, availability, and communication latency) should be explicitly represented in the service models and feasibility sets. Data-driven identification and surrogate modelling offer a promising complement to physics-based modelling, particularly for capturing aggregate DVPP dynamics at scale [101, 102]; however, to be useful for optimization and market clearing, such models should be combined with robustness guarantees or certified error bounds to ensure safety and stability under model mismatch. Incorporating explicit network and protection constraints, as well as location-dependent effects, would further tighten the link between DVPP operation and system-wide stability.

A second direction concerns the interaction between VPPs and the power-system operation and market layers. A central practical question is what information a VPP needs from the TSO/DSO (and what information it can provide in return) in order to reliably deliver dynamic services while respecting network constraints. This includes how network and congestion constraints should be embedded into VPP scheduling, what level of grid observability and data exchange is required (e.g., grid strength and short-circuit levels, voltage sensitivities, protection settings, and local operating limits), and how such requirements can be made compatible with existing regulatory rules and operational procedures. If a VPP participates directly in ancillary-service markets, the coordination problem becomes intertwined with bidding, clearing, and settlement. This motivates ancillary-service product definitions that go beyond capacity and explicitly reflect dynamic performance characteristics, together with practical methodologies for measurement and verification. In parallel, the VPP must quantify and communicate its deliverable capability under uncertainty, for example through service feasibility regions and capability envelopes that account for network dependence, operating-point dependence, and stochastic availability. This naturally suggests a "dual" perspective of coordination: rather than dispatching DERs given a system request, the VPP must infer or approximate aggregate service envelopes from local DER information, limited measurements, and privacy constraints [103, 104]. Future work could therefore develop principled methods for capability estimation and aggregation, e.g., certified inner/outer approximations of feasibility sets, performance-constrained capability envelopes, and location-aware capability maps, and investigate how these abstractions can be integrated into market design so that DER-based VPPs can participate without excessive conservatism or administrative burden.

A third direction is to deepen the socio-technical layer by developing richer behavioural models of DER prosumers and corresponding nudging mechanisms. The commitment-factor model and recommendation-based nudging scheme proposed in this thesis provide a starting point, but future work could consider time-varying and multi-dimensional behavioural parameters that distinguish, for example, responsiveness, risk aversion, fairness perceptions, and fatigue or disengagement over long horizons. Advancing this line of work requires empirical grounding through field pilots and realistic deployment data, enabling identification and validation beyond simulated preferences. Data-driven inverse optimisation remains a particularly suitable tool because it links observed decisions to interpretable behavioural parameters and supports both offline calibration and online adaptation [95, 105]. In parallel, different families of nudges, such as social-norm feedback, goal-setting, default options, or commitment contracts, could be designed and evaluated

within VPP programmes on top of these behavioural models [38, 106]. Such mechanisms should emphasise transparency, fairness, and long-term engagement, while incorporating safeguards that discourage undesirable strategic behaviour (e.g., enrolling but under-delivering, manipulating reference baselines used for performance verification, or responding only in favourable conditions in ways that reduce aggregate reliability).

A fourth direction is algorithmic: strengthening the distributed optimisation and learning methods that underpin VPP coordination under realistic operational conditions. As VPPs scale and diversify, privacy becomes increasingly important; future work could therefore integrate privacy-preserving techniques into distributed coordination, ranging from secure aggregation and cryptographic protocols [107] to differential-privacy mechanisms [108, 109] that limit information leakage while maintaining acceptable regulation performance. Meanwhile, intermittency, forecast updates, and changing DER availability call for coordination schemes that can respond quickly to evolving conditions, converge efficiently to updated operating points, and remain robust to noisy online measurements. This motivates asynchronous, event-triggered, and gossip-based variants of the distributed algorithms [110], as well as uncertainty-aware coordination methods that explicitly capture uncertainty in deliverable service capability and system conditions while providing probabilistic or worst-case performance guarantees. A particularly important next step is to test the framework in a genuine online optimisation setting with continuously updated forecasts and measurements, so that the distributed coordination algorithm can track changing operating conditions in closed loop.

