LESSON BOOK

Distributed Energy Resources as Grid Resources

From DER fundamentals to a worked virtual power plant

Eighteen lessons · 308 pages

Ricardo Marquez, Ph.D.

ricardo-marquez.com

Distributed Energy Resources as Grid Resources

From device physics to virtual power plants

Track B lessons, written up as a book

Ricardo Marquez

October 5, 2026 (work in progress: one chapter per lesson, added as the lessons proceed)

How to read this book

Each chapter of this book expands one lesson of "Track B", a sequence that follows the chain

DER fundamentals → demand response → DERMS → virtual power plants → thermostats, EVs and batteries → fleet aggregation → dispatch → dynamic load shaping → grid services.

The emphasis is on the *physical and operational logic*, not on vocabulary. Where the industry uses a shorthand term, a *Jargon* box expands it once and the text then uses the expanded meaning. Every symbol is defined in the notation chapter that follows and is used with the same meaning in every chapter.

Each chapter closes with a short section that connects the lesson to the TERRA framework, a graph-based way of modeling physical systems as conservation laws on nodes with flows on edges. That framework is a convenient mental model for everything in this book, because a distributed energy resource is, at bottom, a device that exchanges a conserved quantity (electrical energy, thermal energy) across a boundary under constraints.

Contents

How to read this book **ii**

Notation and vocabulary **viii**

1 What makes a DER a grid resource? **1**
1.1 Devices and resources . . . 1
1.2 The device as a dynamical system . . . 1
1.3 Flexibility as a quantity . . . 4
1.4 Constraints and uncertainty . . . 7
1.5 From one device to a fleet . . . 9
1.6 The working definition . . . 11
1.7 Exercises . . . 12
1.8 Connection to the TERRA framework . . . 12

2 Demand response and demand flexibility **14**
2.1 Demand response as a subset of demand flexibility . . . 14
2.2 The demand-response event . . . 14
2.3 The physics of an event . . . 16
2.4 Demand flexibility . . . 17
2.5 Objectives and decisions . . . 19
2.6 From events to continuous operation . . . 20
2.7 The key distinction . . . 22
2.8 Exercises . . . 23
2.9 Connection to the TERRA framework . . . 23

3 What a DERMS does **25**
3.1 Definition and the operational chain . . . 25
3.2 Sensing the fleet . . . 25
3.3 Aggregation . . . 28
3.4 Decision and action . . . 30
3.5 The DERMS as a control system . . . 31
3.6 Architecture and neighbors . . . 32
3.7 The working definition . . . 34
3.8 Exercises . . . 34
3.9 Connection to the TERRA framework . . . 35

4 Virtual power plants **37**
4.1 Many devices and one resource . . . 37
4.2 Capacity . . . 37
4.3 The service product . . . 39
4.4 The VPP as a resource abstraction . . . 42

4.5 Commitment under uncertainty . . . 43
4.6 Operation and forecasting . . . 45
4.7 The working definition . . . 46
4.8 Exercises . . . 47
4.9 Connection to the TERRA framework . . . 47

5 Thermostatically controlled loads 49
5.1 The question . . . 49
5.2 The building and the thermostat . . . 49
5.3 What control does . . . 52
5.4 The state-space model . . . 54
5.5 Fleet behavior . . . 55
5.6 Forecasting and dispatch . . . 57
5.7 What to retain . . . 59
5.8 Exercises . . . 59
5.9 Connection to the TERRA framework . . . 60

6 Batteries as DERs 61
6.1 The question . . . 61
6.2 State and dynamics . . . 61
6.3 Constraints . . . 62
6.4 Dispatch . . . 65
6.5 Fleet and forecasting . . . 66
6.6 Comparison and control formulation . . . 69
6.7 What to retain . . . 70
6.8 Exercises . . . 71
6.9 Connection to the TERRA framework . . . 71

7 EV managed charging 72
7.1 The question . . . 72
7.2 One vehicle's night . . . 72
7.3 The model . . . 73
7.4 Scheduling . . . 74
7.5 Direction, forecasting and risk . . . 76
7.6 Vehicle-to-grid and the fleet . . . 78
7.7 What to retain . . . 81
7.8 Exercises . . . 81
7.9 Connection to the TERRA framework . . . 81

