

Secure and Cost-Effective Operation of Low Carbon Power Systems Under Multiple Uncertainties

Cormac O'Malley



A thesis submitted to Imperial College London,
Department of Electrical and Electronic Engineering,
for the degree of Doctor of Philosophy

January 2023

The material contained within this thesis is my own work, except where other work is appropriately referenced. Any use of the first-person plural is for reasons of clarity.

Cormac O'Malley (2023)

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Licence (CC BY-NC-ND).

Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text.

Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

Firstly, I would like to thank Prof. Goran Strbac, Dr. Fei Teng and Dr. Luis Badesa for their supervision. My PhD straddled the worst of the Covid-19 pandemic, and during those strange and uncertain times, their unerring guidance and company were invaluable to me.

My second thanks are to my family. Thank you, Patricia, Damian, Finn and Ella, for always supporting me. I'm sure my intellectual curiosity stems somewhat from dinner table conversations while growing up. I'm lucky to have parents who provided patient answers to my relentless questions as a child.

Finally, I thank the funding bodies that made this work possible. This research has been supported by the National Grid ESO as part of the project 'Integrated Development of Low-Carbon Energy Systems' (IDLES, Grant EP/R045518/1), and by the Innovate UK project 'e4Future' (104227).

Abstract

Power system decarbonisation is driving the rapid deployment of renewable energy sources (RES) like wind and solar at the transmission and distribution level. Their differences from the synchronous thermal plants they are displacing make secure and efficient grid operation challenging. Frequency stability is of particular concern due to the current lack of provision of frequency ancillary services like inertia or response from RES generators. Furthermore, the weather dependency of RES generation coupled with the proliferation of distributed energy resources (DER) like small-scale solar or electric vehicles permeates future low-carbon systems with uncertainty under which legacy scheduling methods are inadequate. Overly cautious approaches to this uncertainty can lead to inefficient and expensive systems, whilst naive methods jeopardise system security.

This thesis significantly advances the frequency-constrained scheduling literature by developing frameworks that explicitly account for multiple new uncertainties. This is in addition to RES forecast uncertainty which is the exclusive focus of most previous works. The frameworks take the form of convex constraints that are useful in many market and scheduling problems.

The constraints equip system operators with tools to explicitly guarantee their preferred level of system security whilst unlocking substantial value from emerging and abundant DERs. A major contribution is to address the exclusion of DERs from the provision of ancillary services due to their intrinsic uncertainty from aggregation. This is done by incorporating the uncertainty into the system frequency dynamics, from which deterministic convex constraints are derived. In addition to managing uncertainty to facilitate emerging DERs to provide legacy frequency services, a novel frequency containment service is designed. The framework allows a small amount of load shedding to assist with frequency containment during high RES low inertia periods. The expected cost of this service is probabilistic as it is proportional to the probability of a contingency occurring. The framework optimally balances the potentially higher expected costs of an outage against the operational cost benefits of lower ancillary service requirements day-to-day.

The developed frameworks are applied extensively to several case studies. These validate their security and demonstrate their significant economic and emission-saving benefits.

Contents

Acknowledgements	3
Abstract	4
List of Publications	13
1 Introduction	19
1.1 Motivation and Background	19
1.1.1 Unit Commitment	20
1.1.2 Stability Constrained Unit Commitment	22
1.2 Thesis Structure	24
1.3 Original Contributions	27
2 Reinforcement Learning and Mixed Integer Programming for Power Plant Scheduling in Low Carbon Systems: Comparison and Hybridisation	29
2.1 Introduction	29
2.2 The Unit Commitment Problem	32
2.2.1 Unit Commitment	32
2.2.2 Simulation Environment	34
2.3 Scheduling Agents	35
2.3.1 Mixed-Integer Programming Agents	35
2.3.2 Reinforcement Learning Agents	40
2.4 Case Studies and Comparison	42
2.4.1 Cost Comparison	43
2.4.2 Solve Time Comparison	44
2.4.3 Schedule Comparison	45
2.5 Hybridisation	47
2.5.1 Warm Start Method Comparison	48
2.5.2 Concurrent Solves	50
2.6 Conclusion	53
3 Frequency Response from Aggregated V2G Chargers Under Uncertain Electric Vehicle Connections	55
3.1 Introduction	55
3.2 Modelling of Aggregated V2G Chargers in Frequency Dynamics	59
3.2.1 Frequency Security Constraints Under Uncertainty	60
3.2.2 Convex Reformulation of Chance Constraints	62
3.2.3 Comparison to Simplistic Risk-Aware Scheduling Method	66

3.2.4	State of Charge and Reserve	68
3.3	Stochastic Unit Commitment	69
3.3.1	System Constraints	70
3.3.2	Generator and Storage Constraints	70
3.3.3	Rolling Planning	72
3.4	EV Connectivity Forecasting and Data Analysis	72
3.4.1	Forecasting Technique	73
3.4.2	Tests for Ambiguity Set Selection	75
3.5	Case Studies	77
3.5.1	Constraint Reliability	78
3.5.2	Value of Response from V2G and Comparison to State of the Art	80
3.5.3	Constraint Performance with Varied Fleet Types	82
3.5.4	Investigating time-varying FR contribution from EVs	83
3.5.5	Value's Sensitivity to User Risk aversion and Forecast Uncertainty	85
3.5.6	Value's Sensitivity to System Characteristics	87
3.5.7	Impact of Delays on Value of EV's Response	89
3.6	Conclusion and Future Work	90
4	Probabilistic Scheduling of UFLS to Assist With Single Plant Outages in Low Inertia Systems	91
4.1	Introduction	92
4.2	Frequency Security constraints	97
4.2.1	Nadir Constraint with UFLS	98
4.2.2	Nadir Constraint Convexification	99
4.2.3	Optimising UFLS amount	102
4.3	Risk-Constrained UFLS Feeder Allocation Under Uncertain Distributed Generation	104
4.3.1	Deterministic allocation approach	105
4.3.2	Chance-constrained approach	106
4.3.3	Applicability to current and future scenarios of power system operation .	111
4.4	Constraint Verification	114
4.4.1	Nadir Constraint Verification	115
4.4.2	Feeder Allocation Method - Verification and Case Studies	116
4.4.3	Limitations of a deterministic optimisation approach	117
4.4.4	Chance-constrained approach	119
4.5	Case Studies	123
4.5.1	Evaluation of the Conservativeness of Nadir Constraint Approximation .	126
4.5.2	Value of UFLS in reducing FSC	126
4.5.3	Value of Considering Fast FR and a Reduced Largest Loss	128
4.5.4	Impact of the Expected Cost of UFLS	130
4.5.5	Impact of an Increasing Marginal Load-Shed Cost	133
4.5.6	Communication Requirements for Scheduling UFLS	134
4.6	Conclusion	135
5	Conclusions and Future Work	138

Nomenclature

Indices and Sets

- $j \in \mathcal{J}$ Set of frequency response speeds: $1, \dots, J$.
- $i \in \mathcal{I}$ Set of aggregated EV fleets: $1, \dots, I$.
- $g \in \mathcal{G}$ Set of distinct thermal generators: $1, \dots, G$.
- $\hat{g} \in \hat{\mathcal{G}}$ Set of thermal generator types: $1, \dots, \hat{G}$.
- $s \in \mathcal{S}$ Set of storage units: $1, \dots, S$.
- $k \in \mathcal{K}$ Set of UFLS levels: $1, \dots, K$.
- $t \in \mathcal{T}$ Half-hourly time steps: $1, \dots, T$.
- $n \in \mathcal{N}$ Set of quantile-based scenarios for net demand: $1, \dots, N$.
- $l \in \mathcal{L}_g$ Piecewise production cost intervals for thermal generator g : $1, \dots, L_g$.
- $r \in \mathcal{R}$ Set of net-demand scenario realisations for schedule testing: $1, \dots, R$.
- $g \in \mathcal{G}_{on}^1$ Set of thermal generators which are initially committed (on).
- $g \in \mathcal{G}_{off}^1$ Set of thermal generators which are not initially committed (off).

System Parameters

- c^{ls} Cost of load shed (£/MWh).
- c_k^{UFLS} Value of load-shed at UFLS level k (£/MWh).
- c^{wc} Cost of wind curtailment (£/MWh).
- D_t Demand forecast at time t (GW).
- $\eta D_{t,n}$ Demand forecast error at time t for quantile based scenario n (GW).
- $\eta \hat{D}_{t,n}$ Net demand forecast error at time t for quantile based scenario n (GW).
- f_0 Nominal grid frequency (Hz).
- Δf_{max} Maximum admissible frequency deviation (Hz).
- Δf_{trig} Low frequency deviation at which UFLS is triggered (Hz).
- $H_{\hat{g}}$ Inertia constant of generator type $\hat{g}$ (s).
- M Big-M constant.
- $N_{0,i}$ Current number of connected EVs in fleet i .
- $N_{PL_{max}}$ Number of units generating at PL_{max} .

- $\Delta\hat{N}_{i,t}$ EV connections for fleet i between t_d and the start of t .
- p^{out} Probability of PL_{max} outage (occurrences/h).
- P_k^{UFLS} Load-shed amount at level k (GW).
- q_n Quantile of the net demand forecast error for scenario n .
- R_t Spinning reserve requirement at time t (GW).
- $RoCoF_{max}$ Maximum admissible RoCoF (Hz/s).
- S_t Solar forecast at time t (GW).
- t_d Time of scheduling decision (h).
- t_{rec} Disconnection length of UFLS (h).
- t_s Scheduling time period (h).
- T_{del} Delay of FR from EVs (s).
- T_j Delivery speed of frequency response type j (s).
- W_t Wind forecast at time t (GW).
- $\eta W_{t,n}$ Wind forecast error at time t for quantile based scenario n (GW).
- η Charge/discharge efficiency.
- ϵ Acceptable risk of FR under-delivery from V2G.
- μ_i Mean net EV connection forecast for fleet i .
- σ_i Std of net EV connection forecast for fleet i .
- $\Delta\tau$ Time-step length (h).
- ϕ_n Probability of occurrence for quantile-based scenario n .

Thermal Generator Parameters

- c_g^{st} Startup cost of thermal generator g (£).
- c_g^{nl} No load generation cost for generator type g (£).
- c_g^m Marginal power generation cost for generator type g (£/GWh).
- $CP_{g,l}$ Cost of operating at piecewise generation point l for generator g (£).
- DT_g Minimum down time for generator g (h).
- DT_g^1 Number of time periods the unit has been off prior to the first time period for generator g .
- $P_{g,l}$ Power level for piecewise generation point l for generator g (MW); $P_g^1 = \underline{P}_g$ and $P_g^{L_g} = \overline{P}_g$.

$\overline{P}_g$ Maximum power output for generator g (MW).

$\underline{P}_g$ Minimum power output for generator g (MW).

U_g^1 Initial on/off status for generator g , $U_g^1 = 1$ for $g \in \mathcal{G}_{on}^1$, $U_g^1 = 0$ for $g \in \mathcal{G}_{off}^1$.

UT_g Minimum up time for generator g (h).

UT_g^1 Number of time periods the unit has been on prior to the first time period for generator g .

Decision Variables (binary)

Decision variables are emboldened throughout the thesis.

$\mathbf{u}_{g,t}$ Commitment status of thermal generator g at time t .

$\mathbf{v}_{g,t}$ Startup status of thermal generator g at time t .

$\mathbf{w}_{g,t}$ Shutdown status of thermal generator g at time t .

$\mathbf{b}_{t,n}$ Binary variable to relax (3.22) at time t for scenario n .

$\mathbf{m}_{k,t,n}$ Binary variable corresponding to UFLS level k at time t in scenario n .

Decision Variables (continuous)

Decision variables are emboldened throughout the thesis.

$\mathbf{C}_{i,t,n}$ Charge rate of an individual EV in fleet i at time t for scenario n (GW).

$\mathbf{D}_{i,t,n}$ Discharge rate of an individual EV in fleet i at time t for scenario n (GW).

$\mathbf{E}_{i,t,n}$ Aggregate fleet i , state of charge at the end of timestep t for scenario n (GWh).

$\mathbf{E}_{s,t,n}$ Storage unit s , state of charge at the end of timestep t for scenario n (GWh).

$\mathbf{H}_{t,n}$ System inertia at time t for quantile scenario n after the loss of $\mathbf{PL}_{max}$ (GWs).

$\mathbf{p}_{g,t,n}$ Power above minimum for thermal generator g at time t at quantile scenario n (MW).

$\mathbf{p}_{t,n}^{ls}$ Load shed at time t for quantile scenario n (MW).

$\mathbf{p}_{t,n}^{wc}$ Wind curtailment at time t for quantile scenario n (MW).

$\mathbf{P}_{\hat{g},t,n}$ Power output of units $\hat{g}$ at time t for scenario n (GW).

$\mathbf{p}_{s,t,n}$ Power output from storage s at time t for scenario n (GW).

$\mathbf{p}_{i,t,n}^{EV}$ Power output from EV fleet i at time t for scenario n (GW).

$\mathbf{p}_{t,n}^{sc}$ Solar curtailment at time t for scenario n (GW).

$\mathbf{PL}_{max,t,n}$ Largest power infeed at time t for scenario n (GW).

$\mathbf{r}_{g,t}$ Spinning reserve provided by thermal generator g at time t (MW).

$R_{j,t,n}$ Magnitude of the frequency response of type j , scheduled at time t for scenario n .
Superscripts: $\mathbf{G}$, from thermal plants; $\mathbf{ND}$, from non-distributed non-thermal resources.
(GW)

$\bar{R}^{EV}_{j,t,n}$ Magnitude of the frequency response of type j , scheduled at time t for scenario n from the cumulative EV fleets. (GW)

$U_{\hat{g},t,n}$ Number of committed plants of type $\hat{g}$ at time t for quantile scenario n .

$V_{\hat{g},t,n}$ Startup number of type $\hat{g}$ at time t for quantile scenario n .

$\lambda_{g,t,n,l}$ Fraction of power from piecewise generation point l for generator g at time t for quantile scenario n .

Linear Expressions of Decision Variables (Deterministic)

$C_{\hat{g}}(P_{\hat{g},t,n})$ Cost of power produced for thermal generator type $\hat{g}$ at time t for quantile scenario n (£).

$c_g(p_{g,t,n})$ Cost of power produced above minimum for thermal generator g at time t at quantile scenario n (£).

g_i Individual EV FR capacity in fleet i (GW).

$R_j(t)$ Frequency response dynamics of type j . Superscripts: $\mathbf{G}$, from thermal plants; $\mathbf{ND}$, from non-distributed non-thermal resources. (GW)

$x, \psi_i, \omega, \omega'$ Auxiliary expressions for (3.10) and (3.21).

$C_{t,n}^U$ Expected cost of UFLS at time t for scenario n (£).

X, Y, Z Auxiliary expressions for eq. (4.12).

Linear Expressions of Decision Variables (Stochastic)

δ Excess scheduled FR from EVs compared to amount truly available (GW).

$R_j^{EV}(t)$ Frequency response dynamics of type j from EVs. (GW)

$R_{j,i}^{EV}$ Magnitude of available FR of type j from fleet i (GW).

R_j^{EV} Magnitude of total available FR from all system EVs (GW).

y, z Auxiliary expressions for (3.10).

Nonlinear Expressions of Decision Variables (Deterministic)

$\Delta f^u(t)$ Frequency deviation at time t after outage (Hz).

$\hat{t}$ Time after outage that frequency reaches the UFLS trigger level (s).

Nonlinear Expressions of Decision Variables (Stochastic)

$\Delta f(t)$ Frequency deviation at time t after outage (Hz).

t^* Time after outage of frequency nadir (s).

Random Variables

ΔN_i Net EV connections for fleet i between t_d and the time of outage during t_s .

Markov Decision Process

A Set of all actions (action space).

$P(s'|s, a)$ Transition function: probability of transitioning to state s' from state s after taking action a .

$R(s', s, a)$ Reward resulting from transition to state s' from s following action a .

S Set of all states (state space).

$\pi(a|s)$ Policy: probability of taking action a from state s .

Acronyms

ARMA	Autoregressive Moving Average Process
ADR	Aggregated Distributed Resources
CCGT	Combined Cycle Gas Turbine
CFS	Cost of Frequency Security
DER	Distributed Energy Resources
DG	Distributed Generation
DR-CC	Distributionally Robust Chance Constraint
EV	Electric Vehicle
FS	Frequency Services
FSC	Frequency Security Cost
FR	Frequency Response
MIP	Mixed Integer Programming
MISOCP	Mixed-Integer Second Order Cone Programme
NDFE	Net Demand Forecast Error
OCGT	Open Cycle Gas Turbine
PPO	Proximal Policy Optimisation
RES	Renewable Energy Sources
RL	Reinforcement Learning
RoCoF	Rate of Change of Frequency
SOC	Second Order Cone
SUC	Stochastic Unit Commitment
UC	Unit Commitment
UFLS	Under Frequency Load Shedding
V2G	Vehicle to Grid
WAMS	Wide-Area Measurement Systems

List of Publications

Journals

- [1] **C. O'Malley**, L. Badesa, F. Teng, and G. Strbac, "Probabilistic Scheduling of UFLS to Secure Credible Contingencies in Low Inertia Systems," *IEEE Transactions on Power Systems*, vol. 37, no. 4, pp. 2693–2703, 2022. [Online]. Available: <https://arxiv.org/abs/2103.06616>
- [2] **C. O'Malley**, L. Badesa, F. Teng, and G. Strbac, "Frequency Response from Aggregated V2G Chargers with Uncertain EV Connections," *IEEE Transactions on Power Systems*, accepted for publication 2022. [Online]. Available: <https://arxiv.org/abs/2202.08026>
- [3] L. Badesa*, **C. O'Malley***, M. Parajeles, and G. Strbac, "Chance-Constrained Allocation of UFLS Candidate Feeders Under High Penetration of Distributed Generation," *International Journal of Electrical Power Energy Systems*, vol. 147, p. 108782, 2023. [Online]. Available: <https://arxiv.org/abs/2211.02953>

*Cormac O'Malley and Luis Badesa contributed equally to this publication.

- [4] **C. O'Malley**, P. de Mars, L. Badesa, and G. Strbac, "Reinforcement Learning and Mixed-Integer Programming for Power Plant Scheduling in Low Carbon Systems: Comparison and Hybridisation," *Applied Energy*, under review, submitted 12 2022. [Online]. Available: <https://arxiv.org/abs/2212.04824>

Conferences (peer-reviewed)

- [5] **C. O'Malley**, M. Aunedi, F. Teng, and G. Strbac, "Value of Fleet Vehicle to Grid in Providing Transmission System Operator Services," in *2020 Fifteenth International Conference on Ecological Vehicles and Renewable Energies*, Monte-Carlo (Monacco), Conference Proceedings. [Online]. Available: <https://arxiv.org/abs/2201.05520>

Technical Reports

- [6] F. Oldfield, K. Kumpavat, R. Corbett, A. Price, M. Aunedi, G. Strbac, **C. O'Malley**, D. Gardner, D. Pfeiffer, and J.-T. Kamphus, "*White Paper: The Drive Towards a Low-Carbon Grid. Unlocking the Value of Vehicle-to-Grid in Great Britain*," Nissan, Imperial College London, Eon Drive, 2020.

List of Tables

2.1	Sum of mean operating costs for 40 case studies (\$m)	44
2.2	Number of Unfinished Solves	45
3.1	Expressions for $f_P^{-1}(1 - \epsilon)$	66
3.2	Tests on ΔN_i for Normality and Unimodality	77
3.3	Generation and Storage Characteristics	78
3.4	Constraint Value Comparison for Specified 99% Frequency Security and 100,000 Chargers	80
3.5	Dynamic Simulation Parameters	81
3.6	Increase of FR from EVs Using (3.21) over (3.26)	82
3.7	Sensitivity of Normalised V2G Value to FR Provision Delay	89
4.1	Dynamic Simulation parameters	115
4.2	Test system with 20 candidate feeders for the UFLS scheme, with varying levels of uncertainty in net-load of each feeder.	116
4.3	Mathematical description of the different PDFs considered.	117
4.4	Results of allocation of load shedding using a deterministic optimisation approach.	118
4.5	Results of neglecting (Case 1) and considering (Case 2) correlation among feeders' net-load forecasts.	120
4.6	Results of chance-constrained allocation of load shedding assuming Gaussian uncertainty.	122
4.7	Results of distributionally-robust chance-constrained allocation of load shedding under different types of uncertainty.	122
4.8	Generation and Storage Characteristics	125
4.9	Impact of a Variable Marginal UFLS Cost	133

List of Figures

2.1	Illustration of how quantile-based scenarios can be used to approximate the empirical net demand forecast distribution to reduce MIP size and, therefore, increase tractability. Plotted are 200 net demand forecast realisations, approximated by 3 scenarios corresponding to the quantiles: $(q_1 = 0.1, q_2 = 0.5, q_3 = 0.999)$	37
2.2	Operating costs relative to perfect foresight costs for the four optimisation agents. The box encompasses the inter-quartile range, with a line at the median. The whiskers extend to the max and min values, but are limited to 1.5x the inter-quartile range. The left-hand plot shows operating costs for all 40 case studies; the right-hand plot is clipped to relative costs ≤ 1.5 . The extreme operating costs of RL-MF and RL-LA in some instances, indicate unreliable performance and difficulties in generalising to unseen problems. However, median operating costs of the RL agents are competitive with D-MIP up to 30 generators.	44
2.3	Distribution of solve times for the optimisation agents over the 40 case studies. The RL agents have relatively stable solve times with increasing problem size, whereas the solve times for MIP rise super-linearly with the number of generators. Solve times are limited to a time budget of 10 minutes (600 seconds).	46
2.4	Selected 20 generator case studies, solved by the four optimisation agents. A subset of 100 scenarios of demand and wind are shown for each problem. In case A, both MIP methods outperform the two RL methods, with S-MIP producing the lowest operating costs overall. RL-LA outperforms RL-MF and employs small reserve margins. In case B, RL-LA outperforms D-MIP, committing less reserve. Case C is characterised by two large swings in wind generation; in this case, RL-LA is 78% cheaper than RL. In case D, both RL-MF and RL-LA perform poorly as compared with the MIP methods, with large levels of load shedding during the evening peak.	47
2.5	Average schedule cost and MIP gap over the 40 test days for the 3 warm starting methods with either 1-thread or 8-threads per solve. For the RL and rand methods, each of the 40 days was solved 8 times using 8 distinct initial solutions. More computer resources per solve result in lower average MIP gaps, which correlates with lower cost schedules when assessed via (2.13). However, this does not account for concurrency.	49
2.6	Cost vs MIP gap for days where the original solutions had 0% MIP gap with 1-thread. The optimisation was re-run with varied time budgets to produce a range of final MIP gaps. This demonstrates how the misalignment of the objective being optimised (2.18) with the true objective (2.13) results in a noisy correlation between MIP gap and a schedule's true cost, i.e. some solutions with higher gaps produce lower cost schedules.	51

2.7	Concurrently warm starting multiple MIP solves with various RL solutions produces multiple low gap solutions within the time budget. Plotted here are the eight 1-thread solutions for 29/09/16. Selecting the lowest cost schedule according to (2.13) therefore exploits its slight misalignment with the MIP cost function (2.18). This finds schedules that are lower cost than the average of 1-thread and 8-thread solves.	52
2.8	Taking the lowest cost, RL warm started 1-thread solve for each test day results in schedules that are on average 0.3% lower cost than only solving a single 8-thread MIP using vanilla or RL warm start methods.	53
3.1	Graphic representation of the input/output flow for the SUC. To simulate annual system operation, system variables are updated and the optimisation is iterated every hour.	69
3.2	Schematic of a typical scenario tree used in SUC.	72
3.3	Typical Fri-Sat normalised EV connection numbers for the ‘work’ and ‘domestic’ fleets. Nighttime periods are shaded purple.	73
3.4	Example of domestic EV connection values during a weekday morning with $N_0 = 5,119$, $t_d = 07:00$ and $t_s = 07:00 \rightarrow 08:00$ or $t_s = 08:00 \rightarrow 09:00$ or $t_s = 09:00 \rightarrow 10:00$. FR must consider intra-hour connections (5-min) to define ΔN_i , while charging decisions use average hourly values to define $\Delta \hat{N}$	74
3.5	Empirical distributions of ΔN_i for the domestic fleet on weekday, with $t_d = 07:00$ and $t_s = 07:00 \rightarrow 08:00$ or $t_s = 08:00 \rightarrow 09:00$. These are the distribution of possible changes in the number of EVs connected between 07:00 and any 5-minute period (when an outage could occur and activate FR provision) within the 2 FR scheduling hour time periods.	75
3.6	Flow diagram to summarise forecasting approach. Empirical distributions for future intra-hour EV connectivity are assumed to be well characterised by driving behaviour from a similar time of the week the previous year. t_d is the current time and t_s is the future period in which FR from EVs is being scheduled. . . .	76
3.7	The inter-quartile range, median and max/min of hourly nadir security levels under different ambiguity set assumptions (Gaussian, Unimodal, DRO) on δ when scheduling FR via (3.21). The scheduled amount was compared to true EV connectivity, found by sampling Gaussian or Empirical distributions of ΔN_i . . .	79
3.8	Comparison of the frequency evolution post PL_{max} loss given the delivery of the scheduled amount of FR $\bar{R}_1^{EV}$ or of the 1st percentile of the true FR distribution R_1^{EV} for the system conditions shown in Table 3.5. The ‘scheduled’ plots are very similar so for clarity only one is plotted.	82
3.9	Total scheduled FR from three 15,000 EV ‘work’ fleets over the same two-day period, using the ‘joint’ (3.21) and ‘individual’ (3.26) approaches. Nighttime is shaded purple.	83
3.10	Two day example comparing system operation with and without FR from V2G. Nighttime is shaded purple. The top figure plots the two system’s CCGT commitment number as well as the difference in wind and solar curtailment between the two systems. The bottom figure plots the operation of the Domestic fleet virtual battery for the system with V2G FR enabled.	84
3.11	Sensitivity of charger value to FR under-delivery risk aversion.	86
3.12	Sensitivity of charger value to increased standard deviation (std) of ΔN_i	86

3.13	The main value creation mechanism of FR from V2G chargers is the ability to facilitate renewable generation integration by displacing inertia and FR from thermal plants. The value sensitivity to installed system wind capacity and 10GW ('S10') and 20GW ('S20') of solar capacity is plotted here.	87
3.14	Operational cost increase from applying the frequency security constraints to systems with different storage penetrations. The difference between battery and V2G is primarily because the chargers on average only have an EV connected (thus FR capability) $\approx 40\%$ of the time. Adjusting for this gives the 'Average V2G' capacity plots.	89
4.1	Feasible space defined by the nadir constraint (4.12) for various scheduled amounts of UFLS. $\mathbf{Y}$ can be considered as an adjusted outage size and $\mathbf{Z}$ as an adjusted sum of inertia and FR.	100
4.2	Convexification of the nadir constraint (4.12) using a second order cone (4.17). For a fixed UFLS amount, the SOC's gradient (g) and Z intercept (c) are defined by the tangential point $(\hat{Y}, \hat{Z})$. The SOC approximation with the minimum A_{con} is the least conservative.	101
4.3	Value of the left-hand side constant for chance constraints (4.32) and (4.39), over a range of risk preference (ϵ). To accommodate any distribution type with a specified mean and standard deviation, the distributionally robust constraint's constant is larger over all ϵ , resulting in a robust and therefore more conservative solution.	111
4.4	Frequency excursions following a generation loss, for a WAMS-based adaptive UFLS scheme with operational and unavailable communications. When communications fail, the relays selected for the backup scheme disconnect their load once their local frequency reaches the 0.8Hz trigger level. These relays must be selected ahead of time and thus are subject to net-load uncertainty caused by DG.	114
4.5	Frequency evolution after a loss of PL_{max} for the system conditions in Table 4.1.	116
4.6	Heatmap for the correlation coefficient matrices of the uncorrelated and correlated net-load probabilistic forecast.	120
4.7	Sampled distributions of available load in the selected feeders, for the four PDFs considered for underlying net-load uncertainty. Feeders were selected through a chance-constrained optimisation assuming Gaussian uncertainty.	123
4.8	Sampled distributions of available load in the selected feeders, for the four PDFs considered for underlying net-load uncertainty. Feeders were selected through a distributionally-robust chance-constrained optimisation.	124
4.9	The conservative approximation of (4.12) by (4.19) results in the over scheduling of FS, which could secure a stricter nadir (Δf_{trig}^{true}) than that required ($\Delta f_{trig} = 0.8$ Hz). Shown here are the annual inter-quartile range, median and max/min of Δf_{trig} for systems with two levels of UFLS available, 0 GW or the value shown on the x axis.	127
4.10	The annual reduction in frequency security costs and emissions from optimally scheduling UFLS to support frequency security, in systems with 35 GW and 55 GW of installed wind capacity.	128
4.11	Three-day example of the system operation, showcasing the use of UFLS to accomodate more wind.	129

4.12	The expected amount of annual load shed from optimally scheduling UFLS to assist with frequency security, in systems with 35 GW and 55 GW of installed wind capacity.	130
4.13	Annual FSC for systems with 0.6 GW of UFLS, and varying fast FR and reduced largest loss capabilities.	131
4.14	Savings in annual frequency security costs enabled by 0.6 GW of UFLS in systems with varying fast FR and reduced largest loss capabilities.	131
4.15	Sensitivity of frequency security savings from 0.6 GW of UFLS to the cost of the load shed (c^{UFLS}): ‘Free UFLS’ load shed is zero cost; ‘Fixed UFLS’ UFLS is scheduled at all times; ‘Optimal UFLS’ optimiser chooses when to schedule UFLS.	132
4.16	FSC savings from two types of UFLS scheme. Constant UFLS, which allows scheduling of a fixed amount of UFLS shown on the x axis, or UFLS that tracks a set percentage of national demand. The range of UFLS levels available within the set percentage schemes is shown by their coverage of the x axis.	135

Chapter 1

Introduction

This chapter gives an introduction to the thesis by explaining the motivation to undertake this work. It subsequently details the thesis structure, summarising the main results of each chapter followed by a list of the major research contributions.

1.1 Motivation and Background

To prevent the most catastrophic impacts of climate change, the world needs to become carbon neutral by 2050 [7]. Electricity generation currently accounts for 40% of global emissions and will play a key role in decarbonising other sectors such as heating and transport [8]. Recognising this, many countries have committed to dates by which their power systems will be fully decarbonised, such as 2035 for the USA and the UK, 2050 for Indonesia and 2060 for China [9].

Power system decarbonisation is characterised by replacing polluting thermal plants that run on fossil fuels like coal and gas with low-carbon renewable energy sources (RES), mainly wind and solar. Simultaneously, the total power generation capacity must be scaled steeply to accommodate heating and transport electrification. Accordingly, RES are currently being built at unprecedented scales. Over the next 5 years, 1.470 TW of solar and 0.702 TW of wind will be installed globally [10]. This will bring global RES capacity to 5.640 TW, 28% larger than the currently installed fossil fuel generator capacity. Build rates are forecast to accelerate in the longer term future. Consequently, the cost-efficient and secure integration of RES into modern power systems is key to the energy transition and is the topic of this thesis.

The operation of RES-dominated power systems presents significant challenges. These stem from the differences between RES and traditional dispatchable thermal plants, around which the previous operation relied upon. The three main differences are:

1. Variability - RES generation is highly weather dependent. Wind generation is proportional to the cube of the wind speed up to a turbine's rated power and solar to the incident solar irradiance. Hence the power generated is non-dispatchable and varies over all timescales, from seconds to years.
2. Uncertainty - RES power generation cannot be known exactly ahead of time. A good forecast of wind [11] or solar [12] generation will be probabilistic with explicit uncertainty quantification. The uncertainty from thermal generators is dominated by fault probabilities. RES generation also suffer from faults.
3. Stability - Grid-connected RES characteristics differ highly from traditional synchronous thermal plants, thus so does their interaction with the grid. This presents challenges in maintaining traditional grid stability metrics such as frequency, voltage and rotor angle [13]. New instability modes are also emerging due to the interaction between RES power electronic controllers.

Once a RES generator is built, the marginal cost and emission of electricity it generates is close to zero. This, coupled with the fact that RES are non-dispatchable, means that low-carbon power system operational methods focus on scheduling the few dispatchable plants in the system to integrate the RES most efficiently. This is called the unit commitment (UC) problem and is the core optimisation problem within this thesis.

1.1.1 Unit Commitment

The UC is used for cost minimisation in centrally dispatched systems including North American and Australian markets, and a variety of security assessment and intra-day operational tasks in self-dispatching systems like Great Britain. Due to the intrinsic uncertainty of RES power forecasts, the efficient scheduling of thermal plants to satisfy demand is becoming increasingly challenging.

The multi-period UC problem is typically solved using mixed-integer programming (MIP) methods which leverage the powerful algorithms used in off-the-shelf commercial MIP solvers. These achieve near-optimal solutions for most problems but can struggle to solve large problem instances over operational timescales while taking rigorous account of RES forecast uncertainty [14]. MIP methods for UC are categorised as deterministic [15] or stochastic [16]. Deterministic methods deal with uncertainty by setting an explicit reserve requirement, using experience and heuristics. Stochastic MIP accounts for uncertainty by minimising expected costs over a finite set of discrete scenarios. Stochastic programs result in lower costs compared to deterministic MIPs [17, 18], but come at the cost of increased complexity, decreased tractability and the need for explicit knowledge of the forecast uncertainty distributions.