A fifth direction is implementation and validation. Translating the analytical and simulation-based results of this thesis into practice requires real-time experimentation with physical controllers and communication constraints. Hardware-in-the-loop and real-time simulation platforms can be used to validate VPP coordination strategies, including dynamic service tracking and distributed algorithms, under realistic latency, packet loss, and measurement limitations [111]. Digital-twin architectures for VPPs [112], combining detailed device models with live measurements, would enable continuous calibration, predictive analysis, and safe testing of new services and nudging schemes before field deployment. Alongside these efforts, cyber-security, privacy, and interoperability standards will be essential to ensure that VPPs can operate as trusted and resilient components of future low-inertia power systems.

Bibliography

[1] EU. *Regulation (EU) 2021/1119 of the European Parliament and of the Council*. 2021.

[2] United States Government. *The Long-Term Strategy of the United States: Pathways to Net-Zero Greenhouse Gas Emissions by 2050*. 2021.

[3] Internationale Energieagentur. *An Energy Sector Roadmap to Carbon Neutrality in China*. Tech. rep. 2021.

[4] UK Government. *United Kingdom of Great Britain and Northern Ireland's 2035 Nationally Determined Contribution*. 2025.

[5] Baraa Mohandes et al. "A Review of Power System Flexibility with High Penetration of Renewables". In: *IEEE Transactions on Power Systems* 34.4 (2019), pp. 3140–3155.

[6] Md Motinur Rahman et al. "An Overview of Power System Flexibility: High Renewable Energy Penetration Scenarios". In: *Energies* 17.24 (2024), p. 6393.

[7] Rafat Aljarrah et al. "Issues and Challenges of Grid-Following Converters Interfacing Renewable Energy Sources in Low Inertia Systems: A Review". In: *IEEE Access* 12 (2024), pp. 5534–5561.

[8] Dohyuk Kim et al. "Evaluating Influence of Inverter-based Resources on System Strength Considering Inverter Interaction Level". In: *Sustainability* 12.8 (2020), p. 3469.

[9] Nikos Hatziargyriou et al. "Definition and Classification of Power System Stability–Revisited & Extended". In: *IEEE Transactions on Power Systems* 36.4 (2021), pp. 3271–3281.

[10] Comité 28-A. *Informe no confidencial del Comité de Análisis de las circunstancias que concurrieron en la crisis de electricidad del 28 de abril de 2025*. Tech. rep. Madrid, Spain: Ministerio para la Transición Ecológica y el Reto Demográfico, June 2025.

[11] ENTSO-E Expert Panel. *Factual Report on the Grid Incident in Spain and Portugal on 28 April 2025*. Tech. rep. Brussels, Belgium: ENTSO-E, Oct. 2025.

[12] Jose Daniel Lara et al. *April 28th 2025 Iberian Blackout: Analysis of Available Information*. Presentation NREL/PR-6A40-95103. Golden, CO, USA: National Renewable Energy Laboratory, 2025.

[13] C. MacIver, J. V. Milanovic, and Z. Popovic. "An Analysis of the August 9th 2019 GB Transmission System Frequency Incident". In: *Electric Power Systems Research* 190 (2021), p. 106704.

[14] Energy Emergencies Executive Committee (E3C). *GB Power System Disruption on 9 August 2019: Final Report*. Tech. rep. London, UK: Department for Business, Energy and Industrial Strategy, Jan. 2020.

[15] Australian Energy Market Operator. *Black System South Australia 28 September 2016 – Final Report*. Tech. rep. Melbourne, Australia: AEMO, 2017.

[16] Bogdan Marinescu et al. "Dynamic Virtual Power Plant: A New Concept for Grid Integration of Renewable Energy Sources". In: *IEEE Access* 10 (2022), pp. 104980–104995.

[17] Francesco Gulotta et al. "Opening of ancillary service markets to distributed energy resources: A review". In: *Energies* 16.6 (2023), p. 2814.

[18] Hossam HH Mousa, Karar Mahmoud, and Matti Lehtonen. "Recent developments of demand-side management towards flexible DER-rich power systems: A systematic review". In: *IET Generation, Transmission & Distribution* 18.13 (2024), pp. 2259–2300.