8 Baselines and verification 83
8.1 The counterfactual problem . . . 83
8.2 Baseline methods . . . 86
8.3 Verification metrics . . . 89
8.4 Attributing underperformance . . . 91
8.5 Probabilistic and experimental verification . . . 93
8.6 Verification inside the learning loop . . . 96
8.7 What to retain . . . 98
8.8 Exercises . . . 98
8.9 Connection to the TERRA framework . . . 99

9 Fleet aggregation **101**
9.1 The aggregation problem . . . 101
9.2 Representations of a fleet . . . 103
9.3 Diversity and common-mode uncertainty . . . 106
9.4 Location, service and hierarchy . . . 109
9.5 Probabilistic aggregation . . . 111
9.6 The mathematical object . . . 114
9.7 What to retain . . . 116
9.8 Exercises . . . 117
9.9 Connection to the TERRA framework . . . 117

10 Forecasting flexibility **119**
10.1 Load forecasting and flexibility forecasting . . . 119
10.2 Resource-specific forecasts . . . 122
10.3 Probabilistic forecast targets . . . 124
10.4 Features and realizable flexibility . . . 128
10.5 Endogeneity of the forecast target . . . 129
10.6 The forecast target and the interface to optimization . . . 132
10.7 What to retain . . . 133
10.8 Exercises . . . 133
10.9 Connection to the TERRA framework . . . 135

11 Dispatch and optimization **137**
11.1 The dispatch problem . . . 137
11.2 Costs and constraints . . . 139
11.3 Horizon and model predictive control . . . 142
11.4 Dispatch under uncertainty . . . 144
11.5 Coordinating heterogeneous resources . . . 148
11.6 Economics and the value of a forecast . . . 151
11.7 What to retain . . . 153
11.8 Exercises . . . 153
11.9 Connection to the TERRA framework . . . 154

12 Dynamic load shaping **156**
12.1 From events to trajectories . . . 156
12.2 Cohort rotation, rebound and pre-conditioning . . . 159
12.3 Shaping as tracking control . . . 162
12.4 Shapes and services . . . 164
12.5 Feasibility over the horizon and feedback . . . 167
12.6 Objective, trade-offs and controlled recovery . . . 169
12.7 What to retain . . . 172
12.8 Exercises . . . 172
12.9 Connection to the TERRA framework . . . 173

13 Distribution-grid constraints **175**
13.1 System-wide and local flexibility . . . 175
13.2 Equipment and voltage limits . . . 178
13.3 Network-aware dispatch . . . 182
13.4 Location in aggregation and in the merit order . . . 185
13.5 A worked example with uncertainty . . . 187
13.6 Topology, the ADMS, and planning . . . 190
13.7 What to retain . . . 192

13.8 Exercises . . . 193
13.9 Connection to the TERRA framework . . . 193

14 Grid services and markets **195**
14.1 Products, not megawatts . . . 195
14.2 Wholesale products . . . 197
14.3 Demand response, load shaping, and locational services . . . 200
14.4 Value stacking and co-optimization . . . 203
14.5 Requirements differ by service . . . 206
14.6 Commitment economics . . . 208
14.7 What to retain . . . 211
14.8 Exercises . . . 212
14.9 Connection to the TERRA framework . . . 212

15 Customer and device constraints **214**
15.1 Purpose and the intersection of constraints . . . 214
15.2 Customer constraints . . . 216
15.3 Device constraints . . . 220
15.4 Program rules and the economics of an event . . . 222
15.5 Heterogeneity, fairness and participation . . . 224
15.6 The feasible resource . . . 227
15.7 What to retain . . . 230
15.8 Exercises . . . 230
15.9 Connection to the TERRA framework . . . 231

16 DERMS and VPP architecture **232**
16.1 The layered architecture . . . 232
16.2 Sensing: device integration, telemetry and state . . . 235
16.3 Forecasting and aggregation layers . . . 236
16.4 Decision, actuation and monitoring . . . 238
16.5 Verification, learning and cadence . . . 240
16.6 Interfaces, robustness and uncertainty . . . 242
16.7 What to retain . . . 247
16.8 Exercises . . . 248
16.9 Connection to the TERRA framework . . . 248

17 A commercial platform: the architecture in practice **250**
17.1 The platform and its value chain . . . 250
17.2 Telemetry and DER-specific forecasting . . . 252
17.3 From demand response to dynamic load shaping . . . 255
17.4 Objectives, targets and resource classes . . . 258
17.5 Scale and the platform architecture . . . 260
17.6 The forecasting and optimization role . . . 261
17.7 What to retain . . . 266
17.8 Exercises . . . 266
17.9 Connection to the TERRA framework . . . 268