This thesis builds upon the state-of-the-art stochastic unit commitment formulation presented in [19]. A quantile-based scenario tree method is used to discretize the uncertain net demand forecasts. User-defined quantiles of the random variable of net demand (demand net of wind power) are used to construct a scenario tree. The scenario tree branches into n quantile scenarios at the here-and-now root node only. Each of these scenarios extends 24hr into the future with T nodes, each of which has a full set of generation, storage and power balance constraints which are detailed explicitly in the proceeding chapters. Reference [20] demonstrates that well-chosen quantiles that branch only at the root node can discretize the continuous range of potential wind realisations accurately whilst yielding a highly tractable model. A single-bus power system model is used as it is common to solve the commitment and dispatch problems sequentially, adjusting the initial solution to account for line-flows. Here the first stage problem is modelled.

Each node has a complete set of decision variables, subject to the power balance constraint:

$$\sum_{\hat{g} \in \hat{G}} P_{\hat{g},t,n} + \sum_{s \in S} p_{s,t,n} + (W_t + \eta W_{t,n}) - p^{wc}_{t,n} = D_t - p^{LS}_{t,n} \quad \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (1.1)$$

The ability to curtail demand via a load-shedding decision variable provides the optimisation with recourse to ensure feasibility. Load shedding is penalised in the cost function at the value of lost load. The MIP will find the optimal balance between risking load-shed at high net demand, low probability scenarios and over-committing reserve because the probability of reaching a

given scenario weights the cost function:

$$\sum_{n \in \mathcal{N}} \phi_n \sum_{t \in \mathcal{T}} \left(\Delta \tau (c^{LS} \mathbf{p}_{t,n}^{LS}) + \sum_{\hat{g} \in \hat{G}} C_{\hat{g}}(\mathbf{P}_{\hat{g},t,n}) \right) \quad (1.2)$$

The above cost function does not penalise wind curtailment explicitly, instead, the solver will still seek to minimise curtailment due to wind's zero marginal cost.

Due to the impressive success of reinforcement learning (RL) in solving some challenging optimisation tasks under uncertainty [21], interest in its application to the UC is growing. Emerging RL methods for UC [22, 23] offer an alternative to incumbent MIP methods that could potentially address some of the convexity and tractability issues. Before this thesis, in a problem characteristic of the UC field in general [15], there was no level-playing-field comparison between the RL and MIP paradigms. This is a vitally needed contribution in order to assess their relative performances and highlight valuable future research directions [4].

1.1.2 Stability Constrained Unit Commitment

Stability-constrained scheduling extends UC formulations to explicitly guarantee that their solutions satisfy robust security criteria in addition to reserve and energy requirements. A core focus of this thesis is on developing scheduling methods that secure grid frequency following a large power loss cost-effectively. System operators are obliged to keep the grid frequency within a pre-specified range of the nominal frequency. This is vital in order to prevent system damage and costly emergency demand disconnection. In the UK for example, the National Grid ESO is obliged to keep the frequency between 49.2Hz and 50.8Hz with low frequency demand disconnection first occurring at 48.8Hz [24].

The rate of change of frequency (RoCoF) is determined by the instantaneous demand-generation balance. Immediately following an outage, kinetic energy from the rotating parts of synchronous thermal plants can be extracted to cover the transient deficit. This inertial response is proportional to the RoCoF. The maximum RoCoF must be limited below a pre-specified quantity in order to prevent the disconnection of distributed generation with RoCoF loss of mains protection systems. The lose of distributed generation would exacerbate the demand-generation

imbalance and jeopardise frequency security further.

Teng et al [25] show this can be prevented by limiting the minimum inertia level during scheduling:

$$\frac{2|RoCoF_{max}|}{f_0} \cdot \mathbf{H} \geq \mathbf{PL}_{max} \quad (1.3)$$

This constraint is applied to every node in the scenario tree, so the n and t subscripts are dropped here for clarity.

Frequency response (FR) is a net power injection deployed during this transient to re-balance demand and generation to arrest frequency decline. RES resources are currently not able, or not incentivised to provide an inertial response or FR. Thus their displacement of thermal plants makes the frequency more volatile and harder to contain. In the future when RES will provide frequency services [26], the scheduling methods developed in this thesis can increase the efficiency of procurement and therefore decrease operational costs.

Legacy methods to secure frequency require many committed and partially loaded thermal plants whose minimum stable generation can result in severe RES curtailment. This limits RES integration and can potentially increase the cost of energy by as much as 25% in future systems [27]. Thus, work that strives to reduce the reliance on synchronous plants for frequency security is critical for the energy transition. In the frequency-constrained scheduling literature, the main mathematical challenge lies in containing the frequency nadir, as it necessitates incorporating the dynamic equations dictating post-event frequency evolution into the algebraic optimisation problem [28].

A promising method involves approximating the FR provision from droop control as a linear ramp [25, 28, 29, 30] to simplify the analytical formulation. The paper [31] shows that under this assumption the nadir constraint can be formulated as a convex second-order cone that facilitates the co-optimisation of inertia, fast and slow FR and a dynamically reduced largest loss in order to contain the nadir:

$$\left(\frac{\mathbf{H}}{f_0} - \frac{\mathbf{R}_1 \cdot T_1}{4\Delta f_{max}} \right) \frac{\mathbf{R}_2}{T_2} \geq \left(\frac{\mathbf{PL}_{max} - \mathbf{R}_1}{2\sqrt{\Delta f_{max}}} \right)^2 \quad (1.4)$$

These frequency security constraints are highly tractable and convex, useful in a wide range

of market and scheduling contexts.

This thesis significantly improves the frequency-secured scheduling literature by incorporating uncertainty into the frequency dynamics. This is a significant upgrade to the current methods that almost exclusively consider RES forecast error as the only uncertainty [32]. We present a framework to schedule FR from aggregated distributed energy resources (like electric vehicles) whilst maintaining system security guarantees under their intrinsic capacity uncertainty [2]. This allows their fair comparison with legacy providers such as thermal plants and demonstrates the high and unnecessary cost of their current exclusion from the service provision framework.

In addition to managing uncertainty to facilitate aggregated resources to provide legacy frequency services, this thesis designs a novel frequency containment service whose potential cost is intrinsically probabilistic [1]. Current ESO practice is to always procure enough inertia and response to prevent under-frequency demand disconnection from any single loss on the system.

However, the amount of inertia and FR needed to contain the frequency nadir is driven by the square of the size of the largest generator loss [25]. For systems with a few large losses and high RES penetrations (like the UK), it may be economical to risk a small amount of load shedding on the rare occasions that an outage occurs. This can decrease day-to-day operational costs significantly. Based on this logic, the framework presented in this thesis challenges the paradigm of load shedding as a last resort only. Instead, framing moderate amounts of load shedding as an alternative frequency service, empowering the operator to secure the frequency more cost-effectively and to reduce carbon emissions.

1.2 Thesis Structure

Following the first introductory chapter, the thesis is organised as follows. The second chapter compares incumbent MIP methods for solving the stochastic UC problem against emerging reinforcement learning methods. This is important because growth in RES generation increases uncertainty in the load to be served by power plants and makes their efficient scheduling more difficult. Two MIP (deterministic and stochastic) and two RL (model-free and with lookahead)

scheduling methods are extensively tested over a large set of test days and problem sizes, for the first time comparing the state-of-the-art of these two approaches on a level playing field. It is found that deterministic and stochastic MIP consistently produce lower-cost UC schedules than RL, exhibiting better reliability and scalability with problem size. The average operating costs of RL are more than 2 times larger than stochastic MIP for a 50-generator test case, while the cost is 13 times larger in the worst instance.

However, the key strength of RL is the ability to produce solutions practically instantly, irrespective of problem size. This advantage is leveraged to produce various initial solutions for warm-starting concurrent stochastic MIP solves. By producing several near-optimal solutions simultaneously and then evaluating them using Monte Carlo methods, the differences between the true cost function and the discrete approximation required to formulate the MIP are exploited. The resulting hybrid technique outperforms both the RL and MIP methods individually, reducing total operating costs by 0.3% on average.

The second chapter establishes that the incumbent MIP scheduling methods for UC are currently superior to nascent RL techniques. Consequently, the remainder of the thesis focuses on significantly improving existing MIP security-constrained UC literature by incorporating multiple new uncertainties in addition to RES forecast error. Chapter three develops a novel framework that enables the scheduling of FR from aggregate EVs despite the intrinsic capacity uncertainty that stems from their aggregation.

This is useful because fast FR is highly effective at securing frequency dynamics after a generator outage in low inertia systems [28]. EVs equipped with vehicle to grid (V2G) chargers could offer an abundant source of fast FR in future. However, the uncertainty associated with V2G aggregation, driven by the uncertain number of connected EVs at the time of an outage, has not been fully understood and prevents their participation in the existing service provision framework.

To tackle this limitation, this chapter, for the first time, incorporates such uncertainty into system frequency dynamics, from which probabilistic nadir and steady-state frequency requirements are enforced via a derived moment-based distributionally-robust chance constraint. Field data from over 25,000 chargers are analysed to provide realistic parameters and connection

forecasts to examine the value of FR from V2G chargers in the annual operation of the GB 2030 system. The case study demonstrates that the uncertainty of EV connections can be effectively managed through the proposed scheduling framework, which results in annual savings of £6,300 or 37.4 tCO₂ per charger. The sensitivity of this value to renewable capacity and FR delays is explored, with V2G capacity shown to be a third as valuable as the same grid battery capacity.

The fourth chapter presents a MIP framework in which a moderate amount of load shedding can be scheduled to assist with large generator outages. This reframes under frequency load shedding (UFLS) from a last-resort emergency service only, into a service that can drastically reduce normal system operation costs. Crucially this is achieved with a very low impact on power delivery to consumers. This is achieved by deriving a novel constraint from the swing equation to contain the frequency nadir. It does so by optimally scheduling 5 distinct services: fast and slow FR; a dynamically reduced largest loss; inertia and load shedding. The expected cost of UFLS is found probabilistically to facilitate its comparison to the other frequency services.

We demonstrate that this non-convex constraint can be accurately and conservatively approximated for moderate UFLS levels with a second order cone (SOC), resulting in highly tractable convex problems. Case studies performed on a Great Britain 2030 system demonstrate that UFLS as an option to contain single plant outages can reduce annual operational costs by up to £559m, 52% of frequency security costs. The sensitivity of this value to wind penetration, the abundance of alternative frequency services, UFLS amount and cost is explored.

UFLS is provided by disconnecting feeders whose net-demand sums up to the desired amount. The amount can be decided upon by the method presented in chapter 4 or otherwise. The recent integration of distributed generation (DG) increases the uncertainty in the net consumption of feeders. Indeed, many feeders that historically consumed power can become power exporters, during a sunny summer afternoon in an area with many solar panels for example. Consequently, their disconnection exacerbates the demand-generation imbalance rather than ameliorating it.

Traditional deterministic methods are insufficient to account for this variability and uncertainty and therefore risk system damage by under-shedding load and failing to arrest frequency decline. Hence another major contribution of chapter 4 is to present a chance-constrained methodology for selecting feeders with mathematical guarantees for the disconnection of the

required amount of load with a certain pre-defined probability. Relevant case studies demonstrate the applicability of the proposed method and highlight the need for explicit consideration of uncertainty in the UFLS-allocation process.

1.3 Original Contributions

This thesis has four main contributions to improving the operation and scheduling of low-carbon power systems under uncertainty:

- To provide the most extensive comparison to date between RL and MIP for solving UC problems. Furthermore, the best solution of either method is improved through a novel hybrid methodology using RL to warm-start the stochastic MIP solver. We show that this method outperforms the generic warm starting method of a powerful commercial solver by leveraging the UC specific information encoded in the RL agent.
- To propose a novel, convex distributionally robust chance constraint on the maximum scheduled fast FR from aggregated distributed resources. This is used to investigate the impact of EV connection uncertainty on their FR provision by directly incorporating that uncertainty into the system frequency dynamics. Simulation of the GB 2030 system provides insight on the substantial value of FR provision from the abundant EVs and its sensitivity to grid battery and renewable penetration.
- To develop a comprehensive framework to co-optimize five frequency security services: UFLS, fast and slow FR, inertia and a reduced largest loss. This is achieved via a least conservative convexification of a novel nadir constraint derived from the swing equation. The value of the developed model is demonstrated via case studies of a representative GB 2030 system.
- To deduce a convex-optimisation framework for the UFLS-allocation problem under feeder load flow uncertainty caused by distributed generation. This enables the user to meet a predefined level of risk in obtaining the minimum load to be shed. The advantages of the proposed method include providing the optimal selection of feeders for systems with

traditional UFLS schemes, as well as serving as a robust backup for proposed low latency communication-based adaptive UFLS schemes.

Chapter 2

Reinforcement Learning and Mixed Integer Programming for Power Plant Scheduling in Low Carbon Systems: Comparison and Hybridisation

This chapter compares the state-of-the-art mixed integer programming (MIP) and reinforcement learning (RL) methods for solving the unit commitment problem under RES forecast uncertainty. It is based on the journal paper [4]. The second part of this chapter presents a hybrid RL and MIP method that is shown to outperform either one individually.

2.1 Introduction

Unit commitment (UC) is the task of determining the on/off schedules of power plants ahead of time and is one of the fundamental problems in power system optimisation.

A new paradigm of solving UC problems is RL, leveraging powerful deep learning techniques to map problem features to solutions. RL agents learn by trial and error through repeated simulation of training cases. Artificial intelligence practitioners advocate for RL as a methodology capable of learning optimal control strategies in a large class of problems, including the operation

of electricity networks [33, 34]. Compared to MIP, RL offers the advantage of off-loading most of the computational cost to a training phase, with the potential for rapid decision making close to dispatch. RL is well-suited to decision making with imperfect information, learning operational strategies to manage uncertainty (such as reserve allocation) without reliance on heuristics. To date, no rigorous comparison of MIP and RL for solving UC problems has been made, partly due to limited crossover and collaboration between RL and power systems research communities.

In theory, RL can produce close to optimal operating decisions regardless of the underlying convexity of the system. This is a marked advantage over MIP, which must approximate the true system in a convex (often linear) manner. The optimal decisions for the convex approximation can be non-optimal for the true system. The ability to operate non-convex systems accurately also gives RL a theoretical advantage over supervised learning techniques for UC. Such methods focus on either speeding up MIP solvers [35, 36, 37], or estimating optimal solutions based on a training set of previous MIP solves [38]. As such, supervised learning methods require the same convex approximations of non-convex systems as MIP.

Attempts to solve UC problems using RL have mainly focused on model-free methods based on Q-learning [39, 40, 41, 42, 22], focusing exclusively on small-scale power systems of up to 10 generators. Methods combining RL with tree search have outperformed deterministic MIP approaches on problems of up to 30 generators [23, 43]. However, these methods do not reap the benefits of rapid solution generation that is characteristic of conventional model-free RL methods. Furthermore, no existing studies have compared RL with stochastic MIP.

This chapter compares the performance of two RL (model-free and lookahead) and two MIP (deterministic and stochastic) approaches for solving UC problems. The methods are rigorously evaluated on a diverse set of 40 test days for 5 problem sizes (10–50 generators) using an open-source benchmark system formulated in Python. To the best of our knowledge, no previous paper has compared the state-of-the-art of RL and MIP for UC on a level playing field.

It is shown that MIP methods significantly outperform RL for all problem sizes in terms of operating costs and exhibit significantly better scaling properties. In addition, the performance of RL is highly unreliable, with large volumes of load shedding for specific problem instances.

Whereas the solution times of MIP methods scale super-linearly with problem size, the run time of the RL methods remains roughly constant with the number of generators. Rather than discouraging work on RL for UC, the authors hope to focus and encourage future research by presenting a clear-eyed view of its current shortcomings.

One such promising research direction is in the combination of machine learning and MIP techniques [44]. An agent learns useful properties of the UC problem during training and, once trained, can generate multiple feasible solutions rapidly. Leveraging this advantage, we propose a novel hybrid methodology that employs the RL agent to generate multiple warm-start solutions, used to initialise concurrent stochastic MIP solves. This approach outperforms conventional stochastic MIP and a baseline warm-starting technique, reducing the cost of uncertainty by 0.3%, or \$110,000/yr in the 50 generator system. This exemplifies the opportunities for advancement found at the intersection of RL and MIP, with the potential for large absolute savings in large systems such as MISO or PJM, each with over a thousand generators.

The contributions of this chapter are:

1. To provide the most extensive comparison to date between RL and MIP for solving UC problems.
2. To improve the best solution of either method through a novel hybrid methodology, using RL to warm-start the stochastic MIP solver. We show that this method outperforms the generic warm starting method of a powerful commercial solver, by leveraging the UC-specific information encoded in the RL agent.
3. A variant of the hybrid method is also proposed, taking advantage of parallel computing resources to generate several candidate solutions to the UC. It is shown that this approach somewhat compensates for the unavoidable approximation of stochastic MIP in modelling uncertainty, leading to cost savings when applied to operate a power grid.

The chapter is organised as follows: Section 2.2 details the UC problem, followed in Section 2.3 by descriptions of the MIP and RL agents used to solve it. Section 2.4 compares the methods' efficacy via case study results, while Section 2.5 presents the improved results from hybridising the RL and MIP methods. Finally, Section 2.6 gives the conclusions.

2.2 The Unit Commitment Problem

To facilitate the fair comparison of RL and MIP for the UC problem, a unified problem setup is presented. This enables different scheduling agents to be evaluated in a common environment. We begin by describing the method-agnostic problem setup, designed to reflect real-world objectives closely. A practical simulation environment is then presented, which is used to evaluate UC solutions on different scenarios, as well as to train the RL agents on a variety of problem instances.

2.2.1 Unit Commitment

An agent must output a plant commitment schedule at 23:30 for every half-hour time period for the following day. To allow for post-solve feasibility checks and dispatch signals to be sent with sufficient notice, the agent's best solution at 10 minutes is used. The agent aims to produce the schedule that minimises the expected cost of meeting demand. Although the plant commitments are fixed, the dispatch is varied in real-time to accommodate net-demand realisations best. Thus an agent's goal is defined as follows:

$$\min_{\mathbf{u}} : \mathbb{E} \left[\sum_{t \in \mathcal{T}} \left(\Delta\tau (c^{ls} \cdot \mathbf{p}_t^{ls} + c^{wc} \cdot \mathbf{p}_t^{wc}) + \sum_{g \in \mathcal{G}} \left(c_g^{st} \cdot \mathbf{v}_{g,t} + CP_g^1 \mathbf{u}_{g,t} + \Delta\tau \cdot \mathbf{C}_g(\mathbf{p}_{g,t}) \right) \right) \right] \quad (2.1)$$

Where $\mathbf{C}_g(\mathbf{p}_{g,t})$ is a generators piece-wise linear production cost [45]. Decision variables and expressions of decision variables are emboldened. For generator g :

$$\mathbf{C}_g(\mathbf{p}_{g,t}) = \sum_{l \in \mathcal{L}_g} (CP_{g,l} - CP_{g,1}) \lambda_{g,t,l} \quad \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (2.2)$$

$$\mathbf{p}_{g,t} = \sum_{l \in \mathcal{L}_g} (P_{g,l} - P_{g,1}) \lambda_{g,t,l} \quad \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (2.3)$$

Where $P_{g,1}$ is the minimum stable generation of generator g . The variables $\lambda_{g,t,l}$ are continuous and constrained between 0 and 1. Their sum must equal the commitment status of the generator:

$$u_{g,t} = \sum_{l \in \mathcal{L}_g} \lambda_{g,t,l} \quad \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (2.4)$$

The initial generator commitment statuses are pre-specified and enforced via:

$$\sum_{t=1}^{UT_g - UT_g^1} (u_{g,t} - 1) = 0 \quad \forall g \in \mathcal{G}_{on}^1 \quad (2.5)$$

$$\sum_{t=1}^{DT_g - DT_g^1} u_{g,t} = 0 \quad \forall g \in \mathcal{G}_{off}^1 \quad (2.6)$$

$$u_{g,1} - U_g^1 = v_{g,1} - w_{g,1} \quad \forall g \in \mathcal{G} \quad (2.7)$$

Generator minimum up and down times are enforced through:

$$u_{g,t} - u_{g,t-1} = v_{g,t} - w_{g,t} \quad \forall g \in \mathcal{G}, \forall t \in \{2, \dots, T\} \quad (2.8)$$

$$\sum_{i=t-UT_g+1}^t v_{g,i} \leq u_{g,t} \quad \forall g \in \mathcal{G}, \forall t \in \{UT_g, \dots, T\} \quad (2.9)$$

$$\sum_{i=t-DT_g+1}^t w_{g,i} \leq 1 - u_{g,t} \quad \forall g \in \mathcal{G}, \forall t \in \{DT_g, \dots, T\} \quad (2.10)$$

The agents seek the plant commitment schedules that minimise the expected cost of load shed, wind curtailment and generation for the next day. At each timestep generator constraints

eqs. (2.2) to (2.4), eqs. (2.5) to (2.10), and the energy balance must be maintained:

$$\mathbf{p}_t^{ls} - \mathbf{p}_t^{wc} + \sum_{g \in \mathcal{G}} (P_g^1 \mathbf{u}_{g,t} + \mathbf{p}_{g,t}) = D_t - W_t + (\eta D_t - \eta W_t) \quad \forall t \in \mathcal{T} \quad (2.11)$$

The wind forecast error (ηW_t) is defined as an autoregressive moving average process ‘ARMA(p, q)’ of the form:

$$\eta W_t = \sum_{i=1}^p \alpha_i \eta W_{t-i} + \sum_{i=1}^q \beta_i \epsilon_{t-i} + \epsilon_t \quad (2.12)$$

Where ϵ_t is a standard normal variable. The demand forecast error (ηD_t) also follows an ARMA(p, q) process of the same form but with different values for the α_i, β_i parameters.

It is impossible to calculate a closed-form solution of the expectation (2.1) over these stochastic parameters. However, it can be closely approximated by finding the uniformly-weighted sum over a large number (R) of Monte Carlo generated test scenarios. To evaluate a given plant commitment schedule, the wind and demand forecast error realisations for each scenario are found by sampling the ARMA processes. The generator set-points for that scenario are found by solving the economic dispatch problem: this is a convex optimisation problem that can be rapidly solved. This is repeated R times and averaged assuming uniform likelihood:

$$\min_{\mathbf{u}} : \sum_{t \in \mathcal{T}} \left(\sum_{g \in \mathcal{G}} (c_g^{st} \cdot \mathbf{v}_{g,t} + C P_g^1 \mathbf{u}_{g,t}) \right) + \frac{1}{R} \sum_{r \in \mathcal{R}} \left(\Delta \tau (c^{ls} \cdot \mathbf{p}_{t,r}^{ls} + c^{wc} \cdot \mathbf{p}_{t,r}^{wc}) + \sum_{g \in \mathcal{G}} \Delta \tau \cdot \mathbf{C}(\mathbf{p}_{g,t,r}) \right) \quad (2.13)$$

2.2.2 Simulation Environment

The calculation of (2.13) is conducted in a dedicated simulation environment for the UC problem, as described in [23]. The environment is openly available as a Python package¹, and can be used to evaluate the quality of solutions represented as a binary matrix $\mathbf{u}$ generated by any optimisation technique. In addition, the simulation environment is used to realise the Markov decision process formulation of the UC problem described in Section 2.3.2 and train RL agents.

The environment uses generator specifications given in [46], defining quadratic fuel cost curves, minimum up/down times, start-up costs and initial generator commitments for $N = 10$

¹<https://github.com/pwdemars/r14uc>

generators. To create larger problem sizes where $N > 10$, generators are duplicated. To create individual UC problem instances consisting of demand and wind forecasts, publicly-available data from National Grid [47] and Elexon [48] was used. 806 complete days of national demand and wind generation data for Whitelee wind farm were retrieved, covering dates from 2016–2019 inclusive. 40 case studies were randomly selected to form evaluation case studies used in Section 2.4, while the remaining days were used as training data for the RL agents. Wind shedding is penalised at a value of \$40/MWh, while load shedding is penalised at \$10,000/MWh.

2.3 Scheduling Agents

Here the two MIP (deterministic and stochastic) and the two RL (model-free and lookahead) agents used to solve the UC problem presented in the previous Section are defined. We utilise the generator constraints eqs. (2.2) to (2.4) and eqs. (2.5) to (2.10) defined in [49], which [50] concludes are state-of-the-art. The MIP agents are comprised of a cost function and explicit accompanying constraints that are solved for individual problem instances by an off-the-shelf convex solver. The RL agents follow a decision-making policy that converts the environment state variable into an action (e.g. turn generator on/off). During training, the RL agent implicitly learns the UC and generator constraints via repeated interactions with the UC environment.

2.3.1 Mixed-Integer Programming Agents

Stochastic optimisation finds the decision variables that minimise the expected generation cost over the day in question. However, for (2.13) to accurately approximate the expected cost, the formulation must capture the low-probability worst-case scenarios when load shedding may occur at 500 times the cost of generation [16, 20]. This will be accomplished by utilising a large R . However, this can lead to tractability issues because each scenario optimised over introduces a full additional set of the continuous decision variables and their corresponding constraints.

To solve this issue a quantile-based scenario tree [20] is used to accurately approximate the large Monte Carlo generated tree of (2.13) with many fewer scenarios. This section details the scenario generation methodology and then formally defines the deterministic MIP (D-MIP) and

stochastic MIP (S-MIP) agents.

Scenario Tree

A scenario tree is used to represent the R Monte Carlo generated scenarios in (2.13) with many fewer N representative scenarios corresponding to a user-defined set of net demand forecast error (NDFE) quantiles. An advantage of this method is that it allows extreme high and low net demand realisations to be explicitly accounted for in relatively few scenarios, which is important due to the high expense of load shedding due to under-commitment (and to a lesser extent, wind shedding due to over-commitment). This method has been shown to give a good balance between tractability and accuracy [20].

Each scenario's cost is weighted by the probability of that quantile's occurrence (ϕ_n) within the cost function. Unlike in (2.13), scenario probabilities are not homogeneous. Importantly, the binary commitment variables are shared between scenarios. Thus, only 1 commitment schedule is produced per optimisation, regardless of the number of scenarios.

To find one NDFE sample, the wind and demand forecast errors are sampled to give $\eta W_{t,q}, \eta D_{t,q} \forall t$, where q is the sample number index. These two time series are then added to that day's mean forecasts and clipped to be larger than zero as negative wind or demand has no physical meaning:

$$D_t + \eta D'_{t,q} = \max[0, D_t + \eta D_{t,q}] \quad \forall t \in \mathcal{T} \quad (2.14)$$

$$W_t + \eta W'_{t,q} = \max[0, W_t + \eta W_{t,q}] \quad \forall t \in \mathcal{T} \quad (2.15)$$

One net demand realisation is found by subtracting the wind forecast error from demand forecast error:

$$\eta \hat{D}_{t,q} = \eta D'_{t,q} - \eta W'_{t,q} \quad \forall t \in \mathcal{T} \quad (2.16)$$

Within the power balance constraints of each quantile-based scenario, the NDFE values




Figure 2.1: Illustration of how quantile-based scenarios can be used to approximate the empirical net demand forecast distribution to reduce MIP size and, therefore, increase tractability. Plotted are 200 net demand forecast realisations, approximated by 3 scenarios corresponding to the quantiles: $(q_1 = 0.1, q_2 = 0.5, q_3 = 0.999)$.

$(\eta\hat{D}_{t,n} \forall t)$ are equal to the value at which there is a $q_n \times 100\%$ probability of a lower NDFE realisation occurring. These values are found from $\eta\hat{D}_t$'s empirical distribution. This is found by sampling $\eta\hat{D}_t \forall t$ a large number of times ($\approx 100,000$). Collecting these samples of $\eta\hat{D}_{t,q}$ gives the empirical distribution of $\eta\hat{D}_t$.

Fig. 2.1 plots 200 such scenarios added to that day's net demand forecast. Figure 2.1 also plots the quantile scenarios: $(q_1 = 0.1, q_2 = 0.5, q_3 = 0.999)$. The q_1 corresponds to the bottom scenario, where there is a 10% chance that the day's net demand will be less than this timeseries. Similarly, for the q_3 scenario at the top, only 1 in 1000 net demand realisations will exceed the q_3 scenario at each time-step.

As well as assigning appropriate NDFE values for each quantile-based scenario, the probability of each scenario's occurrence (ϕ_n) must also be found. Assuming that the cost function is a continuous function of the net demand forecast quantiles q , the quantile scenario probabilities are found via:

$$\phi_1 = 0.5 \left(\frac{q_2^2}{q_2 - q_1} \right) \quad (2.17a)$$

$$\phi_2 = 0.5 \left(q_3 - q_1 - \frac{q_1^2}{q_2 - q_1} \right) \quad (2.17b)$$

$$\phi_n = 0.5 (q_{n+1} - q_{n-1}) \quad n = 3, \dots, \mathcal{N} - 2 \quad (2.17c)$$

$$\phi_{\mathcal{N}-1} = 0.5 \left(q_{\mathcal{N}} - q_{\mathcal{N}-2} - \frac{(1 - q_{\mathcal{N}})^2}{q_{\mathcal{N}} - q_{\mathcal{N}-1}} \right) \quad (2.17d)$$

$$\phi_{\mathcal{N}} = 0.5 \left(\frac{(1 - q_{\mathcal{N}-1})^2}{q_{\mathcal{N}} - q_{\mathcal{N}-1}} \right) \quad (2.17e)$$

Thus with the probability of each scenario's occurrence known ($\phi_n \forall n$), as well as the half-hourly net demand forecast errors ($\eta \hat{D}_{t,n} \forall t, \forall n$), the updated cost function and demand balance constraints are:

$$\min_{\mathbf{u}} : \sum_{t \in \mathcal{T}} \left(\sum_{g \in \mathcal{G}} c_g^{st} \cdot \mathbf{v}_{g,t} + C P_g^1 \mathbf{u}_{g,t} + \sum_{n \in \mathcal{N}} \phi_n \left(\Delta \tau (c^{ls} \cdot \mathbf{p}_{t,n}^{ls} + c^{wc} \cdot \mathbf{p}_{t,n}^{wc}) + \sum_{g \in \mathcal{G}} \Delta \tau \cdot \hat{\mathbf{C}}_g(\mathbf{p}_{g,t,n}) \right) \right) \quad (2.18)$$

$$\mathbf{p}_{t,n}^{ls} - \mathbf{p}_{t,n}^{wc} + \sum_{g \in \mathcal{G}} (P_g^1 \mathbf{u}_{g,t} + \mathbf{p}_{g,t,n}) = D_t - W_t + \eta \hat{D}_{t,n} \quad \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (2.19)$$

Stochastic and Deterministic MIP Agents

The S-MIP agent is defined as cost function (2.18) along with constraints eqs. (2.5) to (2.10), eqs. (2.2) to (2.4) and (2.19). Thirteen net demand forecast scenarios are used, with corresponding quantiles of: [0.002, 0.01, 0.1, 0.3, 0.5, 0.6, 0.75, 0.9, 0.95, 0.99, 0.993, 0.9967, 0.999]. The distribution is discretized more finely over the high quantile range, as these represent higher than expected net demand which can significantly impact the cost function if load-shedding occurs.

By explicitly considering load-shed as a recourse action, and then taking actions to minimise the expected operating costs, the S-MIP implicitly optimises the reserve margin. However, almost all real-life implementations of MIP for UC are deterministic, where the reserve margin

Algorithm 1 Branch and Bound Algorithm

```

1: find an initial feasible solution and set the upper bound (UB) to its objective value
2: relax integrality constraints
3: while time-limit not exceeded and MIP gap > 0% do
4:   solve resultant continuous linear program(s)
5:   if first iteration then
6:     set lower bound (LB) equal to solution's objective value and go to step 12
7:   end if
8:   if solution is feasible and less than UB then
9:     UB equal to solution's objective value
10:  end if
11:  set LB to lowest leaf node objective value
12:  find MIP gap:  $\frac{UB-LB}{UB} \cdot 100\%$ 
13:  branch problem into sub-problems (nodes) by constraining some intelligently selected
    binaries to be 0 or 1. Nodes yet to be branched are referred to as 'leaves'
14: end while

```

is explicitly set.

The formulation can accommodate a deterministic reserve requirement to define the D-MIP agent. This is done by using only one scenario, with zero net demand forecast error ($N = 1$, $\phi_1 = 1$, $\eta\hat{D}_{t,1} = 0 \forall t \in \mathcal{T}$), alongside an explicit reserve requirement:

$$\sum_{g \in \mathcal{G}} r_{g,t} \geq R_t \quad \forall t \in \mathcal{T} \quad (2.20)$$

$$p_{g,t} + r_{g,t} \leq (\bar{P}_g - \underline{P}_g)u_{g,t} \quad \forall g \in \mathcal{G}, \forall t \in \mathcal{T} \quad (2.21)$$

In this chapter the D-MIP agent sets R_t equal to four times the net-demand forecast error standard deviation for all timesteps t .

MIP Implementation Details

The MIP is formulated in Python using PYOMO [51] and solved by Gurobi 9.5.2 [52]. Gurobi implements a branch-and-bound algorithm for solving MIPs, industry standard for all high-performance commercial solvers. The algorithm is summarised in Algorithm 1.

Step 4 involves solving multiple leaf LPs simultaneously, thus it is well suited to parallelisation across multiple threads.

2.3.2 Reinforcement Learning Agents

My co-author of paper [4], Patrick De Mars, formulated the UC problem as a Markov decision process (MDP) and created the agents to solve it [23]. A brief summary of how the agent works is included here. I deployed these agents and compared their efficacy in solving the RL problem with the equivalent MIP agents.