[19] Michael Z. Liu et al. "Using OPF-Based Operating Envelopes to Facilitate Residential DER Services". In: *IEEE Transactions on Smart Grid* 13.6 (2022), pp. 4494–4504.

[20] Joakim Björk, Karl Henrik Johansson, and Florian Dörfler. "Dynamic Virtual Power Plant Design for Fast Frequency Reserves: Coordinating Hydro and Wind". In: *IEEE Transactions on Control of Network Systems* 10.3 (2023), pp. 1266–1278.

[21] Junru Chen, Muyang Liu, and Federico Milano. "Aggregated Model of Virtual Power Plants for Transient Frequency and Voltage Stability Analysis". In: *IEEE Transactions on Power Systems* 36.5 (2021), pp. 4366–4375.

[22] Zahid Ullah, Arshad Arshad, and Azam Nekahi. "Virtual Power Plants: challenges, opportunities, and profitability assessment in current energy markets". In: *Electricity* 5.2 (2024), pp. 370–384.

[23] Zhi-Yu Xu et al. “Virtual power plant-based pricing control for wind/thermal co-operated generation in China”. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 46.5 (2015), pp. 706–712.

[24] Ahmed Hany Elgamal et al. “Optimization of a multiple-scale renewable energy-based virtual power plant in the UK”. In: *Applied Energy* 256 (2019), p. 113973.

[25] Xiao Liu, Sinan Li, and Jianguo Zhu. “Optimal coordination for multiple network-constrained VPPs via multi-agent deep reinforcement learning”. In: *IEEE Transactions on Smart Grid* 14.4 (2022), pp. 3016–3031.

[26] Zhuocen Dai et al. “Massive coordination of distributed energy resources in vpp: A mean field rl-based bi-level optimization approach”. In: *IEEE Transactions on Cybernetics* 55.3 (2025), pp. 1332–1346.

[27] Saeed Rahmani Dabbagh and Mohammad Kazem Sheikh-El-Eslami. “Risk assessment of virtual power plants offering in energy and reserve markets”. In: *IEEE Transactions on Power Systems* 31.5 (2015), pp. 3572–3582.

[28] Subir Majumder et al. “Chance-Constrained Pre-Contingency Joint Self- Scheduling of Energy and Reserve in VPP”. In: *IEEE Transactions on Power Systems* 39.1 (2024), pp. 245–260.

[29] Han Wang et al. “Optimal virtual power plant operational regime under reserve uncertainty”. In: *IEEE Transactions on Smart Grid* 13.4 (2022), pp. 2973–2985.

[30] Bogdan Marinescu, Florian Dörfler, and Federico Milano. “Dynamic Virtual Power Plants: A New Paradigm for Grid Services and Control”. In: *IEEE Power and Energy Magazine* 23.6 (2025), pp. 81–96.

[31] Navid Vafamand et al. “Design of Dynamic Virtual Power Plants based on Model Predictive Controller”. In: *2025 IEEE International Conference on Environment and Electrical Engineering and 2025 IEEE Industrial and Commercial Power Systems Europe (EEEIC/I&CPS Europe)*. IEEE. 2025, pp. 1–6.

[32] Hêmin Golpıra and Bogdan Marinescu. “Enhanced frequency regulation scheme: An online paradigm for dynamic virtual power plant integration”. In: *IEEE Transactions on Power Systems* 39.6 (2024), pp. 7227–7239.

[33] Xinxin Ge, Ge Wang, and Fei Wang. “Optimal day-ahead dispatching strategy for VPP using dynamic grouping DER aggregation method based on improved virtual battery model”. In: *Energy* 326 (2025), p. 136166.

[34] Wafa Nafkha-Tayari et al. “Virtual power plants optimization issue: A comprehensive review on methods, solutions, and prospects”. In: *Energies* 15.10 (2022), p. 3607.

[35] Verena Häberle et al. “Grid-forming and spatially distributed control design of dynamic virtual power plants”. In: *IEEE Transactions on Smart Grid* 15.2 (2023), pp. 1761–1777.