18 A worked case: design and operate a VPP **269**
18.1 The case and the service definition . . . 269
18.2 Resource models . . . 271
18.3 Finding a feasible plan . . . 274
18.4 Uncertainty, reserves and real-time operation . . . 278

18.5 Attribution and verification 283
18.6 Presenting the reasoning 285
18.7 What to retain 289
18.8 Exercises 290
18.9 Connection to the TERRA framework 291

Notation and vocabulary

The same symbols are used throughout. Power is in kilowatts (kW) or megawatts (MW); energy is in kilowatt-hours (kWh); temperature is in degrees Celsius; time t is a clock time or a discrete step index, and τ is a duration.

Symbols

Symbol	Unit	Meaning
$P_{\text{grid}}(t)$	kW	Net power drawn from the grid at the point where the premises meet the network. Positive means import, negative means export.
$P_{\text{load}}(t)$	kW	Ordinary, uncontrolled consumption of the premises.
$P_{\text{EV}}(t)$	kW	Electric-vehicle charging power (positive when charging).
$P_{\text{bat}}(t)$	kW	Battery power; positive when charging, negative when discharging.
$P_{\text{PV}}(t)$	kW	Solar photovoltaic generation (always ≥ 0).
$P_{\text{AC}}(t)$	kW	Electrical power drawn by an air conditioner or heat pump.
x_t	–	Internal *state* of a device at step t (what the device "remembers").
u_t	–	*Control action* applied at step t (what we are allowed to choose).
w_t	–	*Disturbance*: an external condition we do not choose (weather, occupancy, driving).
$\mathcal{U}(x_t)$	–	The set of control actions that are allowed when the state is x_t.
$f(x, u, w)$	–	State-transition map: $x_{t+1} = f(x_t, u_t, w_t)$.
$g(x, u, w)$	kW	Output map: $P_{\text{grid}} = g(x, u, w)$.
$P^{\max}$	kW	Nameplate (rated) power of a device.
$P^{\text{avail}}(x)$	kW	Power actually available right now, given the state.
E^{avail}	kWh	Energy available to be delivered (or absorbed) before a limit is hit.
P_t^{base}	kW	Baseline: the power the premises *would have* drawn without control.
P_t^{ctrl}	kW	Power actually drawn under control.
$\Delta P_t = P_t^{\text{base}} - P_t^{\text{ctrl}}$	kW	Load reduction delivered by control at step t.
$F(t, \tau)$	kW	Flexibility: the sustainable load change if a response starts at t and must last τ.

Symbol	Unit	Meaning
SOC	% or –	State of charge of a battery (fraction of usable energy stored).
$\mathrm{SOC}^{\mathrm{res}}$	%	Reserve the owner insists on keeping (backup, warranty).
$\mathrm{SOC}^{\mathrm{req}}$	%	State of charge the EV driver requires at departure.
t_{dep}	time	Departure time of an EV.
$T_{\mathrm{in}}, T_{\mathrm{out}}, T_{\mathrm{set}}$	°C	Indoor, outdoor and setpoint temperatures.
$T^{\min}, T^{\max}$	°C	Customer comfort band on T_{in}.
$P^{\mathrm{act}}(t)$	kW, MW	Actual, metered load during an event.
$P^{\mathrm{tgt}}(t)$	MW	Target trajectory the fleet is asked to follow.
$P^{\mathrm{del}}(t)$	MW	Load change actually delivered by the fleet.
$P_{\mathrm{fleet}}(t)$	MW	Net power of the whole fleet.
$\hat{P}$	MW	A forecast of a power quantity (hat denotes an estimate).
$E^{\max}$	kWh	Usable energy capacity of a battery.
η	–	Charging efficiency (fraction of drawn energy that is stored).
c_t	\$/kWh	Price (or marginal cost) of energy at step t.
J	–	Objective function minimized by a dispatch optimizer.
e_t	MW	Tracking error, $P^{\mathrm{tgt}} - P^{\mathrm{del}}$, at step t.
H	steps	Optimization horizon (number of steps looked ahead).
y_t	–	Telemetry: what the platform *observes* at step t.
$h(x)$	–	Measurement map, $y_t = h(x_t) + v_t$.
v_t	–	Measurement noise.
$F_i(t)$	kW	Flexibility of device i; $F(r, t, d)$ indexes by region r and type d.
ρ	–	Correlation coefficient between device flexibilities.
a_i	–	Availability factor of device i, between 0 and 1.
P^{name}	MW	Nameplate (installed) capacity of a fleet, $\sum_i P_i^{\max}$.
P^{commit}	MW	Capacity committed to a utility or market.
P^{enroll}	MW	Nominal flexibility of everything enrolled.
$P_{\mathrm{VPP}}(t)$	MW	Power of the virtual power plant as seen by the grid.
$C(\tau)$	MW	Deliverable capacity for a required duration τ.
p^{ovr}	–	Probability that a customer overrides an event.
$z = \phi(x)$	–	Reduced representation of the fleet state handed to the grid.
R, C	°C/kW, kWh/°C	Thermal resistance and capacitance of a building; RC is its time constant.
Q_{int}, Q_{HVAC}	kW	Internal heat gains; heat removed by the HVAC system.
s_t	–	Compressor state at step t, on (1) or off (0).
D	–	Duty cycle: fraction of time the compressor is on.
$\mathrm{SOC}^{\mathrm{th}}$	–	Thermal state of charge: position of T_{in} inside the comfort band.
E^{shift}	kWh	Energy shifted by an event, $\int \Delta P(t)\,\mathrm{d}t$.