An MDP is defined as a 4-tuple (S, A, P, R) of a state space, action space, transition function and reward function, respectively. At each timestep t , corresponding to decision periods in the UC problem, the agent receives a partial observation of the state $s_t \in S$, and chooses an action $a_t \in A$ sampled from a policy $\pi(a_t|s_t)$. A new state s_{t+1} is returned, sampled from the transition function $P(s_{t+1}|s_t, a_t)$ as well as numeric reward $R(s_{t+1}, s_t, a_t)$. The task of RL is to learn the policy $\pi(a|s)$ that maximises the agent's long-run expected sum of rewards.

The MDP for the UC problem is summarised as follows (for a more detailed description, see [23]):

- States S : current generator up/down times, demand/ wind forecasts, demand/wind forecast errors, current timestep.
- Actions A : binary commitment decisions for the following timestep.
- Transition function $P(s_{t+1}|s_t, a_t)$: update generator up/down times; sample demand/wind forecast errors from ARMA processes; solve economic dispatch problem determining generator setpoints for demand/wind realisations.
- Reward function $R(s_{t+1}, s_t, a_t)$: calculate the operating costs inclusive of fuel, startups, wind shedding and load shedding.

The agent observes all elements of the state except the demand and wind forecast errors: this is in order to preserve the day-ahead property of the problem.

The MDP is realised in the simulation environment described in Section 2.2.2. During training, the RL agent randomly samples a day from the training data set and commits generators through actions a , aiming to maximise rewards r . We apply a log transformation to the rewards during training, dampening the extreme costs of load shedding which can lead agents to adopt overly conservative policies.

The training data includes 806 unique days of demand and wind forecasts. The demand data is taken from the National Grid Demand Data [53] from 2016-2019 and is scaled linearly with the number of system generators. We used wind data from Whitelee onshore wind farm [54] between 2016-2019. These data sources provide the mean demand and wind forecasts, the stochastic parameters of the ARMA processes from equation (2.12) are set manually and are identical for all days. Wind penetration was scaled to be between 20%-70% of daily energy requirements. A sample of 40 days was used to create a test dataset. These were omitted from the training process and then used to fairly assess the merits of the different agents. More detail on the test data and training process are found in [23] as well as sections 4.2.2 and 4.2.3 of reference [55].

Proximal policy optimisation (PPO) [56] is used to train the agent, as in previous research [43, 23]. Actor and critic neural networks are parameterised separately with deep feed-forward neural networks. Since an action a is a bit-string of length N for a system with N generators, the size of the action space is $|A| = 2^N$. As a result, conventional formulations of the actor as a neural network classifier with $|A|$ output nodes are not tractable except for small problem sizes. We formulate the neural network as a binary classifier which sequentially predicts each action bit a_i , corresponding to the commitment decision for generator i . When the minimum up/down constraints enforce a particular commitment decision a_i is fixed and thus i is incremented without classification occurring.

The optimal commitment decision of generator i is dependent on the commitment decisions of the other generators in the system. As such the binary classifier is also conditioned on a 'draft commitment' decision. This draft is a sequence of N bits. For bits a_j where $j < i$, the commitment decisions are already known. For bits a_j where $j > i$ the plants are assumed not to be committed and therefore $a_j = 0$ unless commitment is enforced minimum up time constraints.

Finally, the input vector also uses one-hot encoding to indicate the generator being classified at each forward pass.

Model-Free and Lookahead RL Agents

A trained agent can be used to solve unseen test cases by sampling from the policy $\pi(a|s)$ for new states s . To improve reliability during testing, the RL agent chooses actions using ‘ $\arg \max_a \pi(a|s)$ ’. We call this agent ‘model-free’ RL (RL-MF) as it does not rely on any lookahead strategy in decision-making.

In addition to the RL-MF agent, we also implemented a 1-step lookahead agent, RL-LA. Using the same policy as RL-MF, the RL-LA agent evaluates a subset of probable actions which meet a threshold ρ , subject to $0 < \rho < 1$. That is, all actions where $\pi(a|s) \geq \rho$ are explored. In our experiments, we set $\rho = 0.05$ allowing for a maximum of 20 actions to be explored, in addition to the ‘do nothing’ action which keeps all commitment decisions the same at the following timestep. For each of the actions in the subset, the agent evaluates operating costs over 100 scenarios of demand and wind sampled randomly, and chooses the action giving the lowest operating costs overall.

Due to the evaluation of multiple actions, the computational expense of evaluating the RL-LA agent is significantly larger than RL-MF, which produces solutions almost instantly. The run times of the two optimisation methods are compared in Section 2.4.2.

2.4 Case Studies and Comparison

Section 2.2 described the UC problem setup and simulation environment. Section 2.3 described the four optimisation agents: D-MIP, S-MIP, RL-MF and RL-LA. In this section, I present for the first time a detailed comparison of these solution methods, evaluating their performance in terms of operating costs and run time. In addition, a qualitative analysis of their characteristics is provided with case studies. The agents are evaluated on problem sizes of 10–50 generators using the environment described in Section 2.2.2, with 40 test days each. Due to the system size and generator characteristics (listed in [46]), the unit of power is MW and the unit of money

is \$.

Simulations were run on a 24-core (48 thread) 2.85GHz AMD EPYC 7443 processor, with the maximum threads per solve constrained to 8 and a time budget of 10 minutes.

2.4.1 Cost Comparison

Expected operating costs of solutions produced by each of the 4 methods are estimated as in (2.13), using $R_n = 5000$ scenarios, thus providing an accurate estimate of the true expected costs (2.1). In addition to the four optimisation agents, an agent with perfect foresight was also implemented. The difference between operating costs produced by the optimisation agents and perfect foresight represents the effective cost of uncertainty.

The mean daily operating costs of the four optimisation agents are shown in Table 2.1. Across all problem sizes of 10–50 generators, both MIP agents outperform the RL agents. As expected, S-MIP produces lower operating costs than D-MIP, by between 4.1–4.4%. Furthermore, RL-LA consistently outperforms RL-MF, benefiting from greater foresight and the ability to consider multiple actions. Overall, the S-MIP agent produces the lowest operating costs on all problem sizes.

Figure 2.2 shows the distribution of mean operating costs over the 40 case studies, normalised against the perfect foresight solutions. The RL-MF and RL-LA results are characterised by greater variability in performance across the different problem instances; for example, RL-MF solutions are up to 21 times more expensive than the perfect foresight costs. This is indicative of the limited ability of the RL agents to generalise effectively to unseen problem instances and the characteristic unreliability of RL as compared with MIP.

The outliers significantly skew the mean costs for the RL agents; thus, it is also valuable to analyse median operating costs over the 40 problem instances, more clearly visualised in the right-hand plot of Figure 2.2. RL-MF and RL-LA both outperform the D-MIP in the 10 generator system, and RL-LA performs similarly to D-MIP for 20 and 30 generator systems. While the RL agents outperform the D-MIP in many cases up to 30 generators, the solution quality is highly unreliable.

The RL agents show limited ability to scale to larger problem sizes, as shown by significant




Figure 2.2: Operating costs relative to perfect foresight costs for the four optimisation agents. The box encompasses the inter-quartile range, with a line at the median. The whiskers extend to the max and min values, but are limited to 1.5x the inter-quartile range. The left-hand plot shows operating costs for all 40 case studies; the right-hand plot is clipped to relative costs ≤ 1.5 . The extreme operating costs of RL-MF and RL-LA in some instances, indicate unreliable performance and difficulties in generalising to unseen problems. However, median operating costs of the RL agents are competitive with D-MIP up to 30 generators.

Table 2.1: Sum of mean operating costs for 40 case studies (\$m)

Generators	RL-MF	RL-LA	D-MIP	S-MIP
10	25.07	24.94	16.58	15.86
20	47.58	37.32	32.97	31.52
30	61.26	54.95	49.28	47.22
40	90.19	73.41	65.63	62.93
50	162.26	100.66	81.94	78.62

worsening performance in Figure 2.2. This trend exists to a much lesser extent for both MIP agents due to higher optimality gaps achieved within the 10-minute time budget. The limited ability of the RL agents to scale to large problems is driven by the exponential growth in the size of the action space, which doubles with each additional generator. Intelligently exploring the action space with RL becomes increasingly challenging as the problem size grows, and is not effectively handled by current methods [57].

2.4.2 Solve Time Comparison

With increasing numbers of generators, the solve times of the MIP methods increase significantly, as shown in Figure 2.3. MIP solves which time out after 10 minutes do not reach the 0%

Table 2.2: Number of Unfinished Solves

Method	Generators				
	10	20	30	40	50
Deterministic	0	0	1	5	13
Stochastic	0	1	10	12	18

MIP gap threshold, and hence are further away from optimality than those which complete. By contrast, the run time of RL-MF exhibits linear complexity and RL-LA remains roughly constant for all problem sizes. Run times of both RL agents remain low for all problem sizes, with maximum solve times of 2s and 17s for RL-MF and RL-LA, respectively.

The number of solves (maximum 40) which did not complete within 10 minutes are shown in Table 2.2 for both MIP agents. When the MIP gap is not reached the solver outputs the current feasible solution with the lowest objective value (the current UB, as detailed in Algorithm 1). For the 10 generator case, both D-MIP and S-MIP succeed in solving all 40 case studies with 0% MIP gap. The number of unfinished solves increases with problem size for both D-MIP and S-MIP. At 50 generators nearly half of S-MIP solves do not complete, with a maximum final MIP gap of 0.23%. Since lower MIP gaps correlate with lower cost solutions, the increasing number of unfinished solves with the number of generators indicates worsening relative solution quality as the problem size increases. However, as discussed in Section 2.4.1, the deterioration in solution quality with problem size for RL agents is much more severe than for MIP agents.

2.4.3 Schedule Comparison

The solutions produced by the four optimisation agents exhibit qualitative differences, as demonstrated by the 20 generator case studies in Figure 2.4. Case A shows a typical example of the hierarchy of operating costs, with S-MIP producing the lowest operating costs, followed by D-MIP, RL-LA and RL-MF. RL-MF is consistently the worst performing agent, exhibiting high levels of load shedding in cases C and D. In case B, RL-LA outperforms D-MIP, operating securely with less conservative reserve margins.

The ability to consider multiple actions at each timestep gives RL-LA a significant advantage over RL-MF in case C. The two sharp decreases in wind output around decision periods 10 and




Figure 2.3: Distribution of solve times for the optimisation agents over the 40 case studies. The RL agents have relatively stable solve times with increasing problem size, whereas the solve times for MIP rise super-linearly with the number of generators. Solve times are limited to a time budget of 10 minutes (600 seconds).

42 are managed poorly by RL-MF, resulting in high volumes of load shedding. By contrast RL-LA commits generators earlier during the morning wind drop and manages the later spike in net demand, resulting in 78% lower operating costs as compared with RL-MF.

In case D, both RL agents decommit baseload generation during a period of high wind generation which cannot be recommitted for several hours, resulting in high levels of load shedding during the evening peak. The RL solutions are over 5 times more expensive than MIP in this case.

In general, the RL solutions are characterised by long periods of little or no commitment changes, as shown in cases A and D. By contrast, the MIP agents follow the net demand profile more continuously. The S-MIP agent operates with the smallest reserve margins, accurately tracking the net demand profile and resulting in more efficient use of generation and the lowest




Figure 2.4: Selected 20 generator case studies, solved by the four optimisation agents. A subset of 100 scenarios of demand and wind are shown for each problem. In case A, both MIP methods outperform the two RL methods, with S-MIP producing the lowest operating costs overall. RL-LA outperforms RL-MF and employs small reserve margins. In case B, RL-LA outperforms D-MIP, committing less reserve. Case C is characterised by two large swings in wind generation; in this case, RL-LA is 78% cheaper than RL. In case D, both RL-MF and RL-LA perform poorly as compared with the MIP methods, with large levels of load shedding during the evening peak.

operating costs.

2.5 Hybridisation

We have shown that S-MIP achieves significantly lower operating costs than the RL agents. On the other hand, for larger system sizes S-MIP frequently utilises the full time budget of 10 minutes, whilst RL produces solutions in less than 2s. In this section, we leverage the high accuracy of S-MIP and the fast solve time of RL to produce a hybrid solution method that outperforms either individually.

The trained RL agents are used to provide useful starting points for warm starting the MIP solver. This replaces the generic (‘vanilla’) initial feasible solution method used in step 1 of Algorithm 1 with an improved domain-specific method. This section has two parts: the first shows how more computer resource results in lower MIP gaps, which correlates with lower cost schedules; second, we show the value of concurrently producing multiple distinct solutions with sufficiently low MIP gaps, and how the RL-MF agent can assist in doing so.

To the author’s best knowledge this is the first time that RL has been used to warm-start a UC problem solve. Encouragingly we find that in doing so our method outperforms the vanilla MIP solve over the range of computer resources tested.

2.5.1 Warm Start Method Comparison

Three different methods were used to produce initial feasible points to warm start the solver when solving the 50-generator UC problem:

1. ‘vanilla’: The generic Gurobi method.
2. RL: The pre-trained agents.
3. ‘rand’: Method that generates random feasible points by randomly constraining generators to be on and off throughout the day whilst respecting their constraints.

The RL and rand methods were each used to produce 8 different initial feasible solutions for each of the 40 test days considered. This is possible since each RL policy effectively defines a probability distribution over solutions for each case study. Thus sampling produces a diverse set of intelligently selected solutions that are guaranteed to be feasible.

Figure 2.5 plots the average MIP gap and uncertainty cost over the 40 test days when using the three warm-start methods. Uncertainty cost is the increase in a schedule’s cost over that of the perfect foresight schedule’s cost. For the RL and rand warm-start methods, the MIP was run 8 times for each test day, each time warm-started by a different one of the 8 distinct initial schedules. The points in Fig. 2.5 are the average cost over these 8 schedules for each of the 40 test days.




Figure 2.5: Average schedule cost and MIP gap over the 40 test days for the 3 warm starting methods with either 1-thread or 8-threads per solve. For the RL and rand methods, each of the 40 days was solved 8 times using 8 distinct initial solutions. More computer resources per solve result in lower average MIP gaps, which correlates with lower cost schedules when assessed via (2.13). However, this does not account for concurrency.

The computing resource available for each solve was either 1 or 8 threads. Comparing thread resources reveals the intuitive result that more threads available result in lower MIP gaps. This is due to the performance gains from parallelizing the solving of concurrent leaf LPs in step 4 of Algorithm 1. Lower MIP gaps are equivalent to a lower objective (2.18), which is an approximation of the more accurate cost function (2.13) that is used to calculate the uncertainty cost on the y-axis. Hence a low MIP gap correlates with a lower uncertainty cost.

Comparison between warm-start methods utilising 1-thread reveals the first instance when hybridisation outperforms the vanilla S-MIP alone. The schedules produced from the S-MIP method with RL warm-start are, on average, \$140 or 0.14% less expensive than the vanilla method alone. This indicates that the RL method produces better initial solutions for this problem type than the vanilla method. This improved starting point gives the RL solves an advantage, allowing them to reduce the MIP gaps faster and produce lower average cost schedules.

When 8 threads are available the vanilla and RL methods perform similarly because the increased computing power enables both methods to close the MIP gap faster. This increased

speed minimises the impact of the better initial solutions offered by the RL. However, the next section demonstrates how the multiplicity of solutions generated by concurrent 1-thread RL solves can be leveraged to outperform the vanilla method when all 8-threads are available.

2.5.2 Concurrent Solves

A low MIP gap indicates a lower objective value (2.18) which is an approximation of (2.13) that is used to find the more accurate uncertainty cost of the schedules. This approximation is intrinsic to all stochastic optimisation problems because of the necessary discretisation of the continuous probability distributions. Well-constructed scenario trees will minimise but never completely eliminate this misalignment of objective functions.

A schedule's uncertainty cost is the increase in it's expected cost over the perfect foresight expected cost. The fallout of the misalignment of (2.18) with (2.13) is a noisy correlation between the MIP gap and a schedule's true uncertainty cost. To demonstrate this, solutions with a range of MIP gap values were generated. To do this, for the 11 days where a 0% MIP gap was achieved with 1-thread and 50 generators, the solves were re-run multiple times with time budgets ranging between 10s - 600s. The costs of these solutions were found using (2.13), normalised (divided) by the 0% MIP gap value, and then plotted in Fig. 2.6. The noisy correlation between the MIP gap and the schedule's uncertainty cost can be clearly observed, where although the general trend is a lower MIP gap implying lower cost, there are many instances when a higher MIP gap solution produces a lower cost schedule. Indeed, there are three instances with MIP gaps higher than 0% where the uncertainty cost is below that of the 0% MIP-gap schedule, indicated by the points with normalised uncertainty costs less than 1. This demonstrates that (2.18) is an approximation of the true objective function, for if they were perfectly aligned, a lower MIP gap would always equal a lower cost schedule.

The value of multiple concurrent solutions is illustrated in Fig. 2.7 which plots the eight, 1-thread solutions for RL warm-start for an exemplary test day. Concurrently solving the 8, 1-thread RL and rand solutions produces 8 distinct solutions within the same low MIP gap neighbourhood. After applying (2.13) to the solutions, the cost function misalignment is exploited to give a range of costs, where the minimum (best schedule) is usually less than that




Figure 2.6: Cost vs MIP gap for days where the original solutions had 0% MIP gap with 1-thread. The optimisation was re-run with varied time budgets to produce a range of final MIP gaps. This demonstrates how the misalignment of the objective being optimised (2.18) with the true objective (2.13) results in a noisy correlation between MIP gap and a schedule’s true cost, i.e. some solutions with higher gaps produce lower cost schedules.

of the equivalent 8-thread cost. The ‘vanilla’ method only produces 1 initial feasible point, so cannot exploit the gains of solution multiplicity.

In Fig. 2.7 the solutions all have similar MIP gaps of 0.06–0.09%, but their schedule uncertainty costs range between \$80,200 to \$81,600. If concurrency was not considered, as in Fig. 2.5, then the mean 1-thread cost of \$80,720 would be expected, which is worse than the expected 8-thread cost of \$80,520. However, because all these solutions are produced within the time and computing budget, an operator can use (2.13) to choose the best solution using a more accurate metric than the optimisation’s cost function or MIP gap. Thus, in this case, the operator would choose the lowest cost schedule of \$80,200, which is an improvement over the expected 8-thread solve value despite having a 0.08% larger MIP gap.

Exploiting this multiplicity of solutions for each of the 40 test-days results in an average uncertainty cost over test days of \$99,140, a reduction of 0.3% over the previous best solutions of an individual 8-thread solve using the vanilla or RL warm-start method. These results are plotted in Fig. 2.8.

In summary, a hybrid scheduling method is proposed where the fast solve times of RL are




Figure 2.7: Concurrently warm starting multiple MIP solves with various RL solutions produces multiple low gap solutions within the time budget. Plotted here are the eight 1-thread solutions for 29/09/16. Selecting the lowest cost schedule according to (2.13) therefore exploits its slight misalignment with the MIP cost function (2.18). This finds schedules that are lower cost than the average of 1-thread and 8-thread solves.

leveraged to produce a variety of distinct schedules that are used to warm-start the highly accurate S-MIP. This method produces lower cost schedules than the incumbent ‘vanilla’ S-MIP over the range of computer resources considered. At 1-thread it produces lower MIP gap solutions, which on average perform better. At 8-thread available resources, the RL warm-start method creates value by producing a variety of distinct initial solutions, which are used to concurrently produce a variety of distinct low MIP gap solutions by the solver. So long as the final MIP gaps are sufficiently low, the randomness introduced by the misalignment of the MIP cost function (2.18) and the more accurate cost function (2.13) can be exploited by choosing the lowest cost schedule out of the 8. This reduces average operating costs when compared to the single 8-thread solve.

This shows that RL has good potential as a principled way of producing feasible initial solutions for the S-MIP solver, offering a significant improvement over random warm-start. To the best of our knowledge, this is the first time that RL solutions have been used to warm-start, and improve the performance of a MIP solver for the UC problem. The dependence of the hybrid method’s performance on the quality and variety in the initial RL solutions is a research area




Figure 2.8: Taking the lowest cost, RL warm started 1-thread solve for each test day results in schedules that are on average 0.3% lower cost than only solving a single 8-thread MIP using vanilla or RL warm start methods.

that should be further investigated, given the potential savings for computationally-expensive optimisation problems such as the stochastic UC.

2.6 Conclusion

This chapter has provided a comprehensive comparison of two predominant optimisation methods for a problem of critical importance in power systems, unit commitment. RL can in principle offer rigorous handling of uncertainty and fast decision-making and has conquered several long-standing challenges of artificial intelligence [58, 59]. However, our results show that currently RL is outperformed decisively by MIP methods and that significant improvements are still required to achieve competitive performance. A model-based lookahead strategy outperformed the model-free approach, albeit at greater computational cost, and was competitive with deterministic MIP for small problem sizes, but remained far below the quality of stochastic MIP.

The issues with RL are characterised in terms of 1) scalability and 2) reliability. Regarding

the former, the exponential growth of the action space with the number of generators is not handled well by existing RL methods [57], limiting its application to large problem sizes. Whereas median operating costs of model-free RL were 4% higher than stochastic MIP for a 10-generator problem, this increased to 44% for 50 generators.

Regarding reliability, the performance of the RL agents varied significantly across different problem instances. In the worst case, the operating costs of model-free RL and lookahead RL were 21 times that of stochastic MIP. Such large variability and extreme worst-case costs are unacceptable for practical applications. The MIP techniques consistently produced good quality solutions, and a measurable optimality gap is a significant advantage, providing certainty to system operators. Significant advances in core RL methods capable of handling high-dimensional discrete action spaces in security-critical contexts are required for this to be a practical alternative to MIP.

While the RL agents were not competitive with MIP, their strong performance on some problem instances indicated that they had learned insightful properties of the UC problem. Furthermore, they are able to produce a range of feasible solutions almost instantly regardless of problem size, whereas MIP solutions routinely utilised the full time budget considered. The main methodological contribution of this chapter was to leverage these RL properties to improve the MIP solve through warm starting. This method reduced operating costs by 0.3% as compared with vanilla stochastic MIP, whilst maintaining its high consistency and explainability properties. Interestingly, the improvement relies on exploiting the mismatch between the stochastic MIP cost function and the real objective by providing multiple low MIP gap solutions concurrently. The hybrid approach shows that RL has strong potential as a method for warm starting MIP solvers for UC.

This research highlights the value of knowledge transfer between machine learning and power systems optimisation communities to further the state-of-the-art. Future collaborations are encouraged to accelerate the development of novel UC solution methods capable of handling the high levels of uncertainty and complexity that will characterise future power systems. Consideration of how each method operates with only partial knowledge of the true distributions of the salient stochastic process is also important when considering their practical application.

Chapter 3

Frequency Response from Aggregated V2G Chargers Under Uncertain Electric Vehicle Connections

The previous chapter established that MIP is currently superior to RL for the UC problem. This chapter improves the MIP unit commitment formulation to include frequency security constraints. By explicitly incorporating uncertainty into the frequency dynamics, these constraints enable aggregated electric vehicles to provide frequency response. This unlocks an abundant future resource of ancillary services by providing explicit guarantees on system security. This chapter is based on the work of [5] and [2].

3.1 Introduction

Decarbonised future power systems will be characterised by low inertia due to the displacement of synchronous fossil fuel generators by converter interfaced generation like wind and solar. This makes grid frequency more volatile, thus more challenging to contain within predefined limits.

Post generator outage, frequency response (FR) is activated to provide a net power injection that aims to arrest frequency decline by restoring the power balance. In this chapter, the term FR exclusively refers to this primary FR service within the seconds after a generator loss. The

required amount of FR depends on the level of system inertia [25], which covers the transient power deficit by extracting kinetic energy stored within the rotating masses of synchronous generators. Previous work [28] has demonstrated that fast FR (delivery less than 1s) from converter interfaced resources, like grid batteries, is extremely effective at containing frequency nadir, and thus a vital resource to decouple frequency security from synchronous machines.

Large numbers of electric vehicles (EVs) will be present in future systems, with more than 23 million expected on the road in the UK by 2030 [60]. When paired with vehicle to grid (V2G) chargers, their smart control offers an abundant and valuable [5] source of FR. The small size and distributed nature of V2G chargers necessitate they be aggregated together into fleets. The capacity of FR from each fleet is determined by the number of connected EVs, which cannot be known exactly ahead of time. This means that unlike other inverter-based resources like grid-batteries, a fleet's FR capacity is intrinsically uncertain, thus its inclusion makes the dynamic frequency evolution post-outage uncertain.

Most of the current literature focuses on the optimal scheduling of secondary and tertiary frequency response services from EVs. These are steady state services with slower delivery times that are unaffected by low inertia levels and thus transient dynamics do not need to be considered. However, some deterministic methods for scheduling FR do already exist [5][61][62][63]. Reference [61] calculates the profit of FR from individual V2G-connected EVs by calculating the optimal charging schedule in relation to historical FR (delivery within 10s) prices and real EV connection data in Great Britain (GB). The method in [62] co-optimises an aggregators participation in the day-ahead energy market and the FR market. The only frameworks that unlock the maximum value of FR from EVs by considering fast FR (delivery within 1s) are presented in [5] and [63]. Reference [63] optimises a commercial fleet's trip times to maximise FR revenue, whilst [5] optimises fleet operation to increase renewable integration in low inertia systems. The chapter demonstrates that FR from bidirectional chargers is up to 20 times more valuable than unidirectional chargers.

All these methods assume that the number of EVs that are plugged in at the time of the outage is known. This number determines the power injection capacity of the aggregated fleets and in reality is impossible to know ahead of time. Thus treating it deterministically jeopardises

system dynamic security, risking system damage and load shed. To date, the literature that accounts for EV uncertainty [64][65][66] focuses on hourly energy requirements, which can only be used to manage slow (and less valuable), steady state ancillary services like reserve or dynamic frequency regulation.

This chapter offers significant improvements on the current state of the art scheduling methods by explicitly incorporating the intra-hour EV connection uncertainties (which is characterised from field charging data) into the system frequency dynamics. Two stochastic methods are presented that allow, for the first time, the scheduling of FR under this uncertainty. This contribution unlocks the substantial value of FR from EVs in future systems whilst maintaining guarantees on system dynamic security.

A simple ‘individual’ approach is presented that limits the scheduled FR from each fleet individually, similar to the approach that the UK’s system operator takes for aggregated FR providers [67]. The main contribution of this chapter is a second ‘joint’ approach that schedules an aggregate amount of FR across all the diverse EV fleets on the system. The second approach focuses on ensuring system dynamic security and offers significant improvements over the ‘individual’ approach that we demonstrate mathematically and with case studies.

The proposed ‘joint’ approach schedules FR from fleets of EVs with uncertain plugins using distributionally robust chance constraints (DR-CC). These allow low-probability violation of uncertain constraints for a set of possible probability distributions called an ambiguity set. DR-CC finds the balance between stochastic and robust approaches. It leverages distributional information like moment or unimodality knowledge, to result in less conservative results than robust programs, but requires less precise distributional knowledge than stochastic programs. Furthermore, many useful ambiguity sets facilitate highly tractable analytical convex reformulations.

Ambiguity set construction is generally categorised into two distinct approaches, moment based [65][66][68][69][70][71] and statistical distance based [72, 73], with some recent work seeking to combine the two [64]. Both approaches have seen widespread application within steady state energy system modelling, primarily to deal with renewable power generation forecast uncertainty [64][66][69][70][71] [72] [73].

Recently, moment based DR-CC are also increasingly being employed to deal with the intrinsic uncertainties of aggregated distributed resources (ADRs) [68, 69, 66, 65]. Reference [65] develops a method for distribution systems to mitigate their renewable power forecast uncertainty via aggregated EV charging. A model predictive control scheduling approach is used, with uncertain EV charging demands accounted for via moment based DR-CC. DR-CCs are used in [66] to facilitate the provision of fast reserves from aggregated behind-the-meter loads (including EVs and water heaters). Uncertain energy and power constraints on reserve are considered, with the option to exploit distributional unimodality information to tighten the ambiguity set. Reference [68] applies DR-CCs to schedule reserve from aggregated air-conditioning loads with uncertain reserve capacity limits within an optimal power flow problem. Ambiguity sets considering exact and approximate second moment information are used, which result in a Second-Order Cone (SOC) Program and a Semi-Definite Program respectively. Finally, Bachi *et al.* [69] apply the conic reformulation of a two-sided linear DR-CC with known second order moments, first derived in [71], to line loading and nodal voltage constraints under uncertain renewable outputs and uncertain ADR energy demand, revealing the impact of network constraints on ADR's bidding strategy in the day ahead electricity markets.

However, the above references only utilise the flexible demand of ADRs (such as EVs) to provide reserve for steady state power balancing. Non consider the ability of ADRs to assist in the dynamic problem of securing frequency in the transient period immediately following the loss of a large generator. This chapter establishes a DR-CC method to optimally schedule FR from aggregated EV fleets. The operator can specify frequency security violation probability, allowing the preferred trade-off between system risk and operational cost reduction to be found. To the best of our knowledge, this is the first work to explicitly evaluate the impact of uncertainty on the value of FR from distributed providers. Although this chapter exclusively focuses on EVs, the presented method is also applicable to other ADRs.

The main contributions of this work are:

1. To investigate the impact of EV connection uncertainty on their provision of primary (fast) FR in the transient period immediately following an outage. The uncertainty is incorporated into the system frequency dynamics, from which probabilistic nadir and

steady state security frequency requirements are derived.

2. To propose a novel, convex moment-based DR-CC on the maximum scheduled FR from V2G chargers. This convexifies the probabilistic frequency security constraints whilst enabling the valuable scheduling of FR from V2G-connected EVs in an efficient and risk-limited manner.
3. To analyse field EV fleet connectivity data to provide realistic parameters and EV connection forecasts, as well as guide ambiguity set selection.
4. To provide new insight, based on simulating the yearly operation of the GB 2030 system, on the the value of FR from V2G and its sensitivity against renewable generation penetration, grid battery penetration, communication delays and uncertainty levels.

This chapter is organised as follows: Section 3.2 derives a convex formulation for probabilistic frequency security constraints. Field EV fleet connectivity data is analysed in Section 3.4 to inform EV connectivity forecasting and parameter selection, whilst Section 3.5 presents case studies exploring the value of FR from aggregated V2G chargers. Section 3.6 gives the conclusions.

3.2 Modelling of Aggregated V2G Chargers in Frequency Dynamics

This section derives frequency security constraints from the dynamic swing equation, which are non-deterministic due to the uncertain FR capacity from V2G chargers. The proposed DR-CC formulation for these constraints is presented, along with the virtual battery model of aggregate fleet charging.

These constraints are widely applicable to any convex optimisation problem, for markets, scheduling as otherwise. Consequently, and in order to reduce the notational burden, the t and n indexes used to apply these constraints within the SUC detailed in Section 3.3 are dropped throughout this section.

3.2.1 Frequency Security Constraints Under Uncertainty

System frequency evolution post generator loss is accurately approximated by the single machine swing equation [74]:

$$\frac{2H}{f_0} \frac{d\Delta f}{dt} = \mathbf{R}_1^{EV}(t) + \mathbf{R}_1^{ND}(t) + \mathbf{R}_2^G(t) - \mathbf{P}L_{max} \quad (3.1)$$

Load damping is neglected as the level in future systems dominated by power-electronics will be much reduced [29]. Thermal plants are grouped by generator types. The system inertia is determined by the number of committed thermal plants:

$$\mathbf{H} = \sum_{\hat{g} \in \hat{G}} H_{\hat{g}} \cdot P_{\hat{g}}^{max} \cdot \mathbf{N}_{\hat{g}} \quad (3.2)$$

The formulation is compatible with binary commitment variables for each individual generator. However, previous work [20] has shown that due to the large number of generators considered, grouping the binary commitment variables within each generator group into one continuous commitment variable $\mathbf{N}_{\hat{g}}$ significantly improves problem solve times with a negligible impact on results [20]. This approach is consistent with the literature [25, 28, 5].

FR dynamics are modelled as linear ramps, similar to the work in [29, 25, 28, 31]. Detailed dynamic simulations carried out in Section III of [31] show that droop controls can be accurately and conservatively approximated by a ramp. More detailed dynamic models prohibit closed form solutions to (3.1), necessary in order to derive convex algebraic frequency security constraints.

$$\mathbf{R}_1^{EV}(t) = \begin{cases} \frac{\mathbf{R}_1^{EV}}{T_1} \cdot t & t \leq T_1 \\ \mathbf{R}_1^{EV} & t > T_1 \end{cases}, \mathbf{R}_1^{ND}(t) = \begin{cases} \frac{\mathbf{R}_1^{ND}}{T_1} \cdot t & t \leq T_1 \\ \mathbf{R}_1^{ND} & t > T_1 \end{cases} \quad (3.3)$$

$$\mathbf{R}_2^G(t) = \begin{cases} \frac{\mathbf{R}_2^G}{T_2} \cdot t & t \leq T_2 \\ \mathbf{R}_2^G & t > T_2 \end{cases} \quad (3.4)$$

In this thesis $T_1 < T_2$. The slower speed models governor controlled FR from thermal plants. The faster speed comes from power-electronic devices, including V2G chargers and non-distributed

devices like grid batteries.

Ahead of time, the dynamics of the cumulative FR delivered from V2G chargers is known. However, the magnitude of delivered FR is uncertain because this is proportional to the number of connected EVs which cannot be known ahead of time. Thus, the grid frequency dynamic is also uncertain. For an individual fleet the response capacity is determined by the charging decisions of that fleet (decision variables), and the number of EVs that are connected (random variable):

$$\mathbf{R}_{1,i}^{EV} = \underbrace{(D_{max,i} - \mathbf{D}_i + \mathbf{C}_i)}_{\mathbf{g}_i} \cdot (N_{0,i} + \Delta N_i) \quad (3.5)$$

The cumulative magnitude of FR from all fleets is:

$$\mathbf{R}_1^{EV} = \sum_{i \in I} \mathbf{R}_{1,i}^{EV} \quad (3.6)$$

It is assumed that charging decisions for EVs within the same fleet are uniform. The number of currently connected EVs N_0 is known. The net EVs connected between now and the time of generator outage (ΔN) can be forecast, but not known exactly in advance. ΔN (and expressions of it) are the sole random variable within the frequency dynamics considered here. This chapter presents a stochastic framework to incorporate FR from aggregated V2G chargers, whilst explicitly limiting the risk of a frequency security breach due to the potential under-delivery of FR from EVs.