[36] Ali Abbasi et al. “Optimizing virtual power plants with parallel simulated annealing on high-performance computing”. In: *Smart Cities* 8.2 (2025), p. 47.

[37] Cheng Feng et al. “Update scheduling for ADMM-based energy sharing in virtual power plants considering massive prosumer access”. In: *IEEE Transactions on Smart Grid* 14.5 (2023), pp. 3961–3975.

[38] Richard H Thaler and Cass R Sunstein. *Nudge: Improving decisions about health, wealth, and happiness*. Penguin, 2009.

[39] Tadeusz Skoczkowski et al. “Participation in demand side response. Are individual energy users interested in this?” In: *Renewable Energy* 232 (2024), p. 121104.

[40] Bryony Parrish et al. “A systematic review of motivations, enablers and barriers for consumer engagement with residential demand response”. In: *Energy Policy* 138 (2020), p. 111221.

[41] Haiwei Xie and Jochen L Cremer. “Game-theoretic learning for power system dynamic ancillary service provisions”. In: *IEEE Control Systems Letters* 8 (2024), pp. 1307–1312.

[42] Verena Häberle et al. “Control Design of Dynamic Virtual Power Plants: An Adaptive Divide-and-Conquer Approach”. In: *IEEE Transactions on Power Systems* 37.5 (2022), pp. 4040–4053.

[43] Luis Badesa, Fei Teng, and Goran Strbac. “Simultaneous scheduling of multiple frequency services in stochastic unit commitment”. In: *IEEE Transactions on Power Systems* 34.5 (2019), pp. 3858–3868.

[44] Lexuan Meng et al. “Fast frequency response from energy storage systems—a review of grid standards, projects and technical issues”. In: *IEEE transactions on smart grid* 11.2 (2019), pp. 1566–1581.

[45] Fei Teng, Vincenzo Trovato, and Goran Strbac. “Stochastic scheduling with inertia-dependent fast frequency response requirements”. In: *IEEE Transactions on Power Systems* 31.2 (2015), pp. 1557–1566.

[46] Zhongda Chu, Guoxuan Cui, and Fei Teng. "Scheduling of software-defined microgrids for optimal frequency regulation". In: *IEEE Transactions on Sustainable Energy* 15.3 (2024), pp. 1715–1728.

[47] Cheng Feng et al. "Provision of Contingency Frequency Services for Virtual Power Plants with Aggregated Models". In: *IEEE Transactions on Smart Grid* 14.4 (July 2023), pp. 2798–2811.

[48] Verena Häberle et al. "Grid-Forming and Spatially Distributed Control Design of Dynamic Virtual Power Plants". In: *IEEE Transactions on Smart Grid* 15.2 (2024), pp. 1761–1777.

[49] Hans-Peter Beck and Ralf Hesse. "Virtual synchronous machine". In: *2007 9th international conference on electrical power quality and utilisation*. IEEE. 2007, pp. 1–6.

[50] Mohammad Ebrahimi, S Ali Khajehoddin, and Masoud Karimi-Ghartemani. "An improved damping method for virtual synchronous machines". In: *IEEE Transactions on Sustainable Energy* 10.3 (2019), pp. 1491–1500.

[51] Verena Häberle et al. "Dynamic ancillary services: From grid codes to transfer function-based converter control". In: *Electric Power Systems Research* 234 (2024), p. 110760.

[52] Z Khalik et al. "Vehicle energy management with ecodriving: A sequential quadratic programming approach with dual decomposition". In: *2018 Annual American Control Conference (ACC)*. IEEE. 2018, pp. 4002–4007.

[53] Na Li and Jason R. Marden. "Designing Games to Handle Coupled Constraints". In: *49th IEEE Conference on Decision and Control (CDC)*. 2010, pp. 250–255.

[54] Yile Liang, Feng Liu, and Shengwei Mei. "Distributed real-time economic dispatch in smart grids: A state-based potential game approach". In: *IEEE Transactions on Smart Grid* 9.5 (2017), pp. 4194–4208.