RoCoF Constraint

The maximum RoCoF occurs at the moment of $\mathbf{PL}_{max}$ outage. At this time no response has been delivered so it is deterministic and limited by inertia alone. Constraining the maximum RoCoF is necessary to prevent RoCoF-sensitive protection systems from disconnecting distributed generation and exacerbating the deficit. Setting $t = 0$ in (3.1) results in:

$$\frac{2|\text{RoCoF}_{max}|}{f_0} \cdot \mathbf{H} \geq \mathbf{PL}_{max} \quad (3.7)$$

Steady State

Frequency drop will be arrested if the sum of FR is greater than the largest loss. Ensured to a high certainty with:

$$\mathbb{P}\left[PL_{max} \leq R_1^{ND} + R_1^{EV} + R_2^G\right] \geq 1 - \epsilon \quad (3.8)$$

Nadir Constraint

Here it is assumed that the nadir occurs after T_1 ($\approx 1s$) as the extremely low inertia required to breach the frequency limit Δf_{max} ($\approx -0.8Hz$) would violate the RoCoF constraint (3.7) for realistic power system parameters. The frequency nadir occurs at the instant of zero RoCoF. According to (3.1) this is:

$$t^* = \frac{[PL_{max} - (R_1^{ND} + R_1^{EV})] \cdot T_2}{R_2^G} \quad (3.9)$$

It is shown in [28] that by integrating (3.1) and then substituting in (3.9), the nadir constraint can be formed as a convex rotated SOC. Thus the post outage frequency drop is contained with high assurance via:

$$\mathbb{P}\left[\underbrace{\left(\frac{H}{f_0} - \frac{(R_1^{ND} + R_1^{EV}) \cdot T_1}{4\Delta f_{max}}\right)}_{=z} \underbrace{\frac{R_2^G}{T_2}}_{=x} \geq \underbrace{\left(\frac{PL_{max} - (R_1^{ND} + R_1^{EV})}{2\sqrt{\Delta f_{max}}}\right)^2}_{=y}\right] \geq 1 - \epsilon \quad (3.10)$$

Finally, the power injection from aggregated chargers may be delayed due to communication or frequency measurement lag. An additional term in the nadir constraint can account for this [31]:

$$z = \frac{H}{f_0} - \frac{(R_1^{ND} + R_1^{EV}) \cdot T_1}{4\Delta f_{max}} - \frac{R_1^{EV} \cdot 2T_{del}}{4\Delta f_{max}} \quad (3.11)$$

3.2.2 Convex Reformulation of Chance Constraints

The inclusion of response from aggregated EV's within the frequency dynamics make (3.8) and (3.10) non deterministic. Therefore they cannot be applied to optimisations within scheduling or market contexts. This severely limits their use and motivates the need for their convex and

deterministic reformulation. Making (3.10) deterministic is challenging as no convex analytical reformulation of a chance constrained SOC currently exists.

Overcoming this mathematical dead-end to produce a convex reformulation of (3.10) and (3.8) that maintains the guarantees on frequency security under uncertain FR delivery is the main methodological contribution of this chapter. It is achieved via the insight that, because $\mathbf{R}_1^{EV}$ is the only non deterministic parameter within (3.8) and (3.10), ensuring that they are met with ‘ $(1 - \epsilon)\%$ ’ certainty is equivalent to scheduling an amount of response from EVs ($\bar{\mathbf{R}}_1^{EV}$) that will be delivered with ‘ $(1 - \epsilon)\%$ ’ certainty. This is found via:

$$\mathbb{P}\left[\bar{\mathbf{R}}_1^{EV} \leq \sum_{i \in I} \mathbf{R}_{1,i}^{EV}\right] \geq 1 - \epsilon \quad (3.12)$$

$\bar{\mathbf{R}}_1^{EV}$ replaces $\mathbf{R}_1^{EV}$ in (3.8) and (3.10), making them deterministic. The argument within chance constraint (3.12) is linear, hence if the mean and standard deviation of the forecasted number of connected EVs is known, it can be analytically reformulated into a SOC [70]. Thus a deterministic convex formulation is achieved. Start by defining a new scalar random variable:

$$\delta = \bar{\mathbf{R}}_1^{EV} - \sum_{i \in I} \mathbf{g}_{i\cdot} (N_{0,i} + \Delta N_i) \quad (3.13)$$

Note, δ represents the amount of FR scheduled over the FR that is actually available. It is desirable for this to be negative with a high probability. In a deterministic system ΔN_i is known, $\bar{\mathbf{R}}_1^{EV} = \sum_{i \in I} \mathbf{R}_{1,i}^{EV}$ and thus $\delta = 0$.

Substituting (3.13) into (3.12):

$$\mathbb{P}\left[\delta \leq 0\right] \geq 1 - \epsilon \quad (3.14)$$

Assuming independent EV connections between fleets, the mean and standard deviation of δ are:

$$\mu(\delta) = \bar{\mathbf{R}}_1^{EV} - \sum_{i \in I} \mathbf{g}_{i\cdot} (N_{0,i} + \mu_i), \quad \sigma(\delta) = \sqrt{\sum_{i \in I} (\mathbf{g}_{i\cdot} \sigma_i)^2} \quad (3.15)$$

Subsequently, δ can be scaled to have a zero mean and unit variance via $\delta_n = [\delta - \mu(\delta)]/\sigma(\delta)$:

$$\mathbb{P}\left[\delta_n \leq \frac{-\mu(\delta)}{\sigma(\delta)}\right] \geq 1 - \epsilon \quad (3.16)$$

The cumulative distribution function ($F_{\mathcal{P}}(k)$) of δ_n gives the probability that δ_n takes a value less than or equal to some constant k :

$$F_{\mathcal{P}}(k) = \mathbb{P}[\delta_n \leq k] \quad (3.17)$$

We now consider the DR-CC formulation. An important strength of this method is that the exact and true distribution $\mathcal{P}$ of δ_n does not need to be known. The set of possible distributions that $\mathcal{P}$ might belong to is called an ambiguity set ($\mathbb{A}$), defined by the distributional assumptions made on $\mathcal{P}$. However, because $\mathcal{P}$ is not known exactly, the exact form of its cumulative distribution is also unknown, inhibiting the reformulation of (3.16). Following the method presented in [70], this problem can be overcome by defining a lower bound on $\mathcal{P}$'s cumulative distribution function ($f_{\mathcal{P}}(k)$):

$$f_{\mathcal{P}}(k) = \inf_{\mathcal{P} \in \mathbb{A}} F_{\mathcal{P}}(k) \quad (3.18)$$

Given that $F_{\mathcal{P}}(k) \geq f_{\mathcal{P}}(k) \forall k$, then its substitution into (3.16) maintains the inequality and guarantees that the scheduled response is deliverable with at least the specified certainty level:

$$F_{\mathcal{P}}(k) \geq f_{\mathcal{P}}(k) \geq 1 - \epsilon \quad (3.19)$$

Given that the function of $f_{\mathcal{P}}(k)$ is increasing, it has a well defined inverse $f_{\mathcal{P}}^{-1}(\lambda)$. Thus constraint (3.16) can be written as:

$$\frac{-\mu(\delta)}{\sigma(\delta)} \leq f_{\mathcal{P}}^{-1}(1 - \epsilon) \quad (3.20)$$

Given that the RHS of (3.20) is a constant, this constraint is a convex SOC after substituting

in the moments of δ_n (3.15):

$$\sqrt{\sum_{i \in I} (\underbrace{g_i \cdot \sigma_i}_{\psi_i})^2} \leq \underbrace{\frac{1}{f_{\mathcal{P}}^{-1}(1 - \epsilon)} \cdot \left(\sum_{i \in I} g_i \cdot (N_{0,i} + \mu_i) - \bar{R}_1^{EV} \right)}_{\omega} \quad (3.21)$$

For some instances when risk aversion is high and the standard deviations of ΔN_i are large, constraint (3.21) can be infeasible. To ensure feasibility during simulation, a binary variable is added to the right hand side. Using the big-M technique, it simultaneously relaxes (3.21) and constrains $\bar{R}_1^{EV}$ to zero.

$$\sqrt{\sum_{i \in I} (\psi_i)^2} \leq \underbrace{\omega + M_1 \cdot b}_{\omega'} \quad (3.22)$$

$$\bar{R}_1^{EV} \leq M_2 \cdot (1 - b) \quad (3.23)$$

The exact form of $f_{\mathcal{P}}^{-1}(\lambda)$ depends on the distributional assumptions made about $\mathcal{P}$. Here, all ambiguity sets assume knowledge of the first two moments of δ . According to (3.15), this requires the true mean and standard deviation of forecasted EV connection numbers (μ_i, σ_i) for each fleet to be known. Three distinct ambiguity sets are considered:

1. Distributionally Robust Optimisation (DRO) - When only the mean and standard deviation of δ are known.
2. Unimodal - The distribution of δ_n is assumed to have a single peak. Likely when ΔN_i are themselves unimodal.
3. Gaussian - The distribution of δ_n is assumed to be Gaussian. True when ΔN_i are themselves Gaussian.

More statistical information on $\mathcal{P}$ defines tighter ambiguity sets and thus results in a less conservative chance constraint. The distributional assumptions made on δ define the form of $f_{\mathcal{P}}^{-1}(1 - \epsilon)$, according to the inverse cumulative distribution for the Gaussian case, or probability inequalities for the Unimodal and DRO case. The exact expressions for $f_{\mathcal{P}}^{-1}(1 - \epsilon)$ are shown in Table 3.1.

Table 3.1: Expressions for $f_{\mathcal{P}}^{-1}(1 - \epsilon)$

Gaussian	$f_{\mathcal{P}}^{-1}(1 - \epsilon) = \Phi^{-1}(1 - \epsilon)$
Unimodal	$f_{\mathcal{P}}^{-1}(1 - \epsilon) = \begin{cases} \sqrt{\frac{4}{9\epsilon} - 1} & \text{for } 0 \leq \epsilon \leq \frac{1}{6} \\ \sqrt{\frac{3(1-\epsilon)}{1+3\epsilon}} & \text{for } \frac{1}{6} \leq \epsilon \leq 1 \end{cases}$
DRO	$f_{\mathcal{P}}^{-1}(1 - \epsilon) = \sqrt{\frac{1-\epsilon}{\epsilon}}$

A great strength of our proposed constraint to schedule response from EVs under uncertainty is that it results in a convex programme. Post substitution of $\bar{\mathbf{R}}_1^{EV}$ into (3.22) and (3.10), the RoCoF (3.7) and the steady-state (3.8) constraints are linear. The deterministic nadir constraint (3.10) is a rotated SOC of the form $\mathbf{z} \cdot \mathbf{x} \geq \mathbf{y}^2$. Finally the ‘joint’ constraint to limit $\bar{\mathbf{R}}_1^{EV}$ (3.21) is a standard SOC of the form $\|\boldsymbol{\psi}\|_2 \leq \boldsymbol{\omega}'$ [75], where $\boldsymbol{\psi}$ is a vector of $\boldsymbol{\psi}_i$. Convexity allows it to be applied to a wide range of market and scheduling problems and solved in polynomial time by widely available commercial solvers with guaranteed convergence. It also allows the use of dual variables for shadow pricing.

3.2.3 Comparison to Simplistic Risk-Aware Scheduling Method

This is the first work to present a framework to schedule fast FR from aggregated EVs under connection uncertainty, so a comparison to state-of-the-art is not straightforward. However, some operators already allow the participation of aggregated resources in FR markets [67] under strict individual deliverability guarantees. This is formulated mathematically, for the first time, by applying (3.12) to schedule response from each fleet individually with ‘ $(1 - \epsilon)\%$ ’ deliverability probability:

$$\mathbb{P} \left[\bar{\mathbf{R}}_{1,i}^{EV} \leq \mathbf{g}_{i \cdot} (N_{0,i} + \Delta N_i) \right] \geq 1 - \epsilon \quad \forall i \in \mathcal{I} \quad (3.24)$$

The sum of these equals the cumulative FR from EVs on the system:

$$\bar{\mathbf{R}}_1^{EV} = \sum_{i \in I} \bar{\mathbf{R}}_{1,i}^{EV} \quad (3.25)$$

Due to only one uncertainty (ΔN_i) being present in each constraint of (3.24), according to (3.21) their reformulation is linear:

$$\mathbf{g}_i \cdot \sigma_i \leq \frac{1}{f_P^{-1}(1 - \epsilon)} \cdot \left(\mathbf{g}_i \cdot (N_{0,i} + \mu_i) - \bar{\mathbf{R}}_{1,i}^{EV} \right) \quad \forall i \in \mathcal{I} \quad (3.26)$$

Application of (3.26) is referred to as the ‘individual’ method, and represents the current state of the art.

The primary concern of an operator is to guarantee that the frequency is contained with a high probability. For the ‘individual’ method, the relationship between the choice of ‘(1 - ϵ) %’ (e.g. 99%) for individual fleets and the guarantee on the total system dynamic security is not known. Indeed if they decrease the risk of an individual fleet’s under-delivery of FR past the desirable system level, the guarantee of system dynamic security is lost. For this reason, the same ϵ value is used when comparing the methods.

The improvement of our proposed ‘joint’ method over the incumbent ‘individual’ method can be mathematically quantified by subtracting the scheduled FR ($\bar{\mathbf{R}}_1^{EV}$) in (3.25) from the amount scheduled using (3.21):

$$\Delta \bar{\mathbf{R}}_1^{EV} = f_P^{-1}(1 - \epsilon) \cdot \left(\sum_{i \in I} \mathbf{g}_i \sigma_i - \sqrt{\sum_{i \in I} (\mathbf{g}_i \sigma_i)^2} \right) \quad (3.27)$$

For any set of positive real numbers the root of the sum of squares will always be less than or equal to the sum of those numbers. Thus $\Delta \bar{\mathbf{R}}_1^{EV}$ will always be greater than or equal to zero. In real terms this means that our proposed method will always schedule more response from the system’s EVs than the current state of the art whilst maintaining dynamic security.

Conceptually this is because the ‘joint’ method compared to the ‘individual’ method leverages the fact that on the rare occasions ($\approx 1\%$) when one fleet has many fewer EVs plugged in than forecast and under delivers FR, the other fleets are ($\approx 99\%$) likely to compensate by over delivering their scheduled FR. Thus the ‘joint’ method’s focus on the system dynamic security makes better use of the EV FR resource for more efficient operation.

3.2.4 State of Charge and Reserve

Contrary to when scheduling FR from EVs, scheduling charging uses hourly EV connection numbers. These are assumed deterministic and known. This is justified because: 1) As shown in Section 3.5 and [5], an EV's value in highly renewable systems is dominated by its FR provision. Thus, characterising the impact of uncertainty on EV value is unhindered. 2) Despite constraint (3.21) being fully compatible with charging-under-uncertainty methods (such as that presented in [66]), deterministic EV charging is used here to simplify modelling and increases insight into system operation under (3.21), the core contribution of the chapter.

Here, each EV fleet is modelled as a virtual aggregate battery. All fleets are modelled in the same way, so the subscript i is dropped for notational clarity. A battery's charge rate is equal to the sum of all the individual constituent EVs:

$$\mathbf{p}^{EV} = (N_0 + \Delta\hat{N}_t) \cdot (\mathbf{D} - \mathbf{C}) \quad (3.28)$$

$\Delta\hat{N}_t$ is the net change in EV numbers between now and the beginning of the scheduling period t . It is deterministic and read into the simulation via a time series. EV connections are discretized, occurring at the start of each hour, then remaining constant for the entire timestep. This makes $\Delta\hat{N}_t$ different to ΔN in (3.13), which is a random variable, representing the instantaneous number of EVs connected. ΔN varies within the hour time-step. This is necessary because response capacity is determined by the instantaneous net power injection capacity of fleets whereas a virtual battery's generation and state of charge relate to the averaged connectivity values.

A battery's state of charge at the end of a scheduling period depends on: that period's charge decisions; its parent node's state of charge at the end of the previous timestep ($\mathbf{E}_{t-1}$); and change in charge incurred by EV (dis)connection.

$$\mathbf{E}_t = \mathbf{E}_{t-1} + \Delta\tau(N_0 + \Delta\hat{N}_t) \cdot (\eta\mathbf{C} - \frac{1}{\eta}\mathbf{D}) + \Delta N_t^{in} E^{in} - \Delta N_t^o E^o \quad (3.29)$$

Note that $\Delta N_t^{in}, \Delta N_t^o$ refer to the number of EVs that connect and disconnect at the beginning




Figure 3.1: Graphic representation of the input/output flow for the SUC. To simulate annual system operation, system variables are updated and the optimisation is iterated every hour.

of timestep t respectively. Thus $\Delta \hat{N}_t$ is their cumulation between now and the scheduling time period $\Delta \hat{N}_T = \sum_{t=1}^T (\Delta N_t^{in} - N_t^o)$. An EV's (dis)connection states of charge (E^o, E^{in}) are assumed known. Note $E^o \geq E^{in}$, thus $E^o - E^{in}$ represents the EV driving energy expenditure.

3.3 Stochastic Unit Commitment

A pre-existing MIP UC formulation is significantly enhanced to optimally co-ordinate charge and generation decisions in light of uncertain future EV connections. The model uses a similar SUC formulation to that detailed in Section 2.2.2. Both formulations use the quantile-based scenario tree method presented in [19]. This model branches into distinct scenarios at the root node only which was established to be superior to current RL solutions in the previous chapter.

This model is used for simulations of annual system operation under different constraints on using FR from V2G connected EVs. These simulations demonstrate: the value of our proposed formulation; insight into its impact on system operation (e.g. change in wind curtailment); and validates the frequency security risk guarantees. The SUC model optimally schedules generator and storage actions to provide reserve, response (fast and slow), inertia and energy production under uncertain renewable output over a 24hr period. Fig. 3.1 represents the process graphically.

This section details the advanced stochastic unit commitment (SUC) model. Any differences with the formulation in 2.2.2 are specified and how rolling-planning is used to enable simulation of annual operations rather than the previous static day ahead implementation.

3.3.1 System Constraints

User defined quantiles of the random variable of net demand (demand net of wind power) are used to construct a scenario tree. The demand forecast is considered deterministic due to its significantly lower standard deviation than wind forecast uncertainty. Reference [20] demonstrates that well chosen quantiles that branch only at the root node can discretize the continuous range of potential wind realisations accurately whilst yielding a highly tractable model. The model also includes a small amount of deterministically modelled solar power.

A single-bus power system model is used. The justification for this is two-fold: 1) It is common to solve the commitment and dispatch problems sequentially, adjusting the initial solution to account for line-flows. Here the first stage problem is modelled. 2) Frequency is a system-wide quantity, so insight into the contribution of FR from V2G connected EVs does not change significantly by including multiple buses. Each node has a complete set of decision variables, subject to the power balance constraint:

$$\sum_{\hat{g} \in \hat{G}} \mathbf{P}_{\hat{g},t,n} + \sum_{s \in S} \mathbf{p}_{s,t,n} + \sum_{i \in I} \mathbf{p}^{EV}_{i,t,n} + (W_t + \eta W_{t,n}) - \mathbf{p}^{wc}_{t,n} + S_t - \mathbf{p}^{sc}_{t,n} = D_t - \mathbf{p}^{LS}_{t,n} \quad \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.30)$$

The probability of reaching a given scenario weights the cost function:

$$\sum_{n \in \mathcal{N}} \phi_n \sum_{t \in \mathcal{T}} \left(\Delta \tau (c^{LS} \mathbf{p}^{LS}_{t,n}) + \sum_{\hat{g} \in \hat{G}} C_{\hat{g}}(\mathbf{P}_{\hat{g},t,n}) \right) \quad (3.31)$$

3.3.2 Generator and Storage Constraints

The SUC described in this section is used to run case studies on representative GB 2030 systems with ≈ 140 generators. Generators are grouped by type and their commitment decisions are relaxed to be continuous. This significantly speeds up computation with minimal loss of optimality and feasibility [20]. Commitment decisions are not shared across scenarios. This rolling-planning approach differs from the day-ahead formulation used in Section 2.2.1, where

only 1 commitment schedule was found. However, there is only one (base) node at the current timestep that is shared across all scenarios whose commitment decisions are implemented when the simulation is rolled forwards.

The comprehensive list of generator constraints is found in the Appendix of [25]. The constraints include minimum up and down times and startup delay. A simplified linear generator cost curve replaces the piece-wise cost curve (2.2) in Section 2.2.1 :

$$\mathbf{C}_{\hat{g}}(\mathbf{P}_{\hat{g},t,n}) = c_g^{st} \mathbf{V}_{\hat{g},t,n} + \Delta\tau(c_{\hat{g}}^{nl} \mathbf{U}_{\hat{g},t,n} + c_g^m \mathbf{P}_{\hat{g},t,n}) \quad \forall \hat{g} \in \hat{\mathcal{G}}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.32)$$

$$\mathbf{U}_{\hat{g},t,n} \underline{P}_{\hat{g}} \leq \mathbf{P}_{\hat{g},t,n} \leq \mathbf{U}_{\hat{g},t,n} \overline{P}_{\hat{g}} \quad \forall \hat{g} \in \hat{\mathcal{G}}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.33)$$

Technology agnostic storage is also modelled in the SUC:

$$\underline{E}_s \leq \mathbf{E}_{s,t,n} \leq \overline{E}_s \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.34)$$

$$0 \leq \mathbf{p}_{s,t,n}^D \leq \overline{P}_s \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.35)$$

$$0 \leq \mathbf{p}_{s,t,n}^C \leq \overline{P}_s \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.36)$$

$$\mathbf{p}_{s,t,n} = \mathbf{p}_{s,t,n}^D - \mathbf{p}_{s,t,n}^C \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.37)$$

$$\mathbf{E}_{s,t,n} = \mathbf{E}_{s,t-1,n} + \Delta\tau(\eta_s \mathbf{p}_{s,t,n}^C - \frac{1}{\eta_s} \mathbf{p}_{s,t,n}^D) \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T}, \forall n \in \mathcal{N} \quad (3.38)$$

Where $\mathbf{p}^D$ and $\mathbf{p}^C$ are the storage discharge and charge variables respectively, units GW. Note that the charge decision variables are from the system side, so they are positive when the storage discharges.




Figure 3.2: Schematic of a typical scenario tree used in SUC.

3.3.3 Rolling Planning

An iterative, rolling planning approach is used to simulate annual system operation. Decisions that minimise the expected operating cost over the next 24h period are found. The decisions at the current root node (that are shared between all scenarios) are implemented and the system is rolled forward by an hour, updating system states as well as wind and EV connectivity realisations. With this new information, the scenario tree is updated and the process is iterated. Figure 3.2 outlines the process graphically, where the index n differentiates between scenarios and the t index moves along them temporally on the x axis.

3.4 EV Connectivity Forecasting and Data Analysis

Scheduling FR from aggregated V2G chargers requires forecasting the number of connected EVs, which is equivalent to accurate characterisation of the distributions of ΔN_i . This is important for two main reasons: 1) The analytical reformulations of (3.21) and (3.26) require knowledge of the true mean and standard deviation (std); and 2) The ΔN distributions inform ambiguity set selection.

Constraint (3.21) translates the specified risk of under delivery to a scheduled amount of FR. It is compatible with any forecasting technique that ascertains the mean and std of ΔN . Here, a simple forecasting technique to characterise ΔN s using real data is adopted. The future connectivity is assumed well characterised by data from a similar time of the week during the




Figure 3.3: Typical Fri-Sat normalised EV connection numbers for the ‘work’ and ‘domestic’ fleets. Nighttime periods are shaded purple.

previous year. This is to demonstrate how forecast outputs are translated into operational inputs. It is expected that in actual operation more advanced forecasting techniques will be employed.

3.4.1 Forecasting Technique

Real, open-source EV fleet charging data [76] for 2017 is used to characterise two distinct fleets, ‘domestic’ and ‘work’. The domestic fleet relates to 3.2m charging events across 25,000, 10kW chargers installed in people’s homes. The work fleet relates to 103,000 charging events across 540, 20kW chargers installed in car parks of public sector buildings. Most of the individual chargers only provided data for specific months of the year. To account for this, the number of active chargers each month was found, ‘active’ defined as at least 2 charge events per week. The true fleet size was then taken as the average number of active monthly chargers, 8,500 and 200 chargers for domestic and work, respectively.

A charging event records charge-point identity, EV connection and disconnection time. From these the continuous annual time-series of the number of connected EVs is derived. Fig. 3.3 plots a typical Fri-Sat period for the two fleets. The domestic fleet is characterised by EVs disconnecting in the morning and then reconnecting in the afternoon. The work fleet exhibits the opposite trend during the week. Very few EVs connect to the work chargers on weekends.

The continuous time-series was discretized into average hourly connection and disconnection




Figure 3.4: Example of domestic EV connection values during a weekday morning with $N_0 = 5,119$, $t_d = 07:00$ and $t_s = 07:00 \rightarrow 08:00$ or $t_s = 08:00 \rightarrow 09:00$ or $t_s = 09:00 \rightarrow 10:00$. FR must consider intra-hour connections (5-min) to define ΔN_i , while charging decisions use average hourly values to define $\Delta \hat{N}$.

time series to provide $\Delta \hat{N}, \Delta N^{in}, \Delta N^{out}$, used in the constraints of Section 3.2.4. Average hourly values are appropriate for these constraints as they focus on net energy usage.

FR is procured by ESOs to be provided over a given time period, here we schedule amounts on an hourly basis. An FR activating outage can occur at any point during the scheduling period and the amount of actual FR delivered from EVs will vary depending on plug-ins at that specific time. Thus the variation in intra-hour FR provision capacity from EVs must be considered when limiting the risk of FR under delivery.

The need to consider intra-hour EV connections for FR scheduling is demonstrated in Fig. 3.4 which shows examples of the hourly and 5 minute discretized timeseries for the domestic fleet. EVs are disconnecting throughout the hours of the morning and the system operator wants to ensure the FR scheduled from EVs for each hour period is deliverable with 99% security. Relying on the mean number of connected EVs to deliver FR would result in under delivery when there is an outage in the second half of the scheduling period (50% under-delivery). This risk to system dynamic security is unacceptable. Hence the 5-minute data is used when characterising ΔN for FR scheduling.

ΔN is the change in the number of connected EVs between the scheduling decision time (t_d) and during the scheduling period of interest (t_s). There are 260 weekdays and 105 weekend days in the 2017 data. This means that for a given t_d on a weekday, there are $12 \cdot 260 = 3,140$ different samples of the possible change in the number of EVs between t_d and during each t_s




Figure 3.5: Empirical distributions of ΔN_i for the domestic fleet on weekday, with $t_d = 07:00$ and $t_s = 07:00 \rightarrow 08:00$ or $t_s = 08:00 \rightarrow 09:00$. These are the distribution of possible changes in the number of EVs connected between 07:00 and any 5-minute period (when an outage could occur and activate FR provision) within the 2 FR scheduling hour time periods.

period from the 5-minute discretized timeseries. Assuming all these are equiprobable, when collated together these 3,140 (or 1,260 for weekends) values produce an empirical distribution of ΔN . Given that we are simplistically assuming here that the EV connection patterns do not significantly alter between years, then it is assumed that these historical empirical distributions are the true distributions of ΔN for use in scheduling FR in the future.

Two example distributions for $t_d = 07:00$ and with $t_s = 07:00 \rightarrow 08:00, 08:00 \rightarrow 09:00$ are plotted in Fig. 3.5. These show the distribution of possible changes in the number of EVs connected between 07:00 and any 5-minute period (when an outage could occur) within the 2 FR scheduling hour time periods. The 5-minute time-series from Fig. 3.4 contributes 12 data points to each histogram. The stages of this forecasting approach are summarised in Fig 3.6.

3.4.2 Tests for Ambiguity Set Selection

The mean and standard deviation are necessary for all forms of (3.21) described by Table 3.1, but sufficient for only the most conservative DRO form. If more distributional information of δ can be extracted from the distributions of its constituent ΔN_i , then the ambiguity set can be tightened, and (3.21) made less conservative.

The likelihood that the ΔN_i distributions conform to unimodal or Gaussian distributions




Figure 3.6: Flow diagram to summarise forecasting approach. Empirical distributions for future intra-hour EV connectivity are assumed to be well characterised by driving behaviour from a similar time of the week the previous year. t_d is the current time and t_s is the future period in which FR from EVs is being scheduled.

can be found by applying the Shapiro-Wilk and Hartigan’s dip test respectively. Because δ is the weighted sum of ΔN_i values, if they are established as Gaussian (and independent) with high likelihood, then so too is δ . Indeed, if the ΔN_i s can be shown to fit any distribution type with this summative property, then f_P^{-1} can be accurately defined as the inverse cumulative distribution function. The sums of unimodal distributions are likely to be unimodal [70], but not necessarily unimodal. Despite the lack of this mathematical guarantee, it is shown empirically in Section 3.5 that a unimodal assumption within (3.21) can improve optimality whilst maintaining a significant conservative margin. Although to be guaranteed violation probabilities less than those specified, the DRO form must be implemented.

Table 3.2 shows the results from applying the Shapiro-Wilk and Hartigans dip test to the 24 weekday and 24 weekend ΔN_i distribution where the scheduling period is the hour immediately following t_d (i.e. $t_s = t_d \rightarrow t_d + 1$). These are chosen because the operational scheduling model used here optimally schedules for the entire next 24hr period every hour. Thus, when it makes its scheduling decision, it is only the one for $t_s = t_d \rightarrow t_d + 1$ that impacts reliability post fault, because the other decisions are subsequently revised to account for updated wind and EV connectivity realisations.

Typically the hypothesis (unimodality or normality) is rejected for p-values below 0.05. The hypothesis is neither rejected nor accepted for values above that. Table 3.2 shows that the Gaussian hypothesis is not accepted, obvious from Fig. 3.5 due to the highly skewed distribution. On the other hand, the hypothesis that the distributions are unimodal is not rejected for any of

Table 3.2: Tests on ΔN_i for Normality and Unimodality

P-Value	< 0.05	0.05 : 0.50	0.50 : 0.95	> 0.95
Normality	19	12	12	5
Unimodality	0	0	7	41

the 48 distributions.

3.5 Case Studies

The uncertainty-aware model for V2G proposed in this chapter was incorporated into the SUC model introduced in Section 3.3. Case studies were run to identify how different EV fleet configurations and user risk-aversion impact the annual operational cost of the GB 2030 system. Unless otherwise stated, system thermal generation and storage mix was as listed in Table 3.3.

An advanced statistical wind model [20, 77] was used to characterise wind forecast uncertainty and produce a representative aggregate UK timeseries of wind realisations for the UK system in 2030. Full parametrization can be found in [77]. The average load factor is 35%. The time-series is normalised between 0 and 1 and then multiplied by the chosen GW of installed capacity, set at 40 GW unless otherwise stated. The aggregate UK solar generation time-series utilises the last year available for weather and system operator data [78]. Unless otherwise stated, 20 GW is installed with an average load factor of 11%. The UK 2020 demand time-series was used to represent passive demand, whilst EV fleet demand is detailed in Section 3.4. The annual passive demand ranges between 20:58 GW.

Generator and storage actions were optimally scheduled for 1 month of each season. A scenario tree that branches 7 times at the root node only was used to account for wind forecast uncertainty, which [20] showed to find the appropriate balance between tractability and optimality. Quantiles of 0.005, 0.1, 0.3, 0.5, 0.7, 0.9 and 0.995 were used.

Current GB frequency security standards were used: $f_0 = 50$ Hz, $|\Delta f_{max}| = 0.8$ Hz and $RoCoF_{max} = 1$ Hz/s. The FR time constants are $T_1 = 1$ s, $T_2 = 10$ s, while $c_{LS} = £30,000/\text{MWh}$. Unless otherwise stated, two EV fleets were present on the system. With 85,000, 10 kW ‘Domestic’ V2G chargers and 15,000, 20 kW ‘Work’ V2G chargers. The parameters used were

Table 3.3: Generation and Storage Characteristics

Generation	Nuclear	CCGT	OCGT
Number of Units	4	120	20
Rated Power (GW)	1.8	0.5	0.1
Min Stable Generation (GW)	1.60	0.25	0.05
No-Load Cost (£'000/h)	0.0	4.5	3.0
Marginal Cost (£/MWh)	10	47	200
Startup Cost (£'000)	NA	10	0
Startup Time (h)	NA	3	0
Min up Time (h)	NA	4	0
Inertia Constant (s)	5	4	4
Max Slow FR Capacity (GW)	0.00	0.05	0.04
Storage	Pumped	Battery 1	Battery 2
Capacity (GWh)	10	0.8	12
Dis/Charge Rate (GW)	2.6	0.4	3.0
Max Fast FR Capacity (GW)	0.0	0.8	0.0
Max Slow FR Capacity (GW)	0.5	0.0	0.0
Dis/Charge Efficiency	0.75	0.95	0.95

those derived in Section 3.4, linearly scaled to match the total number. Nadir security was specified at 99%.

An eight-core Intel Xeon 2.40GHz CPU with 64GB of RAM was used to run simulations. The optimisations were solved using XPRESS 8.12 linked to a C++ application via the BCL interface. The mixed-integer program gap was 0.1%.