[55] Na Li and Jason R. Marden. "Decoupling Coupled Constraints Through Utility Design". In: *IEEE Transactions on Automatic Control* 59.8 (2014), pp. 2289–2294.

[56] Yukang Shen et al. "Optimal Allocation of Virtual Inertia and Droop Control for Renewable Energy in Stochastic Look-Ahead Power Dispatch". In: *IEEE Transactions on Sustainable Energy* 14.3 (2023), pp. 1881–1894.

[57] Kemin Zhou and John C. Doyle. *Essentials of Robust Control*. Upper Saddle River, NJ: Prentice Hall, 1998.

[58] Jeff S Shamma. "Game theory, learning, and control systems". In: *National Science Review* 7.7 (2020), pp. 1118–1119.

[59] Jason R Marden. "State based potential games". In: *Automatica* 48.12 (2012), pp. 3075–3088.

[60] Na Li and Jason R Marden. "Designing games for distributed optimization". In: *IEEE Journal of Selected Topics in Signal Processing* 7.2 (2013), pp. 230–242.

[61] Stephen Boyd et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers". In: *Foundations and Trends® in Machine learning* 3.1 (2011), pp. 1–122.

[62] Ali Q Al-Shetwi. "Sustainable development of renewable energy integrated power sector: Trends, environmental impacts, and recent challenges". In: *Science of The Total Environment* 822 (2022), p. 153645.

[63] Abhimanyu Kaushal and Dirk Van Hertem. "An overview of ancillary services and HVDC systems in European context". In: *Energies* 12.18 (2019), p. 3481.

[64] Chengrong Lin et al. "Performance optimization of VPP in fast frequency control ancillary service provision". In: *Applied Energy* 376 (2024), p. 124294.

[65] Hedayat Saboori, M Mohammadi, and R Taghe. "Virtual power plant (VPP), definition, concept, components and types". In: *2011 Asia-Pacific power and energy engineering conference*. IEEE. 2011, pp. 1–4.

[66] Robert Eriksson, Niklas Modig, and Katherine Elkington. "Synthetic inertia versus fast frequency response: a definition". In: *IET renewable power generation* 12.5 (2018), pp. 507–514.

[67] Guo Chen and Jueyou Li. "A fully distributed ADMM-based dispatch approach for virtual power plant problems". In: *Applied Mathematical Modelling* 58 (2018), pp. 300–312.

[68] Chu Sun et al. "Virtual synchronous machine control for low-inertia power system considering energy storage limitation". In: *2019 IEEE Energy Conversion Congress and Exposition (ECCE)*. IEEE. 2019, pp. 6021–6028.

[69] Jialun Zhang et al. "Pricing framework for activation latency of fast frequency response provided by VPP". In: *IEEE Transactions on Energy Markets, Policy and Regulation* (2025), pp. 1–13.

[70] Jinrui Guo et al. “Co-Design of Power Dispatch with Dynamic Power Regulation and Communication Transmission Optimization for Frequency Control in VPPs”. In: *Journal of Modern Power Systems and Clean Energy* 13.4 (2025), pp. 1383–1394.

[71] Smriti Jaiswal, Mausri Bhuyan, and Dulal Chandra Das. “Impact of Communication Delay in a Coordinated Control VPP Model with Demand Side Flexibility: A Case Study”. In: *International Conference on Modeling, Simulation and Optimization*. Springer. 2022, pp. 431–443.

[72] Robin Vujanic et al. “A decomposition method for large scale MILPs, with performance guarantees and a power system application”. In: *Automatica* 67 (2016), pp. 144–156.

[73] Dimitri Bertsekas. *Convex optimization theory*. Vol. 1. Athena Scientific, 2009.

[74] MathWorks. *Optimization Toolbox: User’s Guide (Function `fmincon`)*. The MathWorks, Inc. Natick, MA, USA, 2024.

[75] Eric Hirst and Brendan Kirby. *Electric-power ancillary services*. Oak Ridge National Laboratory Oak Ridge, TN, USA, 1996.

[76] Enrique Lobato Miguélez et al. “An overview of ancillary services in Spain”. In: *Electric Power Systems Research* 78.3 (2008), pp. 515–523.