3.5.1 Constraint Reliability

Fig. 3.7 demonstrates how the specified risk of $\bar{\mathbf{R}}_1^{EV}$ under delivery (ϵ) compares to actual deliverability when using (3.21) to schedule FR from EVs during annual SUC simulations. This varies depending on the ambiguity set assumptions (Gaussian, unimodal or DRO) and the assumed true distribution of ΔN_i (Gaussian or empirical). Hourly nadir security (HNS) is the metric used to evaluate $\bar{\mathbf{R}}_1^{EV}$ deliverability. HNS for a specific hour is found by sampling the ΔN_i distribution for each fleet. When added to the number of currently connected EVs, the actual FR deliverable ($\mathbf{R}_1^{EV}$) if an outage occurred at a random time over the scheduling period can be calculated. This process is repeated 100,000 times within each hour. The HNS is the cumulative ratio of $\bar{\mathbf{R}}_1^{EV} > \mathbf{R}_1^{EV}$. Fig. 3.7 plots the range of HNS for different constraints over




Figure 3.7: The inter-quartile range, median and max/min of hourly nadir security levels under different ambiguity set assumptions (Gaussian, Unimodal, DRO) on δ when scheduling FR via (3.21). The scheduled amount was compared to true EV connectivity, found by sampling Gaussian or Empirical distributions of ΔN_i .

the simulation period. It only shows the t_s in the hour immediately following t_d . Due to the rolling planning approach of the SUC, the $\bar{\mathbf{R}}_1^{EV}$ values for other t_s are revised before the system would experience an outage.

When δ_n is assumed Gaussian and the $\mathbf{R}_1^{EV}$ is found from sampling Gaussian distributed ΔN_i , the HNS exactly equals the specified security level of 99% when constraint (3.21) is tight. The constraint is occasionally not tight during periods of high net demand when inertia and slow FR from thermal plants are sufficient to meet frequency security needs. During these times the HNS takes values above 99%. However, when the true empirical ΔN_i distributions are sampled instead, 50% of hours have a HNS less than the specified 99%. In the worst period, using the Gaussian (3.21) constraint would result in only 96% of outages being contained securely. This unreliability is in line with the analysis of Table 3.2, which showed none of the relevant empirical ΔN_i distributions are Gaussian. Thus using the Gaussian form of (3.21) is over optimistic and risks system security, so is not considered further.

On the other hand, Fig. 3.7 shows that applying (3.21) assuming a unimodal or DRO δ_n distribution results in conservative HNS values for both the Gaussian and empirical distribution sampling. For Unimodal and DRO the worst periods have a HNS of 99.7% 99.9% respectively, with the majority of hours delivering more response than is scheduled with close to 100% probability. Again, this is in line with the analysis of Table 3.2, which showed a high likelihood that ΔN s are unimodal.

Table 3.4: Constraint Value Comparison for Specified 99% Frequency Security and 100,000 Chargers

Constraint	Worst Case Nadir Security	Charger Value (£/yr)	Charger CO ₂ Savings (tons/yr)
No V2G FR	100.0	0	0
Deterministic	50.0	8,040	44.5
Unimodal (3.26)	100.0	5,930	34.1
Unimodal (3.21)	99.7	6,330	37.4
DRO (3.26)	100.0	5,200	31.7
DRO (3.21)	99.9	5,760	33.6

3.5.2 Value of Response from V2G and Comparison to State of the Art

A scheduling method’s success criteria is to schedule as much FR from EVs as possible to improve operational efficiency whilst ensuring that it is deliverable at the specified certainty level or higher. Table 3.4 compares the value of FR from EVs when using our proposed ‘joint’ method (3.21) to the value when using three simple methods that are the current state of the art. It shows it to be more secure than the deterministic approach whilst less conservative than the ‘No V2G’ and ‘individual’ approach.

The ‘No V2G FR’ row in Table 3.4 allows no response from distributed resources. Zero FR is deliverable 100% of the time but offers no operational cost savings. The value of FR from EVs using the other methods is defined as the reduction in system annual operational costs compared to ‘No V2G FR’. The value creation mechanisms are explored in Section 3.5.4.

A system operator could schedule FR using a simple deterministic approach, where it is assumed that the mean forecast number of EV connections will occur. This can be done by making ‘ $\sigma_i = 0 \forall i$ ’ in (3.21), resulting in a linear constraint. By discounting uncertainty in this way, larger amounts of $\bar{\mathbf{R}}_1^{EV}$ are scheduled thus annual system costs are significantly reduced by £8,040/charger. However, because the mean is used, half of the time fewer EVs are available to provide response than is scheduled for. Thus this cost reduction comes at an unacceptable loss of nadir security, with average and worst case HNS being only 50%.

EV connection uncertainty can be considered simplistically via the ‘individual’ method (3.26). As demonstrated in (3.27), our more advanced ‘joint’ approach (3.21) will always schedule more

Table 3.5: Dynamic Simulation Parameters

	$\bar{R}_1^{EV}$ (GW)	$R_{1,1st}^{EV}$ (GW)	R_1^{ND} (GW)	R_1^G (GW)	H (GWs)	PL_{max} (GW)
Unimodal (3.21)	0.23	0.33	0.36	2.27	96.25	1.73
Unimodal (3.26)	0.22	0.42	0.40	2.36	101.49	1.80
Determ	0.28	0.22	0.40	2.23	96.24	1.78

response than (3.26), thus system operation is improved and a charger's value is increased by 6.9% and 10.9% for unimodal and DRO ambiguity set assumptions respectively. Crucially this comes with worst case HNS being maintained well above the specified 99% security requirement.

For 99% of outages the EVs deliver more FR than the first 99th percentile of R_1^{EV} ($R_{1,1st}^{EV}$). Specifying $\epsilon = 0.01$ is equivalent to requiring that the frequency nadir remains higher than 49.2Hz when $R_{1,1st}^{EV}$ GW is delivered. Table 3.5 lists the frequency services scheduled for a typical weekday 09:00 period under four scheduling methods. In whilst Fig. 3.8 a time-domain simulation in MATLAB was used to plot their dynamic frequency evolution when $R_{1,1st}^{EV}$ is delivered instead of the scheduled $\bar{R}_1^{EV}$.

When the scheduled amount of $\bar{R}_1^{EV}$ is delivered, the constraint (3.10) is tight so the nadir is exactly 49.2 Hz. A scheduling method with perfect EV connection forecasts would emulate this when $R_{1,1st}^{EV}$ GW of FR is delivered. Fig 3.8 demonstrates that our 'joint' method is closer to this ideal than the 'individual' approach, that significantly under schedules FR from EVs. This is undesirable as being overly conservative inflates operational costs. The conservativeness of the 'joint' approach is attributed to its robustness in covering the entire unimodal ambiguity set. On the other hand, the deterministic method is overly optimistic and over schedules FR from EVs. Thus 50% of the time the frequency breaches the nadir limit, jeopardising system security and demonstrating the importance of applying a risk-aware scheduling method like the ones derived in this thesis.




Figure 3.8: Comparison of the frequency evolution post PL_{max} loss given the delivery of the scheduled amount of FR $\bar{R}_1^{EV}$ or of the 1st percentile of the true FR distribution R_1^{EV} for the system conditions shown in Table 3.5. The ‘scheduled’ plots are very similar so for clarity only one is plotted.

Table 3.6: Increase of FR from EVs Using (3.21) over (3.26)

Fleets		Unimod		DRO	
Work	Domestic	Value (£)	FR (GW)	Value	FR
1	0	0.0%	0.0%	0.0%	0.0%
1	1	6.9%	5.4%	10.9%	8.5%
2	0	28.0%	22.6%	49.2%	40.3%
3	0	38.6%	32.6%	70.1%	61.0%

3.5.3 Constraint Performance with Varied Fleet Types

A strength of our proposed formulation is its wide applicability to systems with any number and type of fleet so long as their σ_i and μ_i are known. Accordingly, Table 3.6 shows the improvement, in value and amount of scheduled FR, of the ‘joint’ approach over the ‘individual’ approach for a range of fleet setups.

According to (3.27), the improvement of our ‘joint’ method over the state of the art increases with the ratio of $\sum_{i \in I} \mathbf{g}_i \delta_i : \sqrt{\sum_{i \in I} (\mathbf{g}_i \delta_i)^2}$. This ratio will increase when $\mathbf{g}_i \delta_i$ are of a similar size, explaining the increased improvement of our method for the system with two work fleets of 15,000 EV, as opposed to the case with one ‘Domestic’ fleet with 85,000 EVs and one ‘Work’ fleet. Our method will also increasingly outperform the current state of the art as the number of fleets increases, demonstrated by comparing the cases with two and three ‘Work’ fleets.

Fig. 3.9 demonstrates the increased conservativeness of the current state of the art over our




Figure 3.9: Total scheduled FR from three 15,000 EV ‘work’ fleets over the same two-day period, using the ‘joint’ (3.21) and ‘individual’ (3.26) approaches. Nighttime is shaded purple.

proposed method by plotting the total $\bar{R}_1^{EV}$ over the same two-day period using both methods for a case with three ‘Work’ fleets with 15,000 EV each, under DRO ambiguity assumptions. During the highly uncertain period in the morning and evening when EVs are arriving and departing at work, neither method can reliably schedule any FR from the fleets. However, at other times the ‘joint’ method is able to schedule more FR than the individual, with up to 0.38 GW more in the period shown. More FR from EVs translates into reduced operational costs, thus increased charger value.

3.5.4 Investigating time-varying FR contribution from EVs

Fig. 3.10 shows how FR provision from aggregated EV fleets creates system value by significantly reducing wind and solar curtailment during low net-demand periods. It compares the operating conditions of the system with and without FR from EVs enabled over a typical two-day period. The unimodal constraint is used. Net demand is plotted, this is equal to the total wind and solar energy available subtracted from demand, so is the same for both systems.

The number of committed CCGTs for both systems is also plotted. The provision of FR from EVs significantly reduces the need for ancillary services from the thermal plants. Thus the system with FR from EV enabled is able to decommit more CCGTs. One way that this creates value is by increasing the load factor of the remaining CCGTs, improving their efficiency. However, the main value creation mechanism comes from reducing the committed minimum




Figure 3.10: Two day example comparing system operation with and without FR from V2G. Nighttime is shaded purple. The top figure plots the two system’s CCGT commitment number as well as the difference in wind and solar curtailment between the two systems. The bottom figure plots the operation of the Domestic fleet virtual battery for the system with V2G FR enabled.

stable generation. This significantly reduces the curtailment of renewable energy, creating large value. To demonstrate this the additional wind and solar curtailment for the system with no EV FR are plotted. The crucial difference is in how much more of this available zero marginal cost and emissions-free renewable power the system with FR from V2G is able to integrate.

During the first day, net demand is above 10 GW. The inertia from CCGTs committed to serve this load in combination with system FR is enough to fully secure frequency so no

wind shed occurs in either system. However, during the second night demand drops and wind energy increases, resulting in low (and even negative) net-demand for the second day. In the system without FR from V2G, the inertia from thermal plants needed for energy provision alone is insufficient to secure the nadir (3.10). Consequently CCGTs must be committed for their inertia and FR alone. When the sum of the minimum stable generation of online thermal plants is larger than net demand, renewable power must be curtailed to respect the power balance constraint (3.30). The clear correlation between over commitment of CCGTs and wind curtailment is shown in Fig. 3.10, when during the second night around 24 more CCGTs are online for the system without V2G FR. Thus the sum of their 250MW individual minimum stable generation results in roughly 6GW wind curtailment.

Fig. 3.10 also plots the operation and FR provision of the domestic fleet during the same two-day period. During the second night the EVs are charged at a constant rate. The need for full charge by morning synergises with the typically increased FR value overnight caused by lower net-demand. FR is a net power injection thus a charging EV can provide more response via demand alleviation. The approximately 0.6 GW of FR from the fleet during the second night replaces the inertia from approximately 20 additional CCGTs, facilitating up to 11GW of wind integration. Cumulatively over the year the net difference in wind and solar curtailment amounts to 7.55 TWh and 0.40 TWh respectively. In other words, the highly effective FR from V2G connected EVs enables frequency secure operation at very low inertia levels, resulting in 8 TWh less power generated by burning fossil fuels (80 MWh per charger). This accounts for the majority of cost and emission savings in Table 3.4.

3.5.5 Value's Sensitivity to User Risk aversion and Forecast Uncertainty

A useful feature of (3.21) is that it directly translates a user's risk aversion level (ϵ) into a scheduled amount of FR. Fig. 3.11 illustrates that higher risk aversion leads to less response allowed from EVs and thus lowers their value. The tightened ambiguity set from assuming δ_n is unimodal results in a larger $\frac{1}{f_P^{-1}(1-\epsilon)}$ constant in (3.21), and thus more $\bar{R}_1^{EV}$ can be scheduled




Figure 3.11: Sensitivity of charger value to FR under-delivery risk aversion.




Figure 3.12: Sensitivity of charger value to increased standard deviation (std) of ΔN_i .

than the DRO case for the same ϵ . This explains the increased operation cost savings shown in both Table 3.4 and Fig. 3.11.

A large σ_i of ΔN_i implies that the number of connected EVs in the scheduling period of interest is highly variable. This decreases the amount of schedulable FR from the chargers because a high deliverability probability necessitates covering the edge cases where far fewer EVs than the mean are connected at the time of an outage. A reduced std has the opposite effect, increasing the amount of schedulable FR, this is reflected in Fig. 3.12 which shows how the V2G's value varies in relation to a multiplier applied to the stds in equation (3.21). As the stds tend to zero, the cost saving tends towards the deterministic case level of £8,000/yr, with the difference between DRO and unimodal ambiguity set assumptions diminishing. It is




Figure 3.13: The main value creation mechanism of FR from V2G chargers is the ability to facilitate renewable generation integration by displacing inertia and FR from thermal plants. The value sensitivity to installed system wind capacity and 10GW ('S10') and 20GW ('S20') of solar capacity is plotted here.

interesting to observe that even with very large stds the EVs still provide substantial value of approximately £2,000/yr, because their connectivity variability overnight is very low, so FR provision during this valuable period is mostly uninterrupted.

The relationship between a small std and increased V2G value has two main implications for real life application of (3.21). Firstly, more accurate EV connectivity forecasting methods are directly incentivised. Secondly, closer to real time scheduling of FR is desirable, as it reduces uncertainty in forecasting.

3.5.6 Value's Sensitivity to System Characteristics

The value of response from V2G is highly dependant on its ability to facilitate higher renewable integration by displacing the inertia and FR from thermal plants. Consequently, Fig. 3.13 shows that increased renewable generation increases V2G value. With 60 GW of wind and 20 GW solar, one V2G charger has an annual system value of £9,400/yr. This comes from the increased frequency and magnitude of low-inertia periods, where the FR from the 100,000 chargers facilitate a cumulative 14.3 TWh of renewable power integration.

The cost of frequency security (CFS) is the operational cost increase incurred by applying the nadir and RoCoF constraint. We now analyse the CFS's dependence on the amount of

storage on the system. As discussed in Section 3.5.4 this cost increase comes from the need to run CCGTs to provide inertia and FR to satisfy the nadir and RoCoF constraints. Running CCGTs burns fuel at a cost, and often their cumulative minimum stable generation forces the curtailment of zero marginal cost and emission-free renewable energy. In systems with high renewable penetrations this cost increase is very significant. Fig. 3.14 shows that it reaches £2.26bn/yr (22% of total operational cost) for our standard test system of 20GW solar and 40GW wind, when neither V2G or batteries can provide FR.

FR from inverter-based resources reduces the required inertia and FR from thermal plants and thus increases renewable integration and reduces the CFS. A core strength of our proposed formulation is that it allows the abundant distributed resources that will be present in future systems (like V2G-connected EVs) to compete directly with grid batteries to provide this FR, whilst maintaining a user specified guarantee on system dynamic security. Operational costs for batteries and V2G are assumed zero so do not contribute to CFS.

The batteries used here are linearly scaled version of Battery 1 in Table 3.3 so have 2hr tank sizes. Fig. 3.14 shows that the first 2.25 GW of V2G capacity is a third as effective at reducing the CFS as battery capacity. For example, to reduce FSC to £1.5bn/yr requires 0.35 GW of batteries or 0.95 GW of V2G (66,500 ‘Domestic’ and 14,250 ‘Work’ chargers). Reduction to £0.75bn/yr requires 0.75 GW of batteries or 2.30 GW of V2G (161,000 ‘Domestic’ and 34,500 ‘Work’ chargers). Above 2.50 GW of V2G capacity, FR is abundant enough that renewable shedding only occurs during the highly uncertain morning period. During these periods the σ_i values are large enough that $\bar{\mathbf{R}}_1^{EV}$ in (3.21) is forced to be very small or zero, as occurs during the two morning periods in Fig. 3.9. Increased V2G capacity alleviates this slowly, explaining the plateau in CFS reduction from V2G capacity. Whereas battery storage has no uncertainty so the same effect is not observed. However, the marginal value of storage does saturate above 1.2 GW, at which point the minimum inertia for the RoCoF constraint (to which FR does not contribute) dominates the CFS.

The lower value of FR from V2Gs compared to the same capacity of batteries is primarily attributable to a V2G charger only offering FR when an EV is connected. From the fleet parameters derived in Section 3.4, the average charger has an EV plugged in 42% and 26% of




Figure 3.14: Operational cost increase from applying the frequency security constraints to systems with different storage penetrations. The difference between battery and V2G is primarily because the chargers on average only have an EV connected (thus FR capability) $\approx 40\%$ of the time. Adjusting for this gives the ‘Average V2G’ capacity plots.

Table 3.7: Sensitivity of Normalised V2G Value to FR Provision Delay

Delay (s)	0	0.2	0.4	0.6	0.8	1.0
Normalised Value	1.00	0.97	0.93	0.89	0.85	0.80

the time for ‘Domestic’ and ‘Work’ respectively. This is adjusted for within Fig. 3.14 with the average annual V2G capacity available plotted against value created for both the deterministic and unimodal constraints. The small difference between deterministic and the battery capacity is attributable to the EVs time of connection and charge requirements to meet energy needs. The difference between the unimodal and deterministic plots is due to uncertainty, revealing this to be the second most impactful derating factor. Uncertainty has low impact below 1 GW of average capacity, but above this it prohibits frequency security cost reduction again due to the highly uncertain morning periods. This suggests that the addition of a fleet with low morning plugin uncertainty, or a small amount of grid batteries would be valuable at high EV penetrations.

3.5.7 Impact of Delays on Value of EV’s Response

All other sections of this chapter assume FR from EVs begin ramping up the instant of PL_{max} disconnection. In reality, the triggering mechanism may involve delays caused by communication or frequency measurement requirements. The additional term in (3.11) allows the impact of

delays on the value of FR from V2G connected EVs to be analysed within the SUC. The results are shown in Table 3.7. The normalised value decrease is shown for the unimodal case. The value impact for the DRO case is very similar so omitted here for brevity. Delays decrease the efficacy of response in containing frequency drop and thus reduce its value. Although, even when the final FR delivery time is doubled with a 1s delay, the value of FR remains substantial at £5,060/EV/yr for the unimodal case, or a decrease of 20% with respect to the case with no delay.

3.6 Conclusion and Future Work

This chapter proposes a mathematical framework to schedule frequency response from aggregated V2G chargers under uncertainty in EV plug-in times. A distributionally-robust chance constrained formulation for the frequency-security limits is introduced, compatible with any probabilistic forecast for EV connections and providing a convex feasible region. Case studies demonstrate that the proposed scheduling methodology facilitates large cost savings ($\approx$ £6,000/charger) in the future Great Britain low inertia system due to displaced inertia and FR requirements from thermal plants. Crucially, this value is obtained with mathematical guarantees on system frequency security.

In future, a model that accounts for charger location should be developed. Given the significant economic benefits that V2G could provide, it will be key to coordinate V2G power injections to ensure that distribution-side network constraints are respected. Secondly, the design of a market for FR that allows aggregator participation should be investigated. The tool developed here allows market clearing under security guarantees, but the mechanism of specifying, communicating and monitoring aggregator uncertainty characteristics needs development.

Chapter 4

Probabilistic Scheduling of UFLS to Assist With Single Plant Outages in Low Inertia Systems

As shown in the previous chapter, frequency ancillary service requirements are driven by the square of the largest generator loss. This chapter designs a new frequency service based on the intuition that for systems with only a few large possible contingencies, it may be economical to risk a small amount of load shedding on the rare occasions that an outage occurs, in order to decrease day-to-day operational costs significantly. Case studies demonstrate that for systems characterised by high RES penetrations and a few large individual losses (like the UK), the low probability of an outage makes this service valuable.

This chapter also presents a method to allocate the desired load to be shed between candidate feeders. This is an increasingly challenging task due to the power-flow uncertainty introduced by distributed generation. The compatibility of this method with legacy and emerging under frequency load shedding (UFLS) frameworks is expounded. This chapter is based on the journal papers [1] and [3].

4.1 Introduction

Grid frequency needs to be kept within boundaries to prevent system damage and emergency demand disconnection. The rate of change of frequency (RoCoF) is determined by the instantaneous generation-demand imbalance. Post generator outage, the frequency's evolution is dictated by the available frequency services (FS) including inertia and FR.

Synchronous thermal generators provide inertial response from the kinetic energy stored in their rotating turbines and FR from governor droop controls. Decarbonisation replaces synchronous thermal generators with renewable energy sources. This means that frequency security is no longer guaranteed by meeting energy demand alone because currently, converter based generation (like wind) is often not able or not incentivised to provide any FS. Thus explicit frequency security constraints must be included in scheduling to assure system stability. This can result in significant wind curtailment due to the minimum generation from thermal plants that are needed to provide FS [25], negatively impacting costs and emissions. The increase in operational cost is referred to here as a system's frequency security cost (FSC).

This motivates the investigation of other possible tools at the disposal of system operators to contain the frequency more cost effectively. Under frequency load shedding (UFLS) results in a step decrease in demand. It is used traditionally as a last resort containment measure after very large outages (non-credible/unsecured events) [79], and its potential role to support frequency control during normal operation has been overlooked.

The increased deployment of renewable energy sources in electricity grids has happened not only on the transmission level, but also on the distribution side: more than 30% of the solar photovoltaic capacity installed is connected to the distribution system [80]. Distributed generation (DG) resources make the net load flow through feeders more uncertain ahead of time. This presents a challenge to system operators when planning and implementing UFLS schemes, as legacy methods to translate a desired load shed amount into a set of relays to disconnect assume determinism, thus no longer guaranteeing security.

The majority of UFLS literature focuses on how to contain unsecured events with the least load shed [81]. This necessitates a fast nadir prediction following the event, from which the

optimum load to be shed can be calculated online. If the outage size is known, then the single machine swing equation can be used [82]. When outage size is unknown, it can be estimated from initial RoCoF measurements [83]. Alternatively, the future frequency evolution can be predicted directly from frequency measurements post disturbance [84]. Another option is to use artificial neural networks to infer the outage size and subsequent optimum load shed amount [85].

All these methods try to minimize the UFLS used post-event, as a containment measure in a system with set FS and commitment decisions. Accurate nadir predictions will implicitly consider system inertia and FR, but offer no framework to compare their cost effectiveness in securing a given outage against UFLS.

Furthermore, because DG make the net-power flow through feeders stochastic, it will no longer be a simple task to map a desired amount of demand to be shed into a given choice of feeders. Shedding less load than desired will hinder the efficacy of all UFLS schemes and jeopardise system security. Indeed, in extreme cases, DG can make some feeders power exporters at times, in which case their disconnection will exacerbate the problem. Solving this second stage relay-allocation problem is vital to the success of any UFLS scheme, and currently no robust method to deal with the uncertainty exists.

A few limited heuristic methods represent the current state of the art for the feeder allocation problem [86, 87, 88, 89, 90, 91]. References [86], [87] and [88] analyse the effects of disconnecting different feeders on the system's dynamic performance and voltage stability. The prioritisation of feeders to be disconnected is based on certain voltage and frequency stability indexes, linked to active power changes. The online computation of indexes requires real-time measurements, which make these methods highly dependent on measurement and communication devices to perform the UFLS allocation among feeders. The authors of [89] and [90] explore the prioritisation of feeders through the concept of generation-to-consumption ratio, i.e., the total amount of load disconnected compared to the net-power flow disconnected. This metric is used to reduce disruption to customers, as feeders are heuristically classified based on the values of this ratio, prioritising the disconnection of the most load-dense circuits. This classification leads to some of the consumers never being considered for the UFLS scheme, and the selection of feeders is

still dependent on real-time power flow measurements and DG output estimation, therefore no uncertainty is modelled. Finally, the authors in [91] compare current UFLS-allocation methods used in France, based on simple historical measurements of consumption through each feeder, with their proposed method based on clustering feeders that could jointly contribute to minimise deviation of actual load shed from the desired volume of load shed. The clustering is computed using also a time-series of historical data.

In the frequency constrained scheduling literature, the main mathematical challenge lies in containing the frequency nadir, as it necessitates incorporating the dynamic equations dictating post-event frequency evolution into the algebraic optimisation problem [32]. A highly non-linear nadir constraint considering thermal-plant response only, is derived in [92] and implemented in the unit-commitment via piece wise linear functions. Reference [93] augments the formulation of [92] to consider response from batteries by suggesting a control scheme that effectively reduces the size of the loss. Reference [94] derives a nadir constraint incorporating response from grid forming inverters which is used in dynamic simulations to determine the set of secure post-outage plant combinations. Appropriate decision variable are then bounded within the unit commitment to insure commitment lies within this set.

References [28, 25, 29, 30] approximate the FR provision from droop control as a linear ramp to simplify the analytical formulation. These works propose: a linear constraint on FR from turbines in a system with fixed inertia [29]; a linear approximation of the true constraint for systems with variable inertia and single speed FR [25]; and a reformulation of the nadir constraint as a SOC, resulting in a SOC program that optimises over inertia, a reduces largest loss, fast and slow FR [28]. All these papers contain the nadir in order to prevent the triggering of UFLS.

The only framework that considers UFLS as an option to contain secured events is [30], whose nadir constraint bounds the product of the inertia and FR from thermal plants to be larger than a constant. The outage size is assumed constant, and the UFLS amount is discretized into blocks, with one nadir constraint per block. The optimizer can then choose the optimal UFLS level to schedule using binary variables. A piecewise linear approximation is used to apply this constraint within a mixed-integer linear programme framework.

The first of two major contributions of this chapter is to derive a nadir constraint that allows moderate amounts of UFLS to assist in containing the nadir. The expected cost of UFLS is found probabilistically, as the product of the probability of a largest loss occurring in a given hour and the cost of the load-shed it would initiate. This expected hourly cost facilitates the direct comparison of UFLS to other traditional FS in order to secure frequency most cost effectively. Breaking the paradigm of UFLS as a last resort only and demonstrating that when considered alongside other FS, UFLS is a valuable resource to contain outages of the single largest power infeed (credible/secured events) in low inertia systems.

The nadir constraint proposed here is significantly more advanced than the linear one of [30], as it can consider a dynamically reduced largest loss, and multiple speed FR. Both of which [28] has shown to be extremely effective frequency containment measures in low inertia systems, thus vital to consider when accurately evaluating UFLS's value in reducing system FSC. We demonstrate that the proposed nadir constraint can be closely approximated by a single second order cone at each UFLS level. Resulting in a tractable mixed-integer second order cone programme (MISOCP) formulation that fully exploits the capabilities of modern convex optimisation solvers.

None of the previous work, to the best of our knowledge, has addressed the allocation of load shedding while accounting for uncertainty, in a formulation that mathematically guarantees the disconnection of the required load. The second core contribution of this chapter is to propose a method that translates a desired load to be shed into a set of relays that guarantee delivery under DG uncertainty, with a user-specified level of risk for under-delivery in extreme scenarios.

A strength of our proposed method is that it is useful in combination with all types of UFLS schemes, although its application varies depending on when the decision on which relays will be disconnected is made (offline before the outage [95] or online during the transient period following an outage [82, 83, 84]). For the first case [95], where UFLS relays are triggered in a decentralised manner through local frequency measurements, the net-load through the feeders will always be uncertain at the time of outage. However, accurate forecasts can produce information on the expected DG output value, which our method utilises alongside the forecast error to inform relay selections.

For the second case [82, 83, 84], corresponding to adaptive UFLS techniques based on Wide-Area Measurement Systems (WAMS), a powerful monitoring and communications network is necessary. These WAMS-based adaptive UFLS methods use activation signals that are sent to relays in the instants following a contingency. Since adaptive UFLS allocation is updated in real-time, it is unaffected by the uncertainty introduced by DG, as power flows through feeders could be available from the live measurements in WAMS. However, WAMS-based adaptive UFLS schemes rely completely on a communication network, which might fail or be subject to attacks [96]. Given that UFLS is a containment measure that avoids system collapse, robust methods to allocate feeders contributing to UFLS are needed, rather than purely relying on real-time updates. The method proposed in this chapter can act as a back up to WAMS-based adaptive UFLS techniques, ensuring frequency security if the adaptive UFLS scheme cannot provide an acceptable performance due to unavailability of the communication network.

The key contributions of this work are:

1. A comprehensive framework to co-optimize UFLS along with fast and slow FR, inertia and reduced largest loss, to secure frequency in the most cost effective manner.
2. A novel, least conservative convexification of the frequency nadir constraint using SOCs. The computational tractability of this formulation is demonstrated.
3. A convex-optimisation framework for the UFLS feeder-allocation problem under DG uncertainty is proposed, enabling to meet a pre-defined level of risk in obtaining the minimum load to be shed, while accounting for correlation in net-load across the network.
4. The application of the developed model to a representative GB 2030 system. The sensitivity of UFLS value to: UFLS amount and cost; availability of alternative FS; and wind penetration is investigated.

This chapter is organised as follows: Section 4.2 details the formulation of the convex frequency security constraints and section 4.4 confirms their veracity with dynamic simulations and case studies. Case studies exploring the value of UFLS as a service to contain frequency are presented in Section 4.5, whilst section 4.6 gives the conclusion.

4.2 Frequency Security constraints

This chapter applies novel frequency security constraints to the same multi-stage unit commitment problem described in Section 3.3. An additional term is added to the cost function to account for the expected cost of load shedding:

$$\sum_{n \in \mathcal{N}} \phi_n \sum_{t \in \mathcal{T}} \left(\Delta \tau (c^{LS} \mathbf{p}^{LS}_{t,n}) + \mathbf{C}^U_{t,n} + \sum_{\hat{g} \in \hat{G}} \mathbf{C}_{\hat{g}}(\mathbf{P}_{\hat{g},t,n}) \right) \quad (4.1)$$

Each node in the scenario tree has a complete set of system decision variables, so the index t and n are dropped in this section to improve equation clarity. After a generator loss, a system's frequency time-evolution is described by the swing equation [74]:

$$\frac{2H}{f_0} \frac{d\Delta \mathbf{f}^u}{dt} = \mathbf{R}_1(t) + \mathbf{R}_2(t) - \mathbf{P}\mathbf{L}_{max} \quad (4.2)$$

Equation (4.2) neglects load damping because the level will be significantly reduced in future power-electronic dominated systems [29]. As in [29, 25, 30, 28], frequency response is modelled as a linear ramp, fully specified by its magnitude and delivery time. Detailed dynamic simulations in Section III of [31] show that frequency droop controls can be conservatively and accurately approximated by a ramp. More detailed FR dynamic models would impede the algebraic derivation of convex constraints from the swing equation, as no closed form frequency security conditions could be deduced. EVs are not considered in this chapter so all formulations are deterministic.

$$\mathbf{R}_j(t) = \begin{cases} \frac{\mathbf{R}_j}{T_j} \cdot t & t \leq T_j \\ \mathbf{R}_j & t > T_j \end{cases} \quad j = 1, 2 \quad (4.3)$$

In this chapter $T_1 < T_2$. The faster speed can correspond to response from grid batteries, while the slower speed is used to model generator dynamics.

As in Section 3.2.1, the maximum RoCoF occurs the instant of generator disconnection, before any FR has been delivered. The maximum RoCoF must be constrained to prevent distributed generation disconnection (which would further exacerbate the imbalance), derived

directly from (4.2):

$$\frac{2|RoCoF_{max}|}{f_0} \cdot \mathbf{H} \geq \mathbf{PL}_{max} \quad (4.4)$$

4.2.1 Nadir Constraint with UFLS

As shown in (4.2) frequency decline occurs when demand is larger than generation. From this it follows that the lowest frequency (nadir) will occur at the instant when the power deficit is made zero. Scheduling at least enough FR to cover $\mathbf{PL}_{max}$ insures this will happen before T_2 :

$$\mathbf{PL}_{max} < \mathbf{R}_1 + \mathbf{R}_2 \quad (4.5)$$

FR reduces the deficit via power injection from storage or generators. UFLS reduces the deficit via a step demand reduction, initiated when low frequency relays measure the local frequency to have dropped below a predefined threshold [24]. As such, in real terms scheduling UFLS simply means scheduling a reduced combination of other FS that allows the frequency to drop below the threshold frequency. The magnitude of demand disconnected (P^{UFLS}) is equal to the sum of the load through the relays chosen to disconnect upon trigger level breach. This value is discrete by nature. The formulation here allows the system operator to schedule a range of block magnitudes, achievable when relay settings can be updated hourly. The method to choose between varying UFLS amounts is detailed in section 4.2.3, whilst section 4.3 describes my proposed method to solve the second stage problem of feeder allocation under DG uncertainty.