[77] Kankar Bhattacharya and Jin Zhong. “Reactive power as an ancillary service”. In: *IEEE Transactions on Power systems* 16.2 (2001), pp. 294–300.

[78] Brendan Kirby and Eric Hirst. *Ancillary service details: Voltage control*. Tech. rep. Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 1997.

[79] Andreas Kubis, S Ruberg, and Christian Rehtanz. “Development of available short-circuit power in Germany from 2011 up to 2033”. In: *CIRED workshop, Rome, Italy*. 2014, pp. 11–12.

[80] AS Chuang and C Schwaegerl. “Ancillary services for renewable integration”. In: *2009 CIGRE/IEEE PES Joint Symposium Integration of Wide-Scale Renewable Resources Into the Power Delivery System*. IEEE. 2009, pp. 1–1.

[81] Dawei Qiu et al. “Market design for ancillary service provisions of inertia and frequency response via virtual power plants: A non-convex bi-level optimisation approach”. In: *Applied Energy* 361 (2024), p. 122929.

[82] Emiliano Dall’Anese et al. “Optimal regulation of virtual power plants”. In: *IEEE transactions on power systems* 33.2 (2017), pp. 1868–1881.

[83] Pio Lombardi, Michal Powalko, and Krzysztof Rudion. "Optimal operation of a virtual power plant". In: *2009 IEEE power & energy society general meeting*. IEEE. 2009, pp. 1–6.

[84] Moein Esfahani et al. "A distributed VPP-integrated co-optimization framework for energy scheduling, frequency regulation, and voltage support using data-driven distributionally robust optimization with Wasserstein metric". In: *Applied Energy* 361 (2024), p. 122883.

[85] Tine L Vandoorn et al. "Smart microgrids and virtual power plants in a hierarchical control structure". In: *2011 2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies*. IEEE. 2011, pp. 1–7.

[86] Zizheng Ren, Gregor Verbič, and Jaysson Guerrero. "Multi-period dynamic tariffs for prosumers participating in virtual power plants". In: *Electric Power Systems Research* 212 (2022), p. 108478.

[87] Alfiya Kulmukhanova et al. "Mechanism design for virtual power plant with independent distributed generators". In: *IFAC-PapersOnLine* 52.4 (2019), pp. 419–424.

[88] Sambeet Mishra et al. "Smart contract formation enabling energy-as-a-service in a virtual power plant". In: *International journal of energy research* 46.3 (2022), pp. 3272–3294.

[89] Junjie Hu et al. "Transactive control: a framework for operating power systems characterized by high penetration of distributed energy resources". In: *Journal of Modern Power Systems and Clean Energy* 5.3 (2017), pp. 451–464.

[90] Swantje Gährs et al. "Acceptance of Ancillary Services and Willingness to Invest in PV-storage-systems". In: *Energy Procedia* 73 (2015), pp. 29–36.

[91] Bryony Parrish, Phil Heptonstall, and Rob Gross. "The potential for UK residential demand side participation". In: *System Architecture Challenges: Supergen+ for HubNet, HubNet* (2016).

[92] Sen Zhan et al. "Fairness-incorporated online feedback optimization for real-time distribution grid management". In: *IEEE Transactions on Smart Grid* 15.2 (2023), pp. 1792–1806.

[93] Mehran Shakarami, Ashish Cherukuri, and Nima Monshizadeh. "Steering the aggregative behavior of noncooperative agents: a nudge framework". In: *Automatica* 136 (2022), p. 110003.

[94] Ting Yang et al. “A privacy-preserving federated reinforcement learning method for multiple virtual power plants scheduling”. In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 72.4 (2024), pp. 1939–1950.

[95] Timothy CY Chan, Rafid Mahmood, and Ian Yihang Zhu. “Inverse optimization: Theory and applications”. In: *Operations Research* 73.2 (2025), pp. 1046–1074.

[96] Anil Aswani, Zuo-Jun Shen, and Auyon Siddiq. “Inverse optimization with noisy data”. In: *Operations Research* 66.3 (2018), pp. 870–892.