The nadir will occur at the UFLS trigger level if the sum of FR delivered by that time and the load shed amount are larger than the outage:

$$P^{UFLS} \geq \mathbf{PL}_{max} - \mathbf{R}_1 - \frac{\mathbf{R}_2}{T_2} \cdot \hat{t} \quad (4.6)$$

Where $\hat{t}$ is the time after outage that $\Delta \mathbf{f}^u(t) = -\Delta f_{trig}$. Equation (4.6) assumes that $\hat{t}$ occurs within the time interval $T_1 : T_2$. A valid assumption because equation (4.5) insures the nadir will always occur before T_2 . The extremely low inertia required for a system to reach $-\Delta f_{trig} \approx -0.8\text{Hz}$ before $T_1 \approx 1\text{s}$ would violate (4.4) for realistic power system parameters.

An expression for $\hat{t}$ can be found from (4.2). First, integrate between 0 and $\hat{t}$ to remove the time derivative:

$$\frac{2H}{f_0} \cdot \Delta f^u(t) = \frac{R_2}{2T_2} \cdot t^2 - (PL_{max} - R_1) \cdot t - \frac{R_1 T_1}{2} \quad (4.7)$$

By rearranging and completing the square, t for a given frequency can be found. Thus the time of $\Delta f^u(t) = -\Delta f_{trig}$ is given by:

$$\hat{t} = \frac{T_2}{R_2} \cdot \left(\sqrt{\frac{-4H\Delta f_{trig}R_2 + f_0R_1R_2T_1}{f_0T_2} + (PL_{max} - R_1)^2} + (PL_{max} - R_1) \right) \quad (4.8)$$

By substituting (4.8) into (4.6), the nadir constraint can be formulated:

$$\underbrace{\left(\frac{H}{f_0} - \frac{R_1 \cdot T_1}{4\Delta f_{trig}} \right)}_{=z} \underbrace{\frac{R_2}{T_2}}_{=x} \geq \underbrace{\left(\frac{PL_{max} - R_1}{2\sqrt{\Delta f_{trig}}} \right)^2}_{=y} - \underbrace{\left(\frac{P_{UFLS}}{2\sqrt{\Delta f_{trig}}} \right)^2}_{=u} \quad (4.10)$$

Where $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are linear expressions of decision variables and u is a constant. Finally, it is useful to rearrange (4.10) into (4.12) with the following substitutions:

$$\mathbf{Z} = \frac{1}{\sqrt{2}}(\mathbf{x} + \mathbf{z}), \quad \mathbf{X} = \frac{1}{\sqrt{2}}(\mathbf{x} - \mathbf{z}), \quad \mathbf{Y} = \mathbf{y} \quad (4.11)$$

$$\begin{aligned} \left(\underbrace{\frac{1}{\sqrt{2}} \left(\frac{R_2}{2T_2} + \frac{H}{f_0} - \frac{R_1 T_1}{4\Delta f_{trig}} \right)}_{=Z} \right)^2 &\geq \left(\underbrace{\frac{1}{\sqrt{2}} \cdot \left(\frac{R_2}{2T_2} - \frac{H}{f_0} + \frac{R_1 T_1}{4\Delta f_{trig}} \right)}_{=X} \right)^2 \\ &\quad + \underbrace{\left(\frac{PL_{max} - R_1}{2 \cdot \sqrt{\Delta f_{trig}}} \right)^2}_{=Y} - \underbrace{\left(\frac{P_{UFLS}}{2 \cdot \sqrt{\Delta f_{trig}}} \right)^2}_{=u} \end{aligned} \quad (4.12)$$

4.2.2 Nadir Constraint Convexification

The system frequency will not deviate below $-\Delta f_{trig}$ when constraint (4.12) is respected. However, it is non-convex so has limited uses in scheduling and market problems due to low tractability and lack of strong duality properties. The convexification proposed here, stems from the observation that (4.12) is a convex SOC of the classic form $\mathbf{Z}^2 \geq \mathbf{X}^2 + \mathbf{Y}^2$ when there is no

UFLS ($u = 0$). Fig. 4.1 shows the constraint over a range of P^{UFLS} values. A 2D plot of the 3D constraint is appropriate because the constraint is rotationally symmetric around the Z-axis in the XYZ space.




Figure 4.1: Feasible space defined by the nadir constraint (4.12) for various scheduled amounts of UFLS. $\mathbf{Y}$ can be considered as an adjusted outage size and $\mathbf{Z}$ as an adjusted sum of inertia and FR.

By considering $\mathbf{Y}$ as an adjusted outage size and $\mathbf{Z}$ as an adjusted sum of inertia and FR, an intuitive understanding of the constraint's function can be found. As the outage size increases, more FR and inertia is needed to contain the frequency above $(f_0 - \Delta f_{trig})$ Hz. UFLS effectively reduces the outage size, thus diminishing FR and inertia requirements and expanding the feasible space.

Although UFLS makes the feasible space nonconvex, the constraint remains similar to a convex SOC, and indeed tends towards the $u = 0$ SOC as Y increases, because $\mathbf{Y}^2 - u^2 \approx \mathbf{Y}^2$ when $\mathbf{Y} \gg u$. Due to this asymptotic behaviour, any SOC tangential to the nonconvex constraint will never intersect it. Thus the SOC will approximate the true feasible space in a convex and conservative manner, as shown in Fig. 4.2.

Solutions that lie on (4.12) are optimal because they secure the frequency with the least FS, and therefor lowest cost. A convex conservative SOC approximation reduces the feasible space, increasing the required FS and cost. The least conservative SOC minimises this feasible space reduction, thus minimising cost increases. For a given P^{UFLS} , this is found by minimising




Figure 4.2: Convexification of the nadir constraint (4.12) using a second order cone (4.17). For a fixed UFLS amount, the SOC's gradient (g) and Z intercept (c) are defined by the tangential point $(\hat{Y}, \hat{Z})$. The SOC approximation with the minimum A_{con} is the least conservative.

the area between the true constraint and the approximation, labelled A_{con} in Fig. 4.2. We are interested in minimising the reduction in feasible space only, thus A_{con} is further bounded by Z_{min}, Z^{max} planes. The $\mathbf{Z}$ limits are found by considering the upper and lower bounds of $\mathbf{R}_1, \mathbf{R}_2, \mathbf{H}$. These come from (4.4) and the system specific limits on the provision of these FS from plant and battery levels.

Each point on the constraint's surface has only one pair of SOC gradient (g) and zero offset (c) that gives the tangential SOC. Thus due to the symmetry about the origin, these are functions of $\mathbf{Z}$. g is the gradient of the constraint (4.12) at the $\mathbf{Z}$ value of interest ($\hat{Z}$):

$$\hat{Y}(\hat{Z}) = \sqrt{\hat{Z}^2 + u^2} \quad (4.13)$$

$$g(\hat{Z}) = \frac{\delta Z}{\delta Y}(\hat{Y}(\hat{Z})) = \frac{\hat{Y}}{\sqrt{\hat{Y}^2 - u^2}} \quad (4.14)$$

Then c is found as the value that a line with gradient g passing through $(\hat{Y}, \hat{Z})$ crosses the Z axis:

$$c(\hat{Z}) = \hat{Z} - g(\hat{Z}) \cdot \hat{Y} \quad (4.15)$$

The g and c that minimise the area $A_{con}(g, c)$ (labelled here as (g_{opt}, c_{opt})) result in the least conservative SOC approximation because they reduce the feasible space the least. These are

found offline by searching over the feasible $\hat{Z}$ range ($Z_{min} : Z^{max}$). Thus the new convex, least conservative SOC nadir constraint is given by:

$$\mathbf{Z} \geq g_{opt} \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} + c_{opt} \quad (4.16)$$

$$\frac{1}{g_{opt}}(\mathbf{Z} - c_{opt}) \geq \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} \quad (4.17)$$

4.2.3 Optimising UFLS amount

A constant P^{UFLS} facilitates the convexification, but to optimally balance the UFLS as an option against other FS, an optimizer must be able to choose the amount of UFLS to schedule. Binary variables enable this by successively relaxing any nadir constraint with less UFLS, successively increasing the feasible space and the expected UFLS cost. Assuming:

$$P_1^{UFLS} < P_2^{UFLS} < \dots < P_N^{UFLS} \quad (4.18)$$

Then the Big M technique can be implemented with (4.17) to allow UFLS level choice:

$$\mathbf{Z} + M(\mathbf{m}_1 + \mathbf{m}_2 + \dots + \mathbf{m}_K) \geq \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} \quad (4.19a)$$

$$\frac{1}{g_{opt}(P_1^{UFLS})}(\mathbf{Z} - c_{opt}(P_1^{UFLS})) + M(\mathbf{m}_2 + \dots + \mathbf{m}_K) \geq \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} \quad (4.19b)$$

.....

$$\frac{1}{g_{opt}(P_{K-1}^{UFLS})}(\mathbf{Z} - c_{opt}(P_{K-1}^{UFLS})) + M \cdot \mathbf{m}_K \geq \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} \quad (4.19c)$$

$$\frac{1}{g_{opt}(P_K^{UFLS})}(\mathbf{Z} - c_{opt}(P_K^{UFLS})) \geq \sqrt{\mathbf{X}^2 + \mathbf{Y}^2} \quad (4.19d)$$

$$\mathbf{m}_1 + \mathbf{m}_2 + \dots + \mathbf{m}_K \leq 1 \quad (4.20)$$

The above formulation is a convex SOC because the left hand side remains a linear expression of decision variables. Thus the nadir constraint (4.19) can be efficiently solved by any convex solver with MISOCP capabilities, facilitating its tractable application to a wide range of scheduling

and market problems. To reduce solver times, it is recommended to keep M as small as possible and to add the binary variables m_k at each node to a special ordered set.

This nadir constraint is an improvement upon the previous most advanced consideration of UFLS [30], which constrains inertia and FR from thermal plants to secure against a fixed loss. We replace the piece-wise linear constraint at each UFLS level with a SOC. Critically, this allows the largest loss size to be a decision variable, and can consider the contributions of two FR speeds alongside UFLS and inertia. One potential drawback of the above nadir constraint is that the inability to relax the SOC relating to the highest UFLS level may result in overly conservative results for particularly large losses. This did not impact the empirical results obtained in this paper so it is left to future work to find ways to ameliorate this problem such inclusion of binary variables in (4.19d).

The cost of disconnecting P^{UFLS} GW of load depends on the value of the load that is disconnected (c^{UFLS}) and the disconnection length. However, despite being costly when it is utilised, large outages occur rarely so UFLS is seldom triggered, but can substantially reduce the requirements on the other frequency services. To optimally balance these competing effects, the cost of UFLS must be found probabilistically by multiplying the cost of UFLS post-outage by the probability that the outage occurs. This expected cost of UFLS is added to the cost function, and included in the FSC:

$$C^U = t_{rec} \cdot p^{out} \cdot N_{PL_{max}} \left(\sum_{k=1}^K m_k \cdot c_k^{UFLS} \cdot P_k^{UFLS} \right) \quad (4.21)$$

Where $N_{PL_{max}}$ is the number of committed largest generators. This assumes that all of these generators are operating at the same power level, as is the case in our simulations.

Finally, (4.21) does not accurately characterise the expected cost of UFLS if any credible losses smaller than PL_{max} would trigger the demand disconnection. This is possible if inertia and FR levels are sufficiently reduced by large UFLS scheduling, that the smaller loss would cause frequency to breach the trigger level. To prevent this the operator must enforce an additional nadir constraint (4.12), with $P^{UFLS} = 0$ and PL_{max} equal to the next smallest credible loss. This guarantees that UFLS will only be activated for PL_{max} , and thus that (4.21)

reflects the true probabilistic cost of UFLS.

4.3 Risk-Constrained UFLS Feeder Allocation Under Uncertain Distributed Generation

The previous section details a method to schedule a set amount of UFLS (P^{UFLS}) to assist with frequency containment alongside inertia and FR. However, appropriately selecting feeders to ensure P^{UFLS} of load disconnection when the frequency breaches the trigger level is made difficult by the uncertainty introduced into feeder net power flow by DG.

This section presents a method to solve this second-stage problem. The uncertainty is explicitly dealt with using a moment-based distributionally robust chance constraint (DR-CC), a technique also used in Section 3.2.1. The presented method leverages distributional information on the uncertain feeders' net-power flow to select feeders efficiently, whilst providing guarantees on a minimum volume of UFLS that would be delivered in the event of this service being triggered. This method is useful in combination with all types of UFLS schemes, although its application varies depending on when the decision on which relays will be disconnected is made (offline before the outage [95] or online during the transient period following an outage [82, 83, 84]).

Three formulations for the optimisation problem are discussed:

1. Deterministic allocation: represents the current status quo for selecting feeders that contribute toward the UFLS scheme.
2. Chance-constrained under Gaussian uncertainty: assumes that the net-power flows are Gaussian variables with known means and standard deviations. .
3. Distributionally robust chance-constrained: assumes only knowledge of the net-power flows means and standard deviations, with no assumption on their specific probability distributions. Applicable to systems with diverse constituent feeders, where uncertainty arises from different sources such as DG and various loads (e.g. electric vehicles and heat pumps).

4.3.1 Deterministic allocation approach

Depending on the UFLS scheme deployed, the desired load to be shed can be expressed as a set percentage of national demand or a set amount of demand. Here both problems are formulated mathematically and deterministically, before demonstrating their equivalence. This is relevant because it means the learnings from case studies performed for either scheme are transferable.

First, we consider the problem of selecting the optimal feeder combination to deliver a set percentage of national demand ($L^\%$). This is a simplistic deterministic method where the net-load through each feeder is always assumed to follow a certain percentage of national demand. This approach represents the current practice used extensively by system operators such as National Grid ESO in Great Britain [24]. This formulation is defined by the following mixed-integer linear program:

$$\min_{\mathbf{x}} \quad (P^\%)^T \cdot \mathbf{x} \quad (4.22a)$$

$$\text{subject to} \quad L^\% \leq (P^\%)^T \cdot \mathbf{x} \quad (4.22b)$$

$$\mathbf{x}_i \in \{0, 1\} \quad \forall i = 1, \dots, m \quad (4.22c)$$

The objective is to minimise the percentage of load disconnected during the UFLS scheme, whilst ensuring that the minimum threshold for effective load shedding ($L^\%$), is met. Column vector ' $P^\% = [P_1^\%, P_2^\%, \dots, P_m^\%]^T \in \mathbb{R}^m$ ' corresponds to the % of national demand assumed for each of the m participant feeders in the UFLS scheme. The penetration of DG within distribution networks make the forecast of this value uncertain. However, in this simple method that represents the current practice, the stochasticity is ignored and a fixed value is used (e.g. the mean or a fixed percentile). Column vector ' $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m]^T \in \{0, 1\}^m$ ' corresponds to the binary decision variables in the optimisation (i.e., selection of a feeder contributing towards the UFLS scheme or not).

Second, the formulation for a set amount UFLS scheme is given by:

$$\min_{\mathbf{x}} \quad P^T \cdot \mathbf{x} \quad (4.23a)$$

$$\text{subject to} \quad L \leq P^T \cdot \mathbf{x} \quad (4.23b)$$

$$\mathbf{x}_i \in \{0, 1\} \quad \forall i = 1, \dots, m \quad (4.23c)$$

This formulation is the same as in (4.22) except that the column vector ' $P = [P_1, P_2, \dots, P_m]^T \in \mathbb{R}^m$ ' now represents a set amount of UFLS expressed in MW, in the same way as the minimum UFLS requirement ' L ' is defined in MW. The reason for this equivalence is because translating a set amount of load to a set percentage of national demand is trivially achieved by dividing by the total national demand. Due to this equivalence, from here on only the set amount scheme will be considered in mathematical formulations and case studies. The importance and implications of moving to a MW-requirement for UFLS are explained in detail in Section 4.3.3.

Before the rise of distributed RES, this deterministic method was appropriate for scheduling the UFLS services, as the net load through each feeder did follow fairly regular patterns. However, given the stochastic nature of renewable generation, this is no longer the case, and the likelihood of violating requirement (4.23b) could increase significantly. The key parameter in the deterministic formulation presented in this section is the forecast for net-load in each feeder, vector P . As will be shown in Section 4.4.3, this formulation is not suitable for choosing a pre-defined risk in violating the minimum load-shedding threshold L : either an overly optimistic result (i.e. lower total UFLS than L would typically be available when this service is triggered) or overly conservative (i.e. potentially much higher load being disconnected than the minimum amount required) would be achieved with this method. The UFLS-allocation formulations presented here overcome this problem.

4.3.2 Chance-constrained approach

A chance-constrained optimisation approach is proposed here to account for the stochastic nature of feeders' net-load. The formulation in (4.23) is modified by including a chance constraint

which allows the system operator to specify an acceptable risk of UFLS under-delivery:

$$\min_{\mathbf{x}} \quad \mu_{\tilde{P}}^T \cdot \mathbf{x} \quad (4.24a)$$

$$\text{subject to} \quad \mathbb{P}\left(L \leq \tilde{P}^T \cdot \mathbf{x}\right) \geq 1 - \epsilon \quad (4.24b)$$

$$\mathbf{x}_i \in \{0, 1\} \quad \forall i = 1, \dots, m \quad (4.24c)$$

The objective is to minimise the sum of expected load disconnection from all feeders, where column vector ' $\mu_{\tilde{P}} = [\mu_1, \mu_2, \dots, \mu_m]^T \in \mathbb{R}^m$ ' contains the mean values of net-load of each feeder participating in the UFLS scheme (i.e. $\mu_{\tilde{P}} = E[\tilde{P}]$). Constraint (4.24b) ensures that the volume of UFLS disconnected is higher or equal to the minimum threshold ' L ', with a probability equal or greater than ' $1 - \epsilon$ '. In other words, ' ϵ ' represents the risk of violating the minimum UFLS threshold. Column vector ' $\tilde{P} = [\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_m]^T \in \mathbb{R}^m$ ' contains the random variables for net-load in each feeder.

The above formulation is non-deterministic and therefore impossible to solve. However, if the means and covariances of $\tilde{P}$ are known (i.e., the first two moments of the random variables), moment-based distributionally robust chance constraint methods allow its reformulation into a tractable convex second-order cone problem [70]. Here two methods are presented depending on the distributional assumptions made on $\tilde{P}$.

Start by defining a new auxiliary variable:

$$\boldsymbol{\delta} = L - \tilde{P}^T \cdot \mathbf{x} \quad (4.25)$$

This variable represents the cumulative under-delivery of UFLS compared to what is scheduled. It is desirable that this is negative with low probability, ensured when it is substituted into (4.24b).

$$\mathbb{P}(\boldsymbol{\delta} \leq 0) \geq 1 - \epsilon \quad (4.26)$$

The mean and standard deviation of $\boldsymbol{\delta}$ are functions of the mean vector ($\mu_{\tilde{P}}$) and covariance

matrix (Σ) of $\tilde{P}$:

$$\mu_{\delta} = L - \mu_{\tilde{P}}^T \cdot \mathbf{x} \quad (4.27)$$

$$\sigma_{\delta} = \sqrt{\mathbf{x}^T \Sigma \mathbf{x}} = \|\Sigma^{\frac{1}{2}} \mathbf{x}\|_2 \quad (4.28)$$

A covariance matrix is always symmetric and positive semi-definite. Accordingly, $\Sigma^{\frac{1}{2}}$ refers to the unique matrix that is positive semidefinite and such that $\Sigma^{\frac{1}{2}} \Sigma^{\frac{1}{2}} = (\Sigma^{\frac{1}{2}})^T \Sigma^{\frac{1}{2}} = \Sigma$.

The variable δ can be normalised to have zero mean and unity variance (leading to the new variable δ_n), and then substitute it into (4.26) to give:

$$\mathbb{P} \left(\delta_n \leq \frac{-\mu_{\delta}}{\sigma_{\delta}} \right) \geq 1 - \epsilon \quad (4.29)$$

The probability that δ_n is less than or equal to a constant is given by ' $F_{\mathcal{P}}(\cdot)$ ', the cumulative distribution function (CDF) of δ_n . Thus the above is equivalent to:

$$F_{\mathcal{P}} \left(\frac{-\mu_{\delta}}{\sigma_{\delta}} \right) \geq 1 - \epsilon \quad (4.30)$$

The exact form of the CDF depends on the assumed probability distribution of δ , which is referred to as ' $\mathcal{P}$ '. This chapter presents the convex reformulation of (4.30) under the assumption that the underlying uncertainty in net-power flow through feeders follows Gaussian distributions (i.e., $\mathcal{P} = \mathcal{N}(\mu_{\delta}, \sigma_{\delta}^2)$), as well as a less specific distributionally robust assumption. Comparison of the efficiency and appropriateness of each of these assumptions is explored in the case studies.

Convex reformulation of chance constraint under Gaussian uncertainty

Due to the summative property of Gaussian distributions, if it is assumed that the net-power flow uncertainties ($\tilde{P}$) are Gaussian distributed, then so too will be δ_n . Thus, given the well defined CDF of Gaussian variables ' $\Phi(\cdot)$ ', constraint (4.30) can be reformulated:

$$\frac{-\mu_{\delta}}{\sigma_{\delta}} \geq \Phi^{-1}(1 - \epsilon) \quad (4.31)$$

Rearranging and substituting in the moment definitions:

$$\frac{\mu_{\tilde{P}}^T \cdot \mathbf{x} - L}{\Phi^{-1}(1 - \epsilon)} \geq \sqrt{\mathbf{x}^T \Sigma \mathbf{x}} \quad (4.32)$$

Given that the under-delivery of UFLS is desired with low probability (certainly lower than 50%), this equates to assuming ‘ $\epsilon < 0.5$ ’ and thus ‘ $\Phi^{-1}(1 - \epsilon) > 0$ ’. Therefore, eq. (4.32) is a convex second-order cone (SOC) constraint of the standard form:

$$\|A\mathbf{x} + b\|_2 \leq c^T \mathbf{x} + d \quad (4.33)$$

where:

$$A = \Sigma^{\frac{1}{2}} \quad ; \quad b = 0 \quad ; \quad c = \frac{\mu_{\tilde{P}}}{\Phi^{-1}(1 - \epsilon)} \quad ; \quad d = \frac{-L}{\Phi^{-1}(1 - \epsilon)} \quad (4.34)$$

The derived constraint compels the optimiser to find a combination of feeders that will result in at least ‘ L ’ MW of load shed with ‘ $(1 - \epsilon)\%$ ’ probability. If a given feeder ‘ i ’ is chosen by the optimal allocation, making its corresponding binary variable equal unity ($\mathbf{x}_i = 1$), its mean value will increase the left-hand side of (4.32), whilst its variance and covariance coefficients will increase the right-hand side. Thus, the mathematical structure of constraint (4.32) is aligned with the intuitive result that feeders with a high mean and low uncertainty (i.e. low DG penetration) are most desirable.

Convex formulation of a distributionally-robust chance constraint

The expression in (4.32) describes the reformulation of the chance constraint (4.24b) when the probability distribution of $\tilde{P}$, and therefore $\boldsymbol{\delta}$, is Gaussian. However, this is a restrictive assumption that in some situations need not hold. In this section, a second reformulation of the chance constraint is derived. In this case, the objective is to have an expression that guarantees compliance with the constraint when limited information about the characterisation of uncertainty in net-load in each feeder ($\tilde{P}$) is available. It is assumed here that only the mean and standard deviation of each component in $\tilde{P}$ is known, but no information on the PDFs is available. The formulation presented in this section therefore provides a convex reformulation

for the chance constraint for cases when at least some of the feeders' net-load does not follow Gaussian distributions, or information is lacking on their actual distribution.

According to (4.27) and (4.28), knowledge of the mean and standard deviation of the net-load in each feeder implies knowledge of the mean and standard deviation of $\boldsymbol{\delta}$. However, the lack of knowledge on the distributions of $\tilde{P}_i$ makes the true distribution of $\boldsymbol{\delta}$ ambiguous. In other words, the true distribution ' $\mathcal{P}$ ' of $\boldsymbol{\delta}$ could be any within the set of distributions with $\mu = \mu_{\boldsymbol{\delta}}$ and $\sigma = \sigma_{\boldsymbol{\delta}}$, called the ambiguity set ($\mathcal{A}$).

This ambiguity inhibits the use of an exact form of $\boldsymbol{\delta}$'s CDF being used to reformulate (4.30), as was done for the Gaussian case in Section 4.3.2. Instead, (4.26) should be reformulated into convex form such that it always holds, even for the worst case distribution within $\mathcal{A}$. In other words, a distributionally robust formulation is sought.

Defining a lower bound on ' $F_{\mathcal{P}}(\cdot)$ ', the CDF of $\mathcal{P}$ within $\mathcal{A}$, allows the inequality within (4.26) to be maintained and progress to be made. For some positive constant λ :

$$f_{\mathcal{P}}(\lambda) = \inf_{\mathcal{P} \in \mathcal{A}} F_{\mathcal{P}}(\lambda) \quad (4.35)$$

$$F_{\mathcal{P}}(\lambda) \geq f_{\mathcal{P}}(\lambda) \geq 1 - \epsilon \quad (4.36)$$

Following the method presented in [70], the classical Cantelli inequality can be used to find the appropriate form of $f_{\mathcal{P}}(\lambda)$:

$$f_{\mathcal{P}}(\lambda) = \frac{\lambda^2}{1 + \lambda^2} \quad (4.37)$$

With its corresponding well-defined inverse function, given that a CDF is monotonically increasing:

$$f_{\mathcal{P}}^{-1}(\lambda) = \sqrt{\frac{\lambda}{1 - \lambda}} \quad (4.38)$$

Thus by using (4.38) in (4.36) with ' $\lambda = 1 - \epsilon$ ' and rearranging, the distributionally robust form of the chance constraint (4.24b) is found:

$$\sqrt{\frac{\epsilon}{1 - \epsilon}} \cdot (\mu^T \mathbf{x} - L) \geq \sqrt{\mathbf{x}^T \Sigma \mathbf{x}} \quad (4.39)$$




Figure 4.3: Value of the left-hand side constant for chance constraints (4.32) and (4.39), over a range of risk preference (ϵ). To accommodate any distribution type with a specified mean and standard deviation, the distributionally robust constraint's constant is larger over all ϵ , resulting in a robust and therefore more conservative solution.

Comparing between equations (4.32) and (4.39) reveals that they are the same, except for the constant on the left-hand side, thus (4.39) is also a convex SOC with wide applicability. Some expectations of the performance of the constraints can be deduced from their mathematical formulation. For any ϵ :

$$\sqrt{\frac{1-\epsilon}{\epsilon}} > \Phi^{-1}(1 - \epsilon) \quad (4.40)$$

This means that for the distributionally-robust constraint, ' $\mu^T \cdot \mathbf{x} - L$ ' must be a greater quantity than in the Gaussian assumption chance constraint. Thus, the distributionally-robust constraint yields results that are more conservative, i.e. more load could potentially be disconnected in order to assure meeting the required reliability. Figure 4.3 plots the relationship between these two constants and shows that the relative conservativeness of (4.39) over (4.32) increases as the operator seeks higher delivery confidence (i.e. smaller ϵ).

4.3.3 Applicability to current and future scenarios of power system operation

The methods presented above facilitate the selection of the optimal feeder combination to provide the required UFLS amount, given uncertainty in net-power flow through feeders. The applicability of these methods to real power system operation is now discussed, focusing on the

need for a communication network to update UFLS-relay settings with a certain periodicity, as well as the role that the proposed methods would play, either as primary or secondary set-point for the UFLS relays. For this, it is relevant to consider different current and future scenarios for power system operation:

1. Current practice for UFLS, where relays are tuned very infrequently and the goal is to shed a set percentage of national demand in the event of a large frequency excursion [95]: an uncertainty-aware method for selecting the feeders contributing towards the UFLS scheme, such as the one proposed in this chapter, will be necessary as DG penetration increases. This operational practice has traditionally relied on the assumption that electricity consumption is roughly uniform across the network, which will become increasingly inaccurate with the integration of DG. However, this operational practice could continue if combined with a risk-aware UFLS allocation such as the one discussed here, although it could potentially lead to very conservative outcomes: given that the uncertainty in net-load would have to be considered with a very long look-ahead (e.g. one year, the periodicity for re-tuning UFLS relays), the variance could be very large, leading to very conservative UFLS-allocation results (i.e. a large amount of load being disconnected if the UFLS service is triggered).

Given that no communication network is necessary for this operational practice, relays could be updated manually and very infrequently (e.g. once every year). Then, the method introduced in this chapter would find the best combination of feeders to be chosen for the UFLS scheme, given the uncertainty in their instantaneous power consumption for the next year.

2. ‘Fully controllable’ future operation, where a powerful control and communication network has been deployed and therefore adaptive UFLS is widely available: the method proposed in this chapter could be used as a backup set-point for UFLS relays. Given that WAMS-based adaptive UFLS methods rely on demanding communication requirements, since the set-point of UFLS relays must be updated in a timescale of 1 or 2 seconds following a contingency, it would be highly valuable to make these methods robust against failures in the communication network. The backup set-point determined by the uncertainty-aware

UFLS-allocation method proposed here would be based on local frequency measurements (instead of control signals like WAMS-based adaptive UFLS), that are the last resort if the communication network needed for the adaptive UFLS scheme has failed.

The usefulness of our proposed method as a backup to a WAMS-based adaptive UFLS scheme is demonstrated in Figure 4.4, which plots an adaptive UFLS scheme with and without successful communication. When the communication network is fully operational, the controller detects the fault and trips 150 MW of load at 1.4 s. This shedding of load combines with system inertia and frequency response services to arrest frequency decline within the 0.8 Hz limit that would trigger conventional UFLS relays. However, if communication fails, no load is tripped at 1.4 s and thus frequency declines to 0.8 Hz. Relays detect this via local frequency measurements and trip 300 MW of load to restore demand-generation balance and arrest frequency decline. The constituent feeders for this ‘backup’ UFLS must be chosen ahead of time and are subject to the same net-load uncertainty from DG as traditional UFLS schemes, therefore a risk-aware feeder selection method such as the one proposed in this chapter would be needed.

3. ‘Frequently tunable’ future operation, where a communication network is in place, but the system operator does not rely on updating UFLS-relay settings in a sub-second timescale. The framework for co-scheduling UFLS alongside inertia and FR presented in this chapter falls into this category. Our feeder-allocation method would enable the optimal operation of the UFLS scheme in this scenario. The communication network would frequently update the settings of UFLS relays (e.g. every hour), where the risk-aware UFLS-allocation method would be used to select the appropriate feeders, given the uncertainty in net-load through each feeder in the next hour. Based on this set-point updated hourly, local frequency measurements would make the UFLS relays trip if the frequency excursion exceeds the pre-defined threshold.

Furthermore, the chance-constrained allocation method would provide optimal performance when working in tandem with UFLS-scheduling method presented in Section 4.2.2. First, the optimal amount of UFLS to be scheduled would be found following the procedure detailed in




Figure 4.4: Frequency excursions following a generation loss, for a WAMS-based adaptive UFLS scheme with operational and unavailable communications. When communications fail, the relays selected for the backup scheme disconnect their load once their local frequency reaches the 0.8Hz trigger level. These relays must be selected ahead of time and thus are subject to net-load uncertainty caused by DG.

Section 4.2.2 and then this amount of UFLS would be optimally allocated among feeders using the chance-constrained method proposed in this section. Given that the triggering mechanism for UFLS relays would be local measurements (i.e. the relays trip when the local frequency drops below the threshold), the communication network would only be used immediately following scheduling to select the feeders (e.g. every hour or half-hour, each time the energy and ancillary services markets are cleared), and thus this procedure is robust to communication delays or failure during the safety-critical situations that require UFLS. Nevertheless, the updating of UFLS relays could also be done sub-hourly, if information on the net-load through feeders is received with sub-hourly resolution: the closer to real-time the selection of feeders is updated, the lower the forecast error for DG power output will be, therefore the lower the conservativeness of the chance-constrained UFLS allocation will be.

4.4 Constraint Verification

This section uses dynamic simulations from the solution of (4.2), to verify that the nadir constraint (4.19) does contain the system frequency, whilst also offering insight into how the different FS interact to do so. The value of the constraint for system operation is demonstrated by its application to the GB 2030 system in Section 4.5.

Table 4.1: Dynamic Simulation parameters

	P^{UFLS} (GW)	R_1 (GW)	R_2 (GW)	H (GWs)	PL_{max} (GW)
A	0.0	0.2	2.4	130.7	1.6
B	0.0	0.2	2.4	169.8	1.8
C	0.6	0.2	2.4	146.3	1.8
D	0.6	0.3	2.4	127.7	1.8

The second part of this section demonstrates the efficiency and security of the load shedding allocation introduced in the previous section. The optimal selection of feeders to provide 250MW of UFLS (as decided by (4.19) or some other UFLS method) under various power flow distribution and correlation assumptions.

4.4.1 Nadir Constraint Verification

Figure 4.5 plots $\Delta f^u(t)$ following a PL_{max} outage for the system conditions described in Table 4.1, all of which lie on (4.19). The FR time constants are $T_1 = 1s$, $T_2 = 10s$. The nadir will always coincide with the trigger level, because constraint (4.19) guarantees that the triggered UFLS is sufficient to close the generation-demand deficit, arresting frequency drop.

No UFLS is scheduled in cases A and B. Consequently, there is just enough FR and inertia in both cases to ensure that the frequency reaches, but does not cross the UFLS trigger level. However, to secure a loss only 0.2 GW smaller than case B, case A needs 39.1 GWs less inertia. This is equivalent to the inertia from approximately 20 combined cycle gas turbines (CCGTs). Similarly, comparison between cases C and D shows that fast FR is also an effective FS, with a 0.1 GW difference reducing inertia requirement by 18.6 GWs (≈ 9 CCGTs).