[97] Syed Aseem Ul Islam and Dennis S Bernstein. “Recursive least squares for real-time implementation [lecture notes]”. In: *IEEE Control Systems Magazine* 39.3 (2019), pp. 82–85.

[98] Zhongda Chu, Jingyi Wu, and Fei Teng. “Pricing of short circuit current in high IBR-penetrated system”. In: *Electric Power Systems Research* 235 (2024), p. 110690.

[99] Michael Z. Liu et al. “On the Fairness of PV Curtailment Schemes in Residential Distribution Networks”. In: *IEEE Transactions on Smart Grid* 11.5 (2020), pp. 4502–4512.

[100] Shiva Poudel et al. “Fairness-Aware Distributed Energy Coordination for Voltage Regulation in Power Distribution Systems”. In: *IEEE Transactions on Sustainable Energy* 14.3 (2023), pp. 1866–1880.

[101] Ke Yang et al. “Data-driven dynamic modeling for inverter-based resources using neural networks”. In: *Nature Communications* 16 (2025).

[102] Lingling Fan et al. “Data-driven dynamic modeling in power systems: A fresh look on inverter-based resource modeling”. In: *IEEE Power and Energy Magazine* 20.3 (2022), pp. 64–76.

[103] Lin Chen et al. “Feasible operation region estimation of virtual power plant considering heterogeneity and uncertainty of distributed energy resources”. In: *Applied Energy* 362 (2024), p. 123000.

[104] Soumya Kundu, Karanjit Kalsi, and Scott Backhaus. “Approximating flexibility in distributed energy resources: A geometric approach”. In: *2018 Power Systems Computation Conference (PSCC)*. IEEE. 2018, pp. 1–7.

[105] Chaosheng Dong, Yiran Chen, and Bo Zeng. “Generalized inverse optimization through online learning”. In: *Advances in Neural Information Processing Systems* 31 (2018).

[106] Mark A Andor and Katja M Fels. "Behavioral economics and energy conservation–a systematic review of non-price interventions and their causal effects". In: *Ecological economics* 148 (2018), pp. 178–210.

[107] Katrine Tjell and Rafael Wisniewski. "Privacy preservation in distributed optimization via dual decomposition and ADMM". In: *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE. 2019, pp. 7203–7208.

[108] Yu Wang et al. "Differential privacy in linear distributed control systems: Entropy minimizing mechanisms and performance tradeoffs". In: *IEEE Transactions on Control of Network Systems* 4.1 (2017), pp. 118–130.

[109] Zonghao Huang et al. "DP-ADMM: ADMM-based distributed learning with differential privacy". In: *IEEE Transactions on Information Forensics and Security* 15 (2019), pp. 1002–1012.

[110] Shuoguang Yang, Xuezhou Zhang, and Mengdi Wang. "Decentralized gossip-based stochastic bilevel optimization over communication networks". In: *Advances in neural information processing systems* 35 (2022), pp. 238–252.

[111] R Kuffel et al. "RTDS-a fully digital power system simulator operating in real time". In: *Proceedings 1995 international conference on energy management and power delivery EMPD'95*. Vol. 2. IEEE. 1995, pp. 498–503.

[112] Pedro P Vergara et al. "Virtual Power Plant Digital Twins: Ensuring Seamless Deployment via Standard Architectures". In: *IEEE Power and Energy Magazine* 23.6 (2025), pp. 97–110.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my promoters, Prof. Peter Palensky and Dr. Jochen L. Cremer. Peter, as my promoter, has been a role model for me in academia. He has continuously supported me throughout this PhD journey at TU Delft. I am especially grateful for his professionalism, efficiency, and vision in research. Under his leadership, our group has grown into a strong and collaborative team, and it has been a privilege to carry out my work in such an inspiring environment.

Jochen, as my promoter and daily supervisor, is someone I feel very fortunate to know and to work with. Our collaboration began long ago during our time at Imperial College, and I would not have had the courage to embark on this PhD journey without his encouragement and confidence in me. His insightful feedback, patience, and constant support have been invaluable for my growth into an independent researcher. I have greatly benefited from his sharp technical intuition and his genuine care in my development.

I would also like to sincerely thank Dr. Peyman Mohajerin Esfahani for his support and collaboration throughout this project. Peyman's passion for research, deep expertise, precise insights, and clear communication have had a strong impact on both this thesis and my own understanding of research. I am very grateful for the many stimulating discussions we have had and for the guidance he has generously offered along the way.

I would also like to sincerely thank my thesis committee members, Prof. Geert Deconinck, Prof. Florian Dörfler, Prof. Hrvoje Pandžić, and Prof. Bart De Schutter, for their time, careful reading, and insightful comments, which have greatly helped to improve this thesis. Their constructive feedback and expert perspectives are deeply appreciated. I am likewise especially grateful to the chair of my doctoral committee, Prof. Marjan Popov, for his kind support and for his continuous encouragement and guidance throughout the past four years.

I would like to extend my sincere thanks to my collaborators, Dr Amirreza Silani, Shabnam Khodakaramzadeh, Dr Yigu Liu, and Dr Jingwei Dong, for their valuable cooperation and support. I have greatly benefited from our joint work, the many insightful discussions, and their constructive suggestions on our projects and manuscripts. Their contributions have significantly enriched the research presented in this thesis. Special thanks go to Dr

Chenguang Wang, who inspired me the use of the word "orchestration" in the thesis title, and offered me great support and help throughout my PhD journey.

I would like to express my gratitude to all my colleagues in the Intelligent Electrical Power Grids group and the Delft AI Energy Lab for the stimulating academic atmosphere they created.

Last but certainly not least, I would like to thank my family and friends for their care and support. It is hard to name you all here, but I am truly grateful to have had you by my side during this long and sometimes lonely journey. Above all, I am deeply grateful to my parents and my boyfriend, Zhongda, for their love and companionship.

Haiwei Xie
Delft, December 2025

Curriculum Vitæ

Haiwei XIE

11-08-1997 Born in Yangzhou, China.

Education

2015–2019 B.Eng in Electrical Engineering and Automation
Zhejiang University, China

2019–2020 MSc in Future Power Networks
Imperial College London, UK

2022–2026 PhD Candidate in Electrical Sustainable Energy
Delft University of Technology, The Netherlands

List of Publications

Journals

1. **H. Xie**, and J. L. Cremer, "*Game-Theoretic Learning for Power System Dynamic Ancillary Service Provisions*," IEEE Control Systems Letters, Vol. 8, June 2024.

2. **H. Xie**, S. Khodakaramzadeh, P. Mohajerin Esfahani, and J. L. Cremer, "*Enveloping Large Populations of Device Dynamics for Fast Frequency Response in a Dynamic Virtual Power Plant*," submitted.

3. **H. Xie**, A. Silani, and J. L. Cremer, "*Nudging Distributed Energy Resources To Provide Power System Ancillary Services Using Customized Recommendations*," submitted.

4. J. Dong, Y. Liao, **H. Xie**, J. L. Cremer and P. Mohajerin Esfahani, "*Real-time ground fault detection for inverter-based microgrid systems*," IEEE Transactions on Control Systems Technology, Vol. 33, January 2025.

5. Y. Liu, **H. Xie**, A. Presekal, A. Stefanov, and P. Palensky, "*A GNN-Based Generative Model for Generating Synthetic Cyber-Physical Power System Topology*," IEEE Transactions on Smart Grid, vol. 14, no. 6, November 2023.

Peer Reviewed Conferences

1. Y. Liu, **H. Xie**, A. Presekal, A. Stefanov, and P. Palensky, "*A GNN-based Generative Model for Generating Synthetic Cyber-Physical Power System Topology*," 2024 IEEE Power & Energy Society General Meeting (PESGM), Seattle, WA, USA.

2. **H. Xie**, F. Bellizio, J. L. Cremer, and G. Strbac, "*Regularised Learning with Selected Physics for Power System Dynamics*," 2023 IEEE Belgrade PowerTech, Belgrade, Serbia.

Included in this thesis.