Case C has 20.9 GWs (≈ 10 CCGTs) less inertia than case B. This reduction means that the system frequency breaches the trigger level and UFLS is initiated. The step change in demand reverts the frequency gradient to zero, thus the trigger level corresponds to the nadir. In other words, UFLS has reduced the system's inertia requirements at the expense of an increased UFLS expected cost (4.21). The constraints developed in this chapter offers the optimiser a tool to find the optimal balance between these competing effects during scheduling.




Figure 4.5: Frequency evolution after a loss of PL_{max} for the system conditions in Table 4.1.

Table 4.2: Test system with 20 candidate feeders for the UFLS scheme, with varying levels of uncertainty in net-load of each feeder.

Feeder #	$\mu_{\tilde{P}}$ (MW)	$\sigma_{\tilde{P}}$ (MW)	Feeder #	$\mu_{\tilde{P}}$ (MW)	$\sigma_{\tilde{P}}$ (MW)
1	10.00	4.34	11	22.00	2.10
2	21.00	3.06	12	38.00	2.28
3	34.00	4.88	13	28.00	3.48
4	30.00	3.21	14	22.00	4.95
5	16.00	2.62	15	27.00	4.00
6	35.00	3.44	16	40.00	4.51
7	33.00	3.41	17	34.00	3.72
8	20.00	3.92	18	19.00	3.41
9	29.00	1.04	19	18.00	1.51
10	13.00	4.42	20	16.00	1.77

4.4.2 Feeder Allocation Method - Verification and Case Studies

This section demonstrates the efficiency and security of the load-shedding allocation method introduced in the previous section. The test system contains 20 candidate feeders with varying degrees of uncertainty, described in Table 4.2. The uncertainties are assumed independent unless otherwise specified. Minimum permissible load shed L was 250MW. The optimisation problems were modelled using the YALMIP toolbox for MATLAB [97], calling Gurobi as the numerical solver.

The case studies in this section were run by my co-author of the paper [3], Maria Parajeles, as part of her master's project. They are included here to validate and contextualise my theoretical work presented in section 4.3.

Table 4.3: Mathematical description of the different PDFs considered.

	Probability Distribution Function	Parameters	
		Location	Scale
Gaussian	$f(x \mu, \sigma) = \frac{1}{\sigma \cdot \sqrt{2\pi}} \cdot e^{-\left(\frac{x-\mu}{\sigma}\right)^2}$	$\alpha = \mu$	$\beta = \sigma$
Gumbel	$f(x \mu, \sigma) = \frac{1}{\beta} \cdot e^{\frac{x-\mu}{\beta}} - e^{\frac{x-\mu}{\beta}}$	$\alpha = \mu + \beta \cdot \gamma^*$	$\beta = \frac{\sqrt{6}}{\pi} \cdot \sigma$
Laplace	$f(x \mu, \sigma) = \frac{1}{2 \cdot \beta} \cdot e^{\frac{ x-\mu }{\beta}}$	$\alpha = \mu$	$\beta = \frac{\sigma}{\sqrt{2}}$
Student's t	$f(x \nu^{**}) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu \cdot \pi} \cdot \Gamma(\frac{\nu}{2})} \cdot \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}}$	$\alpha = 0,$ for $\nu > 0$	$\beta = \frac{\nu}{\nu-2},$ for $\nu > 2$

* Euler-Mascheroni constant

** Degrees of freedom

Different underlying probability distributions for the uncertainty in feeder's net-load are considered. First, the Gaussian PDF is taken as a reference due to its common use in probabilistic forecasting of both load and DG in power systems [98, 99, 100]. Second, the Gumbel distribution is used to consider higher probability of low net-load values (i.e., greater DG output), which has been applied to modelling uncertainty in electric load and renewable generation [98, 101]. Finally, the Laplace and Student's t distributions were chosen to model heavy-tailed probabilistic forecasts, as these distributions are frequently used in the literature to model uncertainty in solar and wind generation [102, 103].

4.4.3 Limitations of a deterministic optimisation approach

Firstly, the deterministic optimisation approach defined in (4.23) for allocating UFLS under uncertainty is considered, as described in Section 4.3.1. For this approach, a fixed percentile of the net-load forecast in each feeder must be chosen before solving the optimisation, as this formulation does not enable the operator to pre-specify the desired level of risk for the overall UFLS-allocation result. That is, if the system operator is willing to accept a 1% chance of violating the minimum UFLS threshold ' L ', using this deterministic formulation implies that a guess must be made on the percentile of net-load in each feeder that would lead to the desired level of overall risk.

To illustrate this point, six different percentiles of net-load for each feeder described in Table 4.2 are found and input into (4.23). Without loss of generality and for the sake of simplicity, Gaussian uncertainty for the forecast of net-load is assumed for now. The 1%, 10%,

Table 4.4: Results of allocation of load shedding using a deterministic optimisation approach.

Optimisation percentile (%)	Samples below disconnection requirement (%)	Expected load disconnection (MW)
1	0.00	374.02
10	0.01	289.01
20	0.30	282.97
30	5.89	265.96
40	23.71	256.96
50	50.08	249.98

20%, 30%, 40% and 50% percentiles of this PDF are found for each of the 20 feeders, and used as the values within vector P in eqs. (4.23a) and (4.23b). The results of solving the optimisation problem in (4.23) with each of these percentiles are included in Table 4.4, where two metrics are shown: 1) the percentage of occurrences that would violate the load disconnection requirement ‘ L ’, computed by sampling 100,000 times from the Gaussian PDF for each feeder chosen in the optimisation; and 2) the expected load disconnection, which corresponds to the sample mean of the sampling process carried out to obtain the previous metric.

The results demonstrate that using a percentile of 1% for each feeder leads to an extremely conservative allocation of load shedding, since the combined probability of all feeders falling below the 1th percentile forecast error for net-load is almost negligible. To achieve an overall risk of 1%, a percentile between 20% and 30% for each feeder must be chosen. This highlights the unsuitability of the deterministic scheme when dealing with stochastic net-load through feeders in future systems with high DG penetration: a system operator is only concerned with the aggregate load shed being guaranteed above a certain level, but there is no fixed relationship between the individual percentile chosen and this aggregate security level. This means the operator can only guarantee security by being overly conservative and incurring very high levels of potential load shed.

Therefore, a deterministic optimisation approach, which is fairly representative of current practice in most systems, is inadequate for the UFLS-allocation problem under significant DG penetration. Future power systems will require to explicitly account for uncertainty within the UFLS-allocation problem.

4.4.4 Chance-constrained approach

Impact of correlation in net-load probabilistic forecasts across the network

This section demonstrates the performance of the chance-constrained approach in the allocation of load shedding, which not only allows to achieve a pre-defined level of risk, but notably enables consideration of correlation in net-load of the different feeders. Since uncertainty in the net-load of active feeders is largely driven by DG output, which in turn is driven by meteorological conditions, correlation in feeders that are geographically close is an important characteristic that must be accounted for [104]. Again, Gaussian uncertainty is used here for simplicity, while the learnings on the importance of considering correlations are applicable to any other probability distributions.

The chance-constrained optimisation defined in (4.24) is solved while considering two different covariance matrices ' Σ ', in eq. (4.32). This demonstrates the importance of accounting for correlation in forecasts. Covariance matrices are symmetric positive semi-definite. For the case where the net load through feeders are non-correlated, the covariance matrix is a diagonal matrix containing the variances in Table 4.2. To generate the covariance matrix for the correlated case an observation matrix is first generated by sampling the load flow through each feeder 5 times. These observations were then adjusted to have sample means and std close to the true values. From this observation matrix, the MATLAB covariance function was used to ascertain the covariance matrix. This process was iterated until a satisfactory covariance matrix was found, defined here as one where most feeders are positively correlated to reflect the shared impact of weather across them.

Correlation coefficients are numbers between -1 and 1 that represent the strength and direction of a relationship between two variables. This is a useful tool for visualising the relationship between random variables such as the net load flow through feeders. For example, a +ve correlation value between feeders 1 and 2 indicates that if feeder 1 has a higher-than-average sample for net load flow then that is also more likely for feeder 2. As such heatmaps for the correlation coefficient matrices are depicted in Fig 4.6. Table 4.5 shows the results for two cases: Case 1 assumes uncorrelated random variables for the net-load of feeders (i.e., the




Figure 4.6: Heatmap for the correlation coefficient matrices of the uncorrelated and correlated net-load probabilistic forecast.

Table 4.5: Results of neglecting (Case 1) and considering (Case 2) correlation among feeders' net-load forecasts.

	ϵ (%)	Samples below disconnection requirement (%)	Expected load disconnection (MW)
Case 1 (correlation neglected)	1	19.30	275.98
	2	24.61	273.01
Case 2 (correlation considered)	1	0.99	343.07
	2	1.99	328.02

correlation coefficient matrix is diagonal, only ones), while Case 2 does consider correlation in the optimisation problem. The sampling for feeders' net-load performed after solving the optimisation does account for the correlation in both cases, since DG outputs are in fact correlated. The results show that when correlation is neglected as in Case 1, the optimisation is overly optimistic: the UFLS-threshold defined in constraint (4.32) is met for only 19% of the samples, while the accepted risk was set to 1%.

Neglecting the correlation in forecasts would always lead to the same performance as shown in Table 4.5: the level of risk desired would always be violated. This is due to the fact that correlated forecast errors tend to go in the same direction, therefore if any feeder has a high forecast error leading to very low net-load, it is likely that other feeders follow the same behaviour. Ignoring this compounded effect of forecast errors leads to an overly optimistic UFLS-allocation solution. This highlights the importance of computing the covariances of the forecast errors among feeders, and using these as inputs to the chance-constrained optimisation. The results in

Table 4.5 demonstrate that the proposed risk-constrained formulation for UFLS allocation can successfully accommodate correlated forecasts, meeting the desired level of overall risk.

Advantages and limitations of assuming Gaussian uncertainty

In this section, the implications of solving the load shedding allocation using the Gaussian-assumption convex reformulation of the chance constraint are analysed. The optimisation problem in (4.24) is solved by substituting constraint (4.24b) by (4.32), and the results are evaluated by comparing the performance of the selected feeders when the underlying uncertainty actually follows a different distribution. Gumbel, Laplace and Student's t distributions are considered, while in all cases the first two moments are the ones defined in Table 4.2.

Table 4.6 shows the results for two different risk thresholds, $\epsilon = 1\%$ and $\epsilon = 2\%$. The risk of shedding less load than required does indeed remain below 1% or 2%, when uncertainty follows a Gaussian distribution. The advantage of constraint (4.32) is that it provides the least conservative feasible set for the UFLS allocation, if uncertainty can indeed be accurately modelled with Gaussian PDFs. However, when the underlying uncertainty in net-load through feeders follows a different PDF, the reliability requirements are not met. The highest deviation from the desired risk level is obtained when uncertainty follows a Gumbel distribution, which yields a 1.66% risk for a desired value of under 1%. If the distribution is heavily tailed to both extremes, or to the left extreme, as shown in Figure 4.7 for Laplace and Student's t , and Gumbel distributions, respectively, the probability of not shedding the minimum required load would be higher than acceptable.

Distributionally-robust chance-constrained optimisation

The limitations of the Gaussian chance constraint shown in the previous section can be overcome by using instead the distributionally robust chance constraint (4.39). The results of solving the UFLS-allocation problem with this constraint are shown in Table 4.7, which considers the same types of underlying uncertainty as in the previous section.

As expected, the constraint is robust against different distributions and thus the reliability requirements are always met, i.e., the sampled sum of available load in all selected feeders is

Table 4.6: Results of chance-constrained allocation of load shedding assuming Gaussian uncertainty.

ϵ (%)	Underlying distribution	Samples below disconnection requirement (%)	Expected load disconnection (MW)
1	Gaussian	0.90	269.98
	Gumbel	1.65	270.55
	Laplace	1.11	270.01
	Student's t	1.08	270.01
2	Gaussian	1.94	267.99
	Gumbel	2.90	268.54
	Laplace	2.13	267.99
	Student's t	2.07	268.00

Table 4.7: Results of distributionally-robust chance-constrained allocation of load shedding under different types of uncertainty.

ϵ (%)	Underlying distribution	Samples below disconnection requirement (%)	Expected load disconnection (MW)
1	Gaussian	0.00	357.98
	Gumbel	0.00	358.65
	Laplace	0.00	357.98
	Student's t	0.00	358.01
2	Gaussian	0.00	317.99
	Gumbel	0.00	318.59
	Laplace	0.00	318.00
	Student's t	0.00	318.00

above the required load disconnection L in most observations. In fact, the risk of violating the disconnection requirement is roughly zero for all these four PDFs considered, given that the distributionally-robust chance constraint is suitable for any PDF with these same values for the first two moments, and therefore the solution of the optimal UFLS allocation is conservative. Given that future distribution systems are expected to host heterogeneous load profiles, arising from a combination of DG, electric vehicles, heat pumps, combined heat and power plants, and other devices, the significant advantage of the distributionally-robust allocation lies on its ability to accommodate different types of uncertainty. Furthermore, the exact PDFs for net-load through each feeder are not needed, as simply the mean and standard deviation are required for applying the distributionally-robust method.

To conclude, in this section several relevant case studies have been conducted, demonstrating the importance of accounting for correlation in net-load of different feeders (notably due




Figure 4.7: Sampled distributions of available load in the selected feeders, for the four PDFs considered for underlying net-load uncertainty. Feeders were selected through a chance-constrained optimisation assuming Gaussian uncertainty.

to correlated weather forecasts that affect the DG power output). Neglecting this positive correlation has been shown to be a risky practice, since the probability of falling short in total load disconnection would be higher than expected. The proposed chance-constrained formulation can also accommodate varying degrees of information on the characterisation of uncertainty, as a distributionally-robust framework has been deduced, which requires knowledge of only the two first moments of the underlying random variables. This DRO approach is the first proposed method to give explicit security guarantees under heterogeneous power-flow distributions, and has wide utility due to its ability to complement first-stage UFLS methods.

4.5 Case Studies

The SUC model, with cost function (4.1), was used to identify the savings in annual FSCs from UFLS. Each node is subject to: RoCoF constraint (4.4); minimum FR requirement (4.5); nadir constraint (4.19,4.20); demand-power balance (3.30); and the generator and storage constraints listed in Section 3.3.2. The simulation finds the optimal UFLS amount to schedule.




Figure 4.8: Sampled distributions of available load in the selected feeders, for the four PDFs considered for underlying net-load uncertainty. Feeders were selected through a distributionally-robust chance-constrained optimisation.

The secondary problem of mapping this to specific feeders could be achieved by deploying the methods presented in Section 4.3, which Section 4.4.2 proved effective. However, for clarity and simplicity feeder allocation is not explicitly addressed here.

Each simulation corresponds to four months of operation of a representative GB 2030 system. The rolling planning approach described in Section 3.3 is used to sequentially simulate one-month operation from each season. The season dictates the wind and demand levels as well as the stochastic properties of the wind. A scenario tree branching 7 times at the root node only was used to account for net demand forecast uncertainty, as detailed in [20], with quantiles of 0.005, 0.1, 0.3, 0.5, 0.7, 0.9 and 0.995.

A typical time-series of demand was used, with a range of 20:60 GW with daily and seasonal trends. Unless otherwise stated there is an installed wind capacity of 35 GW, and a storage and generation mix detailed in Table 4.8. Solar power was not considered. Frequency security requirements were aligned with GB standard, $f_0 = 50$ Hz, $|\Delta f_{max}| = 0.8$ Hz, $RoCoF_{max} = 1$ Hz/s and $T_s = 0.2$ s (10 cycles at 50 Hz). The FR time constants are $T_1 = 1$ s, $T_2 = 10$ s, while $c_{LS} = £30,000/\text{MWh}$ and $c^{UFLS} = £30,000/\text{MWh}$ unless otherwise stated. UFLS was assumed

Table 4.8: Generation and Storage Characteristics

Generation	Nuclear	CCGT	OCGT
Number of Units	4	120	20
Rated Power (GW)	1.8	0.5	0.1
Min Stable Generation (GW)	1.60	0.25	0.05
No-Load Cost (£'000/h)	0.0	4.5	3.0
Marginal Cost (£/MWh)	10	47	200
Startup Cost (£'000)	NA	10	0
Startup Time (h)	NA	3	0
Min up Time (h)	NA	4	0
Inertia Constant (s)	5	4	4
Max R2 Capacity (GW)	0.00	0.05	0.05
Storage	Pumped	Battery 1	Battery 2
Capacity (GWh)	10	1	12
Dis/Charge Rate (GW)	2.6	0.3	3.0
Max R1 Capacity (GW)	0.0	0.6	0.0
Max R2 Capacity (GW)	0.5	0.0	0.0
Dis/Charge Efficiency	0.75	0.95	0.95

to disconnect load for 1h when utilised, implying that a system re-dispatch would be possible within one hour of the outage. PL_{max} is determined by the power output of nuclear plants, which have an assumed outage rate of 1.8 occurrences/yr.

The maximum UFLS considered is 0.8 GW, chosen because it is always less than the load shed (5% of national demand) at the current first UK UFLS trigger level [24]. The large size difference of at least 0.7 GW, between the nuclear and CCGT plants means that application of an additional nadir constraint, as outlined in Section 4.2.3, is unnecessary here. This is because 0.8 GW or less of UFLS never reduces inertia and FR levels to the point where the smaller loss could trigger UFLS.

Simulations were run in an 8 core Intel Xeon 2.40GHz CPU with 64GB of RAM. XPRESS 8.10 was used to solve the optimisations linked to a C++ application via the BCL interface. The MISOCP duality gap was 0.1%.

4.5.1 Evaluation of the Conservativeness of Nadir Constraint Approximation

Equation (4.19) conservatively approximates the true non convex feasible space. This produces a highly tractable MISOCP formulation that guarantees system security, at the cost of a reduced feasible space potentially decreasing a solution's optimality. In operational terms, this translates into the over-scheduling of FS. The severity is measured by calculating the minimum nadir that the scheduled FS could secure (Δf_{trig}^{true}). The closer this value is to Δf_{trig} , the more optimal the solution, and the better the approximation. The nadir constraint with 0 GW UFLS is not an approximation, so it is not conservative.

Fig. 4.9 shows the distribution of Δf_{trig}^{true} for simulations where 2 UFLS levels are available, 0 GW or the value shown on the x axis. For $P^{UFLS} = 0.6$ GW, the median $\Delta f_{max}^{true} = -0.79$ Hz, and the most conservative is -0.73 Hz. Demonstrating that for this level of load shed, the formulation gives solutions that are always less than 10% conservative, with the vast majority being significantly less so.

The clear trend of increasing conservativeness with P^{UFLS} in Fig. 4.9 can be explained by Fig. 4.1. Larger UFLS levels increase the curvature of the nadir constraint, thus approximating it with a straight line is less accurate. In other words, the minimum area of $A_{con}(g, c)$ becomes larger and the approximation more conservative. However, over the moderate range of available UFLS studied here, the constraint is at the worst 25% conservative when $P^{UFLS} = 0.8$ GW. Even so, when $P^{UFLS} = 0.8$ GW the true nadir was below -0.7 Hz over 95% of the time. Thus most solutions are very close to optimal and insight into the value of UFLS is unhindered.

4.5.2 Value of UFLS in reducing FSC

Fig. 4.10 shows the reduction in annual FSC when two levels of UFLS are available, 0 GW or the value shown on the x axis. The FSC for a system with 55 GW installed wind capacity and no UFLS is £1072 m/yr, 10.8% of total system costs. Allowing 0.8 GW of UFLS offers a 52% reduction of £559m/yr, and 0.4 GW by £115m/yr. These reductions include the increase in expected cost of outages from allowing UFLS, and demonstrate that UFLS can significantly




Figure 4.9: The conservative approximation of (4.12) by (4.19) results in the over scheduling of FS, which could secure a stricter nadir (Δf_{trig}^{true}) than that required ($\Delta f_{trig} = 0.8$ Hz). Shown here are the annual inter-quartile range, median and max/min of Δf_{trig} for systems with two levels of UFLS available, 0 GW or the value shown on the x axis.

improve system operation in systems with high renewable penetration.

The UFLS creates value by reducing the FS required from thermal plants whose minimum generation drives wind curtailment during low net demand periods. An example of this phenomenon is shown in Fig. 4.11, which compares the optimal commitment of CCGT plants over the same 3-day period for two identical systems. One has the option to schedule 0.6 GW of UFLS and the other does not. CCGTs are shown because nuclear plants have a fixed commitment and the open cycle gas turbine (OCGT) output is negligible over the depicted period. During the high net demand period on the first day, no UFLS is scheduled because frequency security requirements are met as a by-product of energy supply. Accordingly, both systems have similar plants commitment levels.

Overnight, the demand drops and wind output increases, staying high for the following days. Over this period, the scheduled UFLS makes containing the frequency possible whilst needing less inertia and FR from the thermal plants, so fewer CCGTs are run. This lowers the sum of their minimum generation, enabling the absorption of up to 4.7 GWh more wind energy per hour. UFLS's ability to facilitate the absorption of more zero marginal emission wind energy explains the CO₂ savings in Fig. 4.10. Furthermore, when there is higher installed wind generation capacity low net demand periods are more frequent, thus FSC and emission savings increase.




Figure 4.10: The annual reduction in frequency security costs and emissions from optimally scheduling UFLS to support frequency security, in systems with 35 GW and 55 GW of installed wind capacity.

Load is only disconnected if a nuclear plants outs whilst UFLS is scheduled. The operator risks load shed in order to reduce FR and inertia requirements. The cost savings shown in Fig. 4.10 include the expected increase in load shedding costs. Importantly, the expected amount of load shed to achieve the cost savings is not significant, hence reliability of supply is maintained. For example, 0.8 GW of UFLS in the system with 55 GW of wind decreases FSC by £559m/yr. Fig. 4.12 shows that the expected load shed to facilitate this is 4.25 GWh/yr, only 0.0013% of the total load served. All other systems simulated here have lower expected load shed amounts, demonstrating that the proposed framework can assist with frequency security without significantly degrading supply reliability.

4.5.3 Value of Considering Fast FR and a Reduced Largest Loss

The proposed framework improves upon the current state of the art nadir constraints. In [30], UFLS, inertia and FR from thermal plants are constrained to contain a fixed largest loss. Whereas [28] optimally chooses between inertia, fast and slow FR to contain a dynamically controlled loss. Constraint (4.19) allows UFLS consideration alongside the other four FS, facilitating the operator to secure frequency in the cost optimal manner.

Figure 4.13 shows the total annual FSC for systems with 0.6 GW of UFLS available. The




Figure 4.11: Three-day example of the system operation, showcasing the use of UFLS to accomodate more wind.

(dis)charge rate of Battery 1 in Table 4.8 was adjusted to simulate varying $\mathbf{R}_1$ capacities. The minimum stable generation of nuclear plants was varied to change the lower bound of $\mathbf{PL}_{max}$.

The system with $\mathbf{R}_1 \leq 1.2\text{GW}$ and $\mathbf{PL}_{max} \geq 1.4\text{GW}$ has a FSC cost of £270m/yr. This is 13.2% of the £2,040m/yr FSC cost for the system with fixed $\mathbf{PL}_{max} = 1.8\text{GW}$ and $\mathbf{R}_1 = 0\text{GW}$. This reduction can be attributed to the high effectiveness of these two services in reducing inertia requirements, as demonstrated in Section 4.4. This lowers plant minimum stable generation, decreasing wind curtailment from 27.3 TWh to 4.3 TWh. Unlike our proposed framework, the constraint of [30] cannot optimise over these two services, so is inhibited from realising the extensive cost savings they offer.

Fig. 4.14 shows the reduction in FSC from the ability to schedule 0.6 GW of UFLS, in systems with varying availability of other FS. As the availability of other FS increases, the value of UFLS decreases from £187m/yr in a system with only UFLS, $\mathbf{H}$ and $\mathbf{R}_2$, to £80m/yr in a system with abundant $\mathbf{R}_1$ and a flexible $\mathbf{PL}_{max}$ power output. This reduction is because




Figure 4.12: The expected amount of annual load shed from optimally scheduling UFLS to assist with frequency security, in systems with 35 GW and 55 GW of installed wind capacity.

the various FS compete to reduce inertia requirements, decreasing each others marginal value. However, UFLS offers significant cost savings even for very flexible systems. Thus, a framework unable to consider it, such as [28], cannot operate a system cost optimally.

Interestingly, Fig. 4.14 shows that when $PL_{max} = 1.8GW$, the value of UFLS is approximately constant with increasing R_1 availability. This is because the nadir requirements at this large loss level cause extreme levels of wind curtailment. There is 28.6 TWh of annual wind curtailment with no UFLS or R_1 , corresponding to a FSC of £2.23bn/yr, which is 25% of total system costs. When FSC is this high, additional R_1 and UFLS can simultaneously reduce it, without infringing on the other service's marginal value.

4.5.4 Impact of the Expected Cost of UFLS

According to [105] the value of load-shed in the UK ranges between 2-45 £'000/MWh depending on the consumer type, importance, season and time of day. Reference [106] estimates the range of value of load-shed in mainland Europe to be slightly lower, between 1.3:20 £'000/MWh. In reality it is likely that lower cost loads would be disconnected to support frequency security, so the assumption of $c^{UFLS} = £30,000/MWh$ used so far in this chapter is a conservative one. Fig. 4.15 details the sensitivity of UFLS value to c^{UFLS} in 3 separate cases: 'Free UFLS' when $c^{UFLS} = 0$; 'Fixed UFLS' when UFLS is scheduled at all times; 'Optimal UFLS' when the optimiser can choose when to schedule UFLS.




Figure 4.13: Annual FSC for systems with 0.6 GW of UFLS, and varying fast FR and reduced largest loss capabilities.




Figure 4.14: Savings in annual frequency security costs enabled by 0.6 GW of UFLS in systems with varying fast FR and reduced largest loss capabilities.

Fig. 4.15 shows that allowing the scheduling of UFLS creates FSC savings even at extreme UFLS costs of £140,000/MWh, which is more than 3 times the most expensive value of load-shed in the UK. Thus, this chapter's framework will provide value when applied to the majority of large power systems. As the expected cost decreases, the penalty for scheduling UFLS decreases,




Figure 4.15: Sensitivity of frequency security savings from 0.6 GW of UFLS to the cost of the load shed (c^{UFLS}): ‘Free UFLS’ load shed is zero cost; ‘Fixed UFLS’ UFLS is scheduled at all times; ‘Optimal UFLS’ optimiser chooses when to schedule UFLS.

so it is scheduled more frequently and the FSC savings increase. For example, if c^{UFLS} is halved to £15,000/MWh, then Fig. 4.15 shows that the FSC savings from UFLS increase by £41m/yr.

The expected cost of UFLS ($\mathbf{C}^U$) is the product of the probability that an outage will occur and the cost of load shedding if it does. As shown in (4.21), $\mathbf{C}^U$ is linearly dependent on c^{UFLS} , as well as $t_{rec}, p^{out}, N_{PLmax}$. So the sensitivity results in Fig. 4.15 are applicable to adjusting those values also. In other words, if the nuclear plants outed half as often, or the demand was disconnected for half as long, savings of £41m/yr would be observed as those changes are equivalent to reducing the cost of load shedding from £30,000/MWh to £15,000/MWh.

The ‘Free UFLS’ value in Fig. 4.15 corresponds to a system with a permanently reduced FS requirement. As such it represents the maximum FSC reduction from UFLS possible, and is approached as c^{UFLS} decreases. On the other hand, as c^{UFLS} increases the ‘Optimised UFLS’ value trends towards 0 because the larger cost of outage makes UFLS a less attractive option to contain the frequency. Above $c^{UFLS} = £60,000/\text{MWh}$ ‘Fixed UFLS’ shows that always scheduling UFLS begins to increase system costs. This is because the gains from reduced FS requirements, facilitating higher wind absorption, are entirely offset by the severe cost when an outage occurs. The difference between ‘Fixed UFLS’ and ‘Optimal UFLS’ is the value in allowing the optimiser to choose via binary variables (4.19), to secure the nadir with UFLS or not.

Table 4.9: Impact of a Variable Marginal UFLS Cost

	UFLS Steps (GW)	Marginal UFLS Cost (£'000/MWh)	Time at UFLS level (%)
α	0.0, 0.6	30	49.3, 50.7
β	0.0, 0.2, 0.4, 0.6	30, 30, 30	50.2, 0.5, 0.3, 49.0
γ	0.0, 0.2, 0.4, 0.6	1, 10, 79	21.0, 0.2, 48.9, 29.9
	FSC Savings (£m/yr)	Run Time (h)	
α	170.4	5.7	
β	170.4	9.6	
γ	180.9	7.2	

4.5.5 Impact of an Increasing Marginal Load-Shed Cost

Comparison between case α and β in Table 4.9 shows the impact of discretising the UFLS into smaller steps. The options to schedule 0.2 GW and 0.4 GW alongside 0 GW and 0.6 GW are introduced. This produces no additional FSC savings, while computation time increased by 68%. This is because P^{UFLS} has increasing marginal gains, due to its quadratic factor in (4.12), but only linearly increasing costs (4.21). Accordingly, it is optimal to schedule UFLS levels smaller than 0.6 GW, less than 1% of the year. This justifies the preceding case studies only considering one UFLS level other than 0 GW. This phenomenon also explains the quadratic increase in FSC savings depicted in Fig. 4.10.

However, in reality the marginal c^{UFLS} often increases due to load being shed in the order of value. The formulation can account for this. For case γ in table 4.9, the first 0.2 GW of scheduled load shed costs only £1,000/MWh, the second 0.2 GW costs £10,000/MWh and the final costs £79,000/MWh. In this case it is optimal to schedule 0.4 GW of UFLS for 48.9% of the year, a dramatic increase compared to case β . Explicitly recognising this cost variation facilitates a more optimal system operation, resulting in an annual FSC reduction of £10.6m, despite the fact that the average cost of loadshed is £30,000/MWh when 0.6GW of UFLS is scheduled for both cases β and γ .

4.5.6 Communication Requirements for Scheduling UFLS

This chapter’s framework allows the operator to schedule the optimal amount of UFLS from a set of discrete levels. As demonstrated in Section 4.4, UFLS is scheduled by simply scheduling lower amounts of inertia and FR. This causes the post-fault frequency to drop below the UFLS trigger level. The active relays detect this via local frequency measurement, which initiates load disconnection and arrests frequency drop. No communication between the operator and the relays is needed during the fault. However, the net load through relays changes over time, so to maintain a constant UFLS amount (e.g. 0.6 GW) the operator must adjust the set of active relays that will disconnect when the frequency trigger level is breached. This requires an hourly communication between the operator and the relays. When multiple non-zero UFLS levels are available (as in Section 4.5.5), the set of active relays each hour is chosen so that their net load equals the chosen UFLS level.

The presented framework is also compatible with traditional UFLS schemes, that have no communication with relays, and instead disconnect a set percentage of the system demand [24]. This is done by setting the constant P^{UFLS} in (4.12) equal to the desired fraction of national demand at each timestep, and then calculating c_{opt} and g_{opt} accordingly. Fig. 4.16 compares the annual FSC savings from 0.5%, 1.0% or 1.5% UFLS schemes to those from fixed UFLS amounts. Significant FSC savings of £127m/yr are available from the 1.5% case, where the available UFLS amounts range between 0.3:0.9 GW. However, this is £187m/yr less than the savings from a constant 0.8 GW of UFLS. The lower average amount of UFLS available contributes to this, but it is compounded by the fact that UFLS is most valuable during periods of low net-demand (as shown in Fig. 4.11), which corresponds to when the set percentage UFLS schemes offer the smallest UFLS amounts.

The majority of this section focuses on constant UFLS amounts, despite Fig. 4.16 demonstrating that the framework is compatible with traditional UFLS systems with no communication infrastructure. There are three main justifications for this:

1. For the same or lower maximum UFLS amount, a constant UFLS amount offers significantly more value (up to £187m/yr). Thus there is a large incentive to invest in the requisite




Figure 4.16: FSC savings from two types of UFLS scheme. Constant UFLS, which allows scheduling of a fixed amount of UFLS shown on the x axis, or UFLS that tracks a set percentage of national demand. The range of UFLS levels available within the set percentage schemes is shown by their coverage of the x axis.

communication infrastructure.

2. UFLS that tracks a set percentage of demand is predicated on the assumption that relays' net demand are proportional to the national demand. This assumption will break down in the future due to the large deployment of distributed generation and active demand (e.g. electric vehicles). Thus frequent communication will be required to maintain a set percentage of demand, which could be better utilised to provide fixed UFLS amounts.
3. The hourly relay updates required by this methodology are entirely feasible with modern communication equipment. Indeed, the communication requirements for the broad range of adaptive UFLS schemes proposed in the literature [81, 82, 83, 84, 85] are much more demanding, often requiring sub-second post fault communications between operator and relays. The proposed framework's local triggering of UFLS makes it robust against communication failures and delays.

4.6 Conclusion

The first major contribution of this chapter is to analytically derive a constraint on the frequency nadir from the swing equation. This constraint ensures that there is sufficient fast and slow FR,

inertia and UFLS scheduled to contain the largest system loss, which itself can be dynamically reduced to improve operation. We show that when the UFLS amount is discretized, the nonconvex constraint can be accurately and conservatively approximated by a convex second order cone. Case studies were run to explore how scheduled UFLS translates into a reduced requirement on other FS, translating into significantly enhanced integration of wind energy due to the lower response requirements from thermal plants. We demonstrate that drastic cost reductions are achievable by breaking the current paradigm of UFLS as a last resort containment measure. Furthermore the constraint offers insight into the synergies and conflicts between the diverse FS considered. A key take away is that UFLS remains valuable even in highly flexible systems.

Figure 4.9 shows that the nadir constraint is more conservative with increasing UFLS. This is acceptably small for the moderate UFLS amounts investigated here, but could significantly reduce optimality for future systems with abundant UFLS from smart devices. Methods to reduce conservativeness should be investigated, such as: 1) Piece-wise SOC approximation of the true constraint (4.12). 2) Searching within a convex inner approximation of the feasible space around the initial solution from (4.19) [107]. This process can be iterated around each subsequent optimal point until convergence occurs.

Finally, this paper only considers how UFLS can assist with frequency containment after the loss of the single largest plant. However, inertia and FR levels also impact the UFLS needed to contain unsecured events (e.g. the simultaneous disconnection of a nuclear plant and a CCGT). The presented framework could be used to explicitly value their role in containing larger losses, thus facilitating a more cost effective FS procurement.

The second major contribution of this chapter is a risk-constrained methodology for selecting feeders contributing to a given UFLS scheme. Convex formulations for a Gaussian chance constraint and distributionally-robust chance constraint have been introduced for this purpose, enabling the system operator to explicitly consider uncertainty in net-load of each feeder during the UFLS-allocation process. This method is applicable both for systems with and without a communication network in place: for the latter, due to the inability to update UFLS relays close to real-time, considering uncertainty in net-load allows to find the optimal combination of

feeders to be chosen for the UFLS scheme, without risking under-delivery; while for the former (i.e. WAMS-based adaptive UFLS schemes), the set-point for relays found through the proposed UFLS-allocation methodology can serve as a security layer to be used in case the communication network fails.

As future work, the impacts of UFLS on other types of stability should be analysed, notably voltage stability. Since DG are already providing reactive power services to the transmission system in some countries, an unintended consequence of the activation of UFLS could be a lack/excess of reactive power. Therefore, the UFLS-allocation algorithm should be enhanced to include voltage stability metrics, avoiding an inadequate reactive power flow in the post-UFLS stage. Moreover, analysing the effect of varying cost for different feeders within the UFLS-allocation optimisation should be studied, given that some loads might be willing to participate in UFLS contracts in exchange of a fixed fee, while other loads would be subject to the value-of-lost-load penalty if eventually disconnected.

Chapter 5

Conclusions and Future Work

This thesis focuses on developing novel mathematical frameworks to improve the scheduling of power plants in low-carbon systems. This is a problem of utmost importance due to the remarkable forecasted build rates of wind and solar over the coming years, and the reliance on electrification to decarbonise the transport sector. The key advancements made in this thesis focus on incorporating new uncertainties into the scheduling framework. This unlocks value from emerging resources like EVs, or creates new value from legacy assets like UFLS. Crucially, expanding the consideration of uncertainty allows explicit guarantees of system security and can result in less conservative formulations than deterministic methods.

The frameworks developed in chapters 3 and 4 take the form of MIP constraints whose continuous relaxations are convex and therefore easily solved by off-the-shelf commercial solvers. The second chapter establishes that MIP is currently a superior scheduling tool to nascent RL methods. These constraints can be applied within many essential power system optimisations for operations and markets. Extensive case studies of the future UK grid validated the constraints' usefulness and explored the sensitivity of the created value to relevant system characteristics such as RES or grid battery penetration. This work, therefore, addresses operational barriers that arise in RES-dominated power systems relating to their characteristic differences from fossil fuel synchronous plants, namely variability, uncertainty and stability (i.e. low inertia).

Nevertheless, there is still room for improvement that could further this work and better facilitate the integration of RES. These points complement and draw together the conclusions

within the individual chapters.

- *Scheduling with only partial knowledge of relevant probability distributions.*

In chapter two, the true distributions of the stochastic processes (demand forecast error and wind forecast error) are both assumed to be known. This enables the formation of a MIP with an accurate scenario tree that efficiently covers the possible net demand realisations. The decisions that minimise the expected cost over this tree will be close to optimal for the real system. The RL agent also exploits the true stochastic distributions during training within the environment. The agent takes an action based on the current state. When the environment translates this into a new state and reward, it must sample the stochastic variables. Consequently, if the environment reflects reality accurately, a well-trained agent will perform optimally when deployed.

In practice, the operator has only partial knowledge of the distributions of the salient stochastic processes. This may be because they do not conform to one, or there is insufficient data to define them accurately. Further investigation of how MIP and RL agents compare under only partial knowledge of the true distributions is needed. Within MIP the field of distributionally robust programming addresses this exact question by optimising for the worst possible distribution within an ambiguity set [108]. The distributionally robust chance constraints used in chapters 4 and 5 are appropriate for their single-stage applications. The multi-stage distributionally robust formulations required for forecast error uncertainty in the UC are challenging to implement and require further development [109]. Within RL, the literature on out-of-distribution generalisation environments is relevant. The work tries to find optimal agent policies when the training and testing environments differ [110], this is a promising field. Still, more work is needed as currently the vast majority of RL literature test and train from the same distributions, and no comparison to MIP ambiguity methods exists.

- *Forecasting*

Incorporating uncertainties into scheduling frameworks relies on accurately characterising the stochastic processes via forecasting. Methods for forecasting RES [111] are well

established as this is the dominant uncertainty in future power systems. However, the main contributions of this thesis are to incorporate other uncertainties into system operation, such as EV availability to provide ancillary services, or the net load flow through feeders that provide UFLS. The forecasting of these quantities is less developed [112], and improvement is needed.

Forecasts for the DR-CC on FR from EVs in chapter 3 must output a distribution of the power capabilities of the EV fleet during a scheduling time period. We present a simplistic method where the stochastic process is assumed stationary in time and well described by a similar time of the week from a previous year. More advanced methods can reduce the conservativeness of the constraint by reducing standard deviations or facilitating better ambiguity set selection. Currently, forecasting relating to EVs only focuses on their hourly energy requirements [113] which is insufficient for power-centric FR provision.

- *Implementation within markets*

The chance constraint limiting the scheduled fast FR from EVs in chapter 3 and the nadir constraint including UFLS in chapter 4 are applied to a centrally dispatched system. This is useful to verify the frameworks and inform on their system value. However, many systems worldwide (including the UK) are not centrally dispatched [114]. Instead, energy and ancillary services are acquired through a market. Simulations of a centrally dispatched system inform on an optimal baseline cost which well-structured markets can approach. However, there are some practical considerations before the frameworks can be applied within markets.

The chance constraint of chapter 3 translates a user's risk aversion level into an amount of FR to schedule from EVs. To find this amount the operator must at least know the mean and standard deviation of the EV connectivity forecast. More research is needed into how the fleets report these quantities and how best to penalise inaccurate submissions. Secondly, the found FR amount is a system aggregate, which needs to be disaggregated, initially to the fleets and then again to their individual EVs. Some work exists on the later disaggregation problem to individual EVs [115], but research to solve the prior fleet

disaggregation problem is still needed.

The nadir constraint of chapter 4 allows the co-optimisation of UFLS alongside fast and slow FR, inertia and dynamically reduced largest loss in securing a loss of generation. Availability and utilisation payments for the different frequency services are not considered. As such, considering UFLS as a service within a market context would allow the investigation of pricing schemes that reflect the appropriate value of different services. Simultaneously clearing energy and multiple ancillary services markets was shown to decrease costs significantly in reference [116]. However, the shadow pricing mechanism used requires continuous convex formulations. The discrete nature of UFLS could be incorporated by solving multiple formulations with different fixed UFLS amounts simultaneously, but more research is needed to find the best method.

Bibliography

- [1] C. O'Malley, L. Badesa, F. Teng, and G. Strbac, "Probabilistic Scheduling of UFLS to Secure Credible Contingencies in Low Inertia Systems," *IEEE Transactions on Power Systems*, vol. 37, no. 4, pp. 2693–2703, 2022. [Online]. Available: <https://arxiv.org/abs/2103.06616>
- [2] C. O'Malley, L. Badesa, F. Teng, and G. Strbac, "Frequency Response from Aggregated V2G Chargers with Uncertain EV Connections," *IEEE Transactions on Power Systems*, 2022. [Online]. Available: <https://arxiv.org/abs/2202.08026>
- [3] L. Badesa, C. O'Malley, M. Parajeles, and G. Strbac, "Chance-Constrained Allocation of UFLS Candidate Feeders Under High Penetration of Distributed Generation," *International Journal of Electrical Power & Energy Systems*, vol. 147, p. 108782, 5 2023. [Online]. Available: <https://arxiv.org/abs/2211.02953>
- [4] C. O'Malley, P. de Mars, L. Badesa, and G. Strbac, "Reinforcement Learning and Mixed-Integer Programming for Power Plant Scheduling in Low Carbon Systems: Comparison and Hybridisation," *Applied Energy, Under Review*, 12 2022. [Online]. Available: <http://arxiv.org/abs/2212.04824>
- [5] C. O'Malley, M. Aunedi, F. Teng, and G. Strbac, "Value of Fleet Vehicle to Grid in Providing Transmission System Operator Services," in *2020 Fifteenth International Conference on Ecological Vehicles and Renewable Energies*, 2020. [Online]. Available: <https://arxiv.org/abs/2201.05520>

- [6] F. Oldfield, K. Kumpavat, R. Corbett, A. Price, M. Aunedi, G. Strbac, C. O'Malley, D. Gardner, D. Pfeiffer, and J.-T. Kamphus, “*White Paper: The Drive Towards a Low-Carbon Grid. Unlocking the Value of Vehicle-to-Grid Fleets in Great Britain*,” Nissan, Imperial College London, Eon Drive, Tech. Rep., 2020.
- [7] H.-O. Pörtner, D. Roberts, M. Tignor, E. Poloczanska, K. Mintenbeck, M. C. A. Alegría, S. Langsdorf, S. Löschke, V. Möller, A. Okem, and B. Rama, “Climate Change 2022: Impacts, Adaptation, and Vulnerability,” *Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press*, 2022.
- [8] National Grid ESO, “Future energy scenarios 2022,” *Report*, 7 2022. [Online]. Available: <https://www.nationalgrideso.com/future-energy/future-energy-scenarios>
- [9] IEA, “Net Zero by 2050,” *IEA, Paris. License: CC BY 4.0*, 2021. [Online]. Available: <https://www.iea.org/reports/net-zero-by-2050>
- [10] IEA, “Renewables 2022,” *IEA, Paris. License: CC BY 4.0*, 2022. [Online]. Available: <https://www.iea.org/reports/renewables-2022>
- [11] J. Zhang, B. M. Hodge, E. Gomez-Lazaro, A. L. Lovholm, E. Berge, J. Miettinen, H. Holttinen, N. Cutululis, M. Litong-Palima, P. Sorensen, and J. Dobschinski, “Analysis of variability and uncertainty in wind power forecasting: An international comparison: Preprint,” *NREL Tech Rep*, 10 2013. [Online]. Available: <https://www.osti.gov/biblio/1107451>
- [12] E. Ela, V. Diakov, E. Ibanez, and M. Heaney, “Impacts of Variability and Uncertainty in Solar Photovoltaic Generation at Multiple Timescales,” *NREL Tech Rep*, 5 2013. [Online]. Available: <https://www.osti.gov/biblio/1081387>
- [13] N. Hatziargyriou, J. V. Milanovic, C. Rahmann, V. Ajjarapu, C. Canizares, I. Erlich, D. Hill, I. Hiskens, I. Kamwa, B. Pal, P. Pourbeik, J. J. Sanchez-Gasca, A. M. Stankovic, T. V. Cutsem, V. Vittal, and C. Vournas, “Definition and classification of power system

- stability - Revisited & Extended,” *IEEE Transactions on Power Systems*, pp. 1–1, 2020.
[Online]. Available: <https://ieeexplore.ieee.org/document/9286772/>
- [14] A. Papavasiliou, S. S. Oren, and B. Rountree, “Applying high performance computing to transmission-constrained stochastic unit commitment for renewable energy entegration,” *IEEE Transactions on Power Systems*, vol. 30, no. 3, pp. 1109–1120, 2014.
- [15] Y. Chen, F. Pan, F. Qiu, A. S. Xavier, T. Zheng, M. Marwali, B. Knueven, Y. Guan, P. B. Luh, L. Wu, B. Yan, M. A. Bragin, H. Zhong, A. Giacomoni, R. Baldick, B. Gisin, Q. Gu, R. Philbrick, and F. Li, “Security-constrained unit commitment for electricity market: Modeling, solution methods, and future challenges,” *IEEE Transactions on Power Systems*, pp. 1–14, 2022.
- [16] M. Håberg, “Fundamentals and recent developments in stochastic unit commitment,” *International Journal of Electrical Power and Energy Systems*, vol. 109, pp. 38–48, 2019.
- [17] P. A. Ruiz, C. R. Philbrick, E. Zak, K. W. Cheung, and P. W. Sauer, “Uncertainty management in the unit commitment problem,” *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 642–651, 2009.
- [18] S. Takriti, J. Birge, and E. Long, “A stochastic model for the unit commitment problem,” *IEEE Transactions on Power Systems*, vol. 11, no. 3, pp. 1497–1508, 1996.
- [19] A. Sturt, “Stochastic Scheduling of Wind-Integrated Power Systems,” Tech. Rep., 2011.
- [20] A. Sturt and G. Strbac, “Efficient Stochastic Scheduling for Simulation of Wind-Integrated Power Systems,” *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 323–334, feb 2012.
- [21] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [22] J. Qin, N. Yu, and Y. Gao, “Solving unit commitment problems with multi-step deep reinforcement learning,” in *2021 IEEE International Conference on Communications*,

- Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE, 2021, pp. 140–145.
- [23] P. de Mars and A. O’Sullivan, “Applying reinforcement learning and tree search to the unit commitment problem,” *Applied Energy*, vol. 302, Nov 2021.
- [24] Western Power Distribution, “Low Frequency Demand Disconnection.” *Report*, June 2017.
- [25] F. Teng, V. Trovato, and G. Strbac, “Stochastic Scheduling with Inertia-Dependent Fast Frequency Response Requirements,” *IEEE Transactions on Power Systems*, vol. 31, no. 2, pp. 1557–1566, mar 2016.
- [26] N. Miller, D. Lew, and R. Piwko, “Technology capabilities for fast frequency response,” *GE Energy Consulting, Tech. Rep*, vol. 3, pp. 10–60, 2017.
- [27] M. Aunedi and G. Strbac, “Whole-system cost of variable renewables in future gb electricity system,” *Imperial College, Tech. Rep*, 11 2016.
- [28] L. Badesa, F. Teng, and G. Strbac, “Simultaneous Scheduling of Multiple Frequency Services in Stochastic Unit Commitment,” *IEEE Transactions on Power Systems*, vol. 34, no. 5, pp. 3858–3868, sep 2019.
- [29] H. Chavez, R. Baldick, and S. Sharma, “Governor Rate-Constrained OPF for Primary Frequency Control Adequacy,” *IEEE Transactions on Power Systems*, vol. 29, no. 3, pp. 1473–1480, 2014.
- [30] F. Teng and G. Strbac, “Full Stochastic Scheduling for Low-Carbon Electricity Systems,” *IEEE Transactions on Automation Science and Engineering*, vol. 14, no. 2, pp. 461–470, apr 2017.
- [31] L. Badesa, F. Teng, and G. Strbac, “Optimal Portfolio of Distinct Frequency Response Services in Low-Inertia Systems,” *IEEE Transactions on Power Systems*, vol. 35, no. 6, pp. 4459–4469, 2020.
- [32] J. Luo, F. Teng, and S. Bu, “Stability-Constrained Power System Scheduling: A Review,” *IEEE Access*, 2020.

- [33] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown *et al.*, “Tackling climate change with machine learning,” *ACM Computing Surveys (CSUR)*, vol. 55, no. 2, pp. 1–96, 2022.
- [34] A. Kelly, A. O’Sullivan, P. de Mars, and A. Marot, “Reinforcement learning for electricity network operation,” *arXiv preprint arXiv:2003.07339*, 2020.
- [35] Á. S. Xavier, F. Qiu, and S. Ahmed, “Learning to Solve Large-Scale Security-Constrained Unit Commitment Problems,” *INFORMS Journal on Computing*, vol. 33, no. 2, pp. 739–756, 2021.
- [36] S. Pineda, J. M. Morales, and A. Jiménez-Cordero, “Data-driven screening of network constraints for unit commitment,” *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3695–3705, 2020.
- [37] V. Nair, S. Bartunov, F. Gimeno, I. von Glehn, P. Lichocki, I. Lobov, B. O’Donoghue, N. Sonnerat, C. Tjandraatmadja, P. Wang *et al.*, “Solving mixed integer programs using neural networks,” *arXiv preprint 2012.13349*, 2020.
- [38] S. Pineda and J. M. Morales, “Is learning for the unit commitment problem a low-hanging fruit?” *Electric Power Systems Research*, vol. 207, p. 107851, 2022.
- [39] E. Jasmin, T. I. Ahamed, and V. Jagthy Raj, “Reinforcement learning solution for unit commitment problem through pursuit method,” in *2009 International Conference on Advances in Computing, Control, and Telecommunication Technologies*, 2009, pp. 324–327.
- [40] E. Jasmin, T. I. Ahamed, and T. Remani, “A function approximation approach to reinforcement learning for solving unit commitment problem with photo voltaic sources,” in *2016 IEEE International Conference on Power Electronics, Drives and Energy Systems (PEDES)*, 2016, pp. 1–6.
- [41] N. K. Navin and R. Sharma, “A fuzzy reinforcement learning approach to thermal unit commitment problem,” *Neural Computing and Applications*, vol. 31, no. 3, pp. 737–750, 2019.

- [42] F. Li, J. Qin, and W. X. Zheng, “Distributed Q-learning-based online optimization algorithm for unit commitment and dispatch in smart grid,” *IEEE Transactions on Cybernetics*, vol. 50, no. 9, pp. 4146–4156, 2019.
- [43] P. de Mars and A. O’Sullivan, “Reinforcement learning and A* search for the unit commitment problem,” *Energy and AI*, vol. 9, p. 100179, 2022.
- [44] G. Ruan, H. Zhong, G. Zhang, Y. He, X. Wang, and T. Pu, “Review of learning-Assisted power system optimization,” *CSEE Journal of Power and Energy Systems*, vol. 7, no. 2, pp. 221–231, 2021.
- [45] S. Sridhar, J. Linderoth, and J. Luedtke, “Locally ideal formulations for piecewise linear functions with indicator variables,” *Operations Research Letters*, vol. 41, no. 6, pp. 627–632, 2013.
- [46] S. A. Kazarlis, A. Bakirtzis, and V. Petridis, “A genetic algorithm solution to the unit commitment problem,” *IEEE Transactions on Power Systems*, vol. 11, no. 1, pp. 83–92, 1996.
- [47] “National Grid Demand Data,” <https://www.nationalgrideso.com/data-explorer>.
- [48] “Balancing Mechanism Reporting Service,” <https://www.bmreports.com>.
- [49] G. Morales-España, J. M. Latorre, and A. Ramos, “Tight and compact MILP formulation of start-up and shut-down ramping in unit commitment,” *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1288–1296, 2013.
- [50] B. Knueven, J. Ostrowski, and J.-P. Watson, “On mixed-integer programming formulations for the unit commitment problem,” *INFORMS Journal on Computing*, vol. 32, pp. 857–876, 6 2020.
- [51] M. L. Bynum, G. A. Hackebeil, W. E. Hart, C. D. Laird, B. L. Nicholson, J. D. Sirola, J.-P. Watson, and D. L. Woodruff, *Pyomo—optimization modeling in Python*, 3rd ed. Springer Science & Business Media, 2021, vol. 67.

- [52] Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2022.
- [53] “National grid demand data.” [Online]. Available: <https://www.nationalgrideso.com/data-explorer>
- [54] “Balancing mechanism reporting service.” [Online]. Available: <https://www.bmreports.com>
- [55] P. de Mars, “Reinforcement learning and tree search methods for the unit commitment problem,” 2022. [Online]. Available: <https://arxiv.org/abs/2212.06001>
- [56] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint 1707.06347*, 2017.
- [57] G. Dulac-Arnold, R. Evans, H. van Hasselt, P. Sunehag, T. Lillicrap, J. Hunt, T. Mann, T. Weber, T. Degris, and B. Coppin, “Deep reinforcement learning in large discrete action spaces,” *arXiv preprint 1512.07679*, 2015.
- [58] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, “Mastering the game of Go without human knowledge,” *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [59] A. Fawzi, M. Balog, A. Huang, T. Hubert, B. Romera-Paredes, M. Barekatain, A. Novikov, F. J. R Ruiz, J. Schrittwieser, G. Swirszcz *et al.*, “Discovering faster matrix multiplication algorithms with reinforcement learning,” *Nature*, vol. 610, no. 7930, pp. 47–53, 2022.
- [60] Climate Change Committee, “The uk’s transition to electric vehicles,” 2020, [Online]. Available: <https://www.theccc.org.uk/wp-content/uploads/2020/12/The-UKs-transition-to-electric-vehicles.pdf>.
- [61] A. Thingvad, L. Calearo, P. B. Andersen, M. Marinelli, M. Neaimeh, K. Suzuki, and K. Murai, “Value of V2G frequency regulation in Great Britain considering real driving data,” in *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe 2019)*, 2019.

- [62] S. Gao, H. Li, J. Jurasz, and R. Dai, “Optimal Charging of Electric Vehicle Aggregations Participating in Energy and Ancillary Service Markets,” *IEEE Journal of Emerging and Selected Topics in Industrial Electronics*, vol. 3, no. 2, pp. 270–278, aug 2021.
- [63] A. Blatiak, F. Bellizio, L. Badesa, and G. Strbac, “Value of optimal trip and charging scheduling of commercial electric vehicle fleets with Vehicle-to-Grid in future low inertia systems,” *Sustainable Energy, Grids and Networks*, vol. 31, p. 100738, 2022.
- [64] A. Hajebrahimi, I. Kamwa, M. M. A. Abdelaziz, and A. Moeini, “Scenario-wise distributionally robust optimization for collaborative intermittent resources and electric vehicle aggregator bidding strategy,” *IEEE Transactions on Power Systems*, vol. 35, pp. 3706–3718, 2020.
- [65] X. Lu, K. W. Chan, S. Xia, X. Zhang, G. Wang, and F. Li, “A model to mitigate forecast uncertainties in distribution systems using the temporal flexibility of EVAs,” *IEEE Transactions on Power Systems*, vol. 35, pp. 2212–2221, 5 2020.
- [66] M. Amini and M. Almassalkhi, “Optimal Corrective Dispatch of Uncertain Virtual Energy Storage Systems,” *IEEE Transactions on Smart Grid*, vol. 11, no. 5, pp. 4155–4166, 2020.
- [67] National Grid ESO, “Reintroduction of aggregation at GSP Group for DC,” 2022, [Online]. Available: <https://www.nationalgrideso.com/electricity-transmission/document/234901/download>.
- [68] Y. Zhang, S. Shen, and J. L. Mathieu, “Distributionally robust chance-constrained optimal power flow with uncertain renewables and uncertain reserves provided by loads,” *IEEE Transactions on Power Systems*, vol. 32, pp. 1378–1388, 2017.
- [69] A. Bagchi, R. Best, and D. J. Morrow, “Investigating impacts of storage devices on distribution network aggregator’s day-ahead bidding strategy considering uncertainties,” *IEEE Access*, vol. 9, pp. 120 940–120 954, 2021.

- [70] L. Roald, F. Oldewurtel, B. Van Parys, and G. Andersson, “Security Constrained Optimal Power Flow with Distributionally Robust Chance Constraints,” pp. 1–8, 2015. [Online]. Available: <http://arxiv.org/abs/1508.06061>
- [71] W. Xie and S. Ahmed, “Distributionally robust chance constrained optimal power flow with renewables: A conic reformulation,” *IEEE Transactions on Power Systems*, vol. 33, pp. 1860–1867, 2018.
- [72] Y. Chen, Q. Guo, H. Sun, Z. Li, W. Wu, and Z. Li, “A distributionally robust optimization model for unit commitment based on kullback-leibler divergence,” *IEEE Transactions on Power Systems*, vol. 33, pp. 5147–5160, 9 2018.
- [73] A. Zhou, M. Yang, M. Wang, and Y. Zhang, “A linear programming approximation of distributionally robust chance-constrained dispatch with wasserstein distance,” *IEEE Transactions on Power Systems*, vol. 35, pp. 3366–3377, 9 2020.
- [74] P. Kundur, *Power System Stability and Control*, 1st ed. McGraw-Hill Education, 1994.
- [75] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge university press, 2004.
- [76] UK Department for Transport, “Electric chargepoint analysis 2017: Domestic,” Feb. 2018, doi: <https://www.gov.uk/government/statistics/electric-chargepoint-analysis-2017-domestics>.
- [77] A. Sturt and G. Strbac, “A times series model for the aggregate GB wind output circa 2030,” in *IET Conference on Renewable Power Generation (RPG 2011)*, 2011, pp. 1–6.
- [78] S. Pfenninger and I. Staffell, “Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data,” *Energy*, vol. 114, pp. 1251–1265, 2016.
- [79] “National Electricity Transmission System Security and Quality of Supply Standard,” The SQSS Review Panel, Tech. Rep. v2.4, April 2019.
- [80] International Energy Agency (IEA), “Renewables 2019: Market analysis and forecast from 2019 to 2024,” Available: <https://www.iea.org/reports/renewables-2019/distributed-solar-pv>, International Energy Agency (IEA), Tech. Rep., 2019.

- [81] H. Haes Alhelou, M. E. Hamedani Golshan, T. C. Njenda, and N. D. Hatziargyriou, “An Overview of UFLS in Conventional, Modern, and Future Smart Power Systems: Challenges and Opportunities,” *Electric Power Systems Research*, vol. 179, p. 106054, 2020.
- [82] M. Sanaye-Pasand and H. Seyedi, “New Centralised Adaptive Load-Shedding Algorithms to Mitigate Power System Blackouts,” *IET Generation, Transmission & Distribution*, vol. 3, no. 1, pp. 99–114, jan 2009.
- [83] V. V. Terzija, “Adaptive Underfrequency Load Shedding Based on the Magnitude of the Disturbance Estimation,” *IEEE Transactions on Power Systems*, vol. 21, no. 3, pp. 1260–1266, aug 2006.
- [84] U. Rudez and R. Mihalic, “WAMS-Based Underfrequency Load Shedding with Short-Term Frequency Prediction,” *IEEE Transactions on Power Delivery*, vol. 31, no. 4, pp. 1912–1920, aug 2016.
- [85] R. Hooshmand and M. Moazzami, “Optimal Design of Adaptive Under Frequency Load Shedding Using Artificial Neural Networks in Isolated Power System,” *International Journal of Electrical Power and Energy Systems*, vol. 42, no. 1, pp. 220–228, nov 2012.
- [86] S. Li, F. Tang, Y. Shao, and Q. Liao, “Adaptive under-frequency load shedding scheme in system integrated with high wind power penetration: Impacts and improvements,” *Energies*, vol. 10, no. 9, 2017.
- [87] H. Atighechi, P. Hu, S. Ebrahimi, J. Lu, G. Wang, and L. Wang, “An effective load shedding remedial action scheme considering wind farms generation,” *International Journal of Electrical Power and Energy Systems*, vol. 95, pp. 353–363, 2018.
- [88] A. Chandra and A. K. Pradhan, “An adaptive underfrequency load shedding scheme in the presence of solar photovoltaic plants,” *IEEE Systems Journal*, vol. 15, no. 1, pp. 1235–1244, 2021.
- [89] S. De Boeck and D. Van Hertem, “Integration of distributed PV in existing and future UFLS schemes,” *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 876–885, 2018.

- [90] K. Das, A. Nitsas, M. Altin, A. D. Hansen, and P. E. Sørensen, “Improved load-shedding scheme considering distributed generation,” *IEEE Transactions on Power Delivery*, vol. 32, no. 1, pp. 515–524, 2017.
- [91] B. Potel, F. Cadoux, L. De Alvaro Garcia, and V. Debusschere, “Clustering-based method for the feeder selection to improve the characteristics of load shedding,” *IET Smart Grid*, vol. 2, no. 4, pp. 659–668, 2019.
- [92] H. Ahmadi and H. Ghasemi, “Security-Constrained Unit Commitment with Linearized System Frequency Limit Constraints,” *IEEE Transactions on Power Systems*, vol. 29, no. 4, pp. 1536–1545, 2014.
- [93] Y. Wen, W. Li, G. Huang, and X. Liu, “Frequency Dynamics Constrained Unit Commitment with Battery Energy Storage,” *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 5115–5125, nov 2016.
- [94] M. Paturet, U. Markovic, S. Delikaraoglou, E. Vrettos, P. Aristidou, and G. Hug, “Stochastic Unit Commitment in Low-Inertia Grids,” *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3448–3458, sep 2020.
- [95] C. R. Rowland, D. W. Smaha, and J. W. Pope, “Coordination of load conservation with turbine-generator underfrequency protection,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-99, no. 3, pp. 1137–1150, 1980.
- [96] I. A. Tøndel, J. Foros, S. S. Kilskar, P. Hokstad, and M. G. Jaatun, “Interdependencies and reliability in the combined ICT and power system: An overview of current research,” *Applied Computing and Informatics*, vol. 14, no. 1, pp. 17–27, 2018.
- [97] J. Lofberg, “YALMIP: a toolbox for modeling and optimization in MATLAB,” in *2004 IEEE International Conference on Robotics and Automation*, 2004, pp. 284–289.
- [98] K. N. Hasan, R. Preece, and J. V. Milanović, “Existing approaches and trends in uncertainty modelling and probabilistic stability analysis of power systems with renewable generation,” *Renewable and Sustainable Energy Reviews*, vol. 101, pp. 168–180, 2019.

- [99] A. Soroudi, *Large Scale Renewable Power Generation*. Singapore: Green Energy and Technology, 2014, ch. Taxonomy of Uncertainty Modelling Techniques in Renewable Energy System Studies.
- [100] A. Zakaria, F. B. Ismail, M. H. Lipu, and M. Hannan, “Uncertainty models for stochastic optimization in renewable energy applications,” *Renewable Energy*, vol. 145, pp. 1543–1571, 2020.
- [101] X. Li, W. Wang, H. Wang, J. Wu, X. Fan, and Q. Xu, “Dynamic environmental economic dispatch of hybrid renewable energy systems based on tradable green certificates,” *Energy*, vol. 193, p. 116699, 2020.
- [102] J. Dong, M. M. Olama, T. Kuruganti, A. M. Melin, S. M. Djouadi, Y. Zhang, and Y. Xue, “Novel stochastic methods to predict short-term solar radiation and photovoltaic power,” *Renewable Energy*, vol. 145, pp. 333–346, 2020.
- [103] W. Xie and S. Ahmed, “Distributionally robust chance constrained optimal power flow with renewables: A conic reformulation,” *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 1860–1867, 2018.
- [104] S. Zhang, H. Cheng, K. Li, N. Tai, D. Wang, and F. Li, “Multi-objective distributed generation planning in distribution network considering correlations among uncertainties,” *Applied Energy*, vol. 226, pp. 743–755, 2018.
- [105] London Economics, “The Value of Lost Load (VoLL) for Electricity in Great Britain: Final report for OFGEM and DECC,” *OFGEM and DECC*, no. July, pp. 1–225, 2013.
- [106] “Study on the Estimation of the Value of Lost Load of Electricity Supply in Europe,” ACER, Tech. Rep., July 2018.
- [107] D. Lee, K. Turitsyn, D. K. Molzahn, and L. A. Roald, “Feasible Path Identification in Optimal Power Flow with Sequential Convex Restriction,” *IEEE Transactions on Power Systems*, vol. 35, no. 5, pp. 3648–3659, 2020.

- [108] H. Rahimian and S. Mehrotra, “Frameworks and results in distributionally robust optimization,” *Open Journal of Mathematical Optimization*, vol. 3, pp. 1–85, jul 2022. [Online]. Available: <https://doi.org/10.5802%2Fojmo.15>
- [109] R. Huang, S. Qu, X. Yang, and Z. Liu, “Multi-stage distributionally robust optimization with risk aversion,” *Journal of Industrial and Management Optimization*, vol. 17, no. 1, pp. 233–259, 2021. [Online]. Available: </article/id/e9b7e234-ad7e-4b39-92a8-69760a30027f>
- [110] R. Kirk, A. Zhang, E. Grefenstette, and T. Rocktäschel, “A survey of generalisation in deep reinforcement learning,” *CoRR*, vol. abs/2111.09794, 2021. [Online]. Available: <https://arxiv.org/abs/2111.09794>
- [111] T. Hong, P. Pinson, Y. Wang, R. Weron, D. Yang, and H. Zareipour, “Energy forecasting: A review and outlook,” *IEEE Open Access Journal of Power and Energy*, vol. 7, pp. 376–388, 2020.
- [112] F. Petropoulos et al, “Forecasting: Theory and practice,” *International Journal of Forecasting*, vol. 38, pp. 705–871, 7 2022, i think this wil be useful for my job.
- [113] L. Buzna, P. D. Falco, S. Khormali, D. Proto, and M. Straka, “Electric vehicle load forecasting: A comparison between time series and machine learning approaches,” *SyNERGY MED 2019 - 1st International Conference on Energy Transition in the Mediterranean Area*, 5 2019.
- [114] V. Ahlqvist, P. Holmberg, and T. Tangerås, “Central- versus self-dispatch in electricity markets,” *SSRN Electronic Journal*, 12 2018. [Online]. Available: <https://papers.ssrn.com/abstract=3302569>
- [115] F. L. Müller, J. Szabó, O. Sundström, and J. Lygeros, “Aggregation and disaggregation of energetic flexibility from distributed energy resources,” *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 1205–1214, 2019.
- [116] L. Badesa, F. Teng, and G. Strbac, “Pricing inertia and frequency response with diverse dynamics in a mixed-integer second-order cone programming formulation,” *Applied*

Energy, vol. 260, p. 114334, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306261919320215>