Symbol	Unit	Meaning
η_{c}, η_{d}, η_{RT}	–	Charging, discharging and round-trip efficiency, $\eta_{\mathrm{RT}} = \eta_{\mathrm{c}}\,\eta_{\mathrm{d}}$.
$\mathrm{SOC}^{\min}$, $\mathrm{SOC}^{\max}$	–	Operating limits on state of charge (protection, degradation, reserve).
P^{ch}, P^{dis}	kW	Charging and discharging power limits (magnitudes).
$F^{\uparrow}$, $F^{\downarrow}$	kW	Upward grid support (reduce net load / inject) and downward support (increase net load).
a_t	–	Plug-in indicator of an EV at step t: 1 if connected, 0 if away.
T^{arr}, T^{dep}	time	Arrival and departure times of an EV (random in general).
E^{need}	kWh	Energy an EV must receive before departure, $(\mathrm{SOC}^{\mathrm{req}} - \mathrm{SOC}_0)\,E^{\max}$.
E^{rem}, T^{rem}	kWh, h	Energy still to deliver and time left before departure.
U_i	–	Urgency of EV i: charging time needed divided by time left.
β_k	–	Coefficients of a regression baseline or forecast model.
RMSE	MW	Root-mean-square tracking error over an event.
ATE	MW	Average treatment effect: control-group load minus treated-group load.
R (response)	MW	Verified response of the fleet; distinguished from forecast F and request P^{tgt}.
n_k	–	Number of devices in state bin k of a histogram representation.
$\mathcal{F}$	MW	A feasible set of fleet power (instantaneous) or of fleet trajectories.
$\mathbf{H}$	–	Sensitivity matrix from DER power changes to network quantities, $\Delta z = \mathbf{H}\,\Delta P$.
S, $S^{\max}$	kVA	Apparent power through a transformer and its rating.
V_i, I_ℓ	V, A	Voltage at node i; current on branch ℓ.
Q	kvar	Reactive power.
r, ℓ	–	Electrical region or location index.
α	–	Reliability level of a commitment or forecast quantile (also a derating factor in dispatch).
π_C, π_P, π_t	\$/MW, \$/MWh	Price of committed capacity, penalty for shortfall, energy price at step t.
λ_k	–	Weight of cost term k in a dispatch objective.
N^{events}	–	Remaining event budget of a demand-response program.
r_i	–	Learned reliability score of device i.

Vocabulary

DER (distributed energy resource)
A physical device connected on the customer side of the distribution network that

generates, stores or consumes electricity and can, under some conditions, change how much it exchanges with the grid. Batteries, electric vehicles, thermostats with the heating/cooling plant behind them, rooftop solar, electric water heaters and controllable commercial loads are all DERs.

DERMS (DER management system)
The software layer that monitors DERs, forecasts what they can do, decides what to ask of them and sends the commands. Chapter 3 of the lesson sequence.

VPP (virtual power plant)
A fleet of DERs operated together so that, from the grid's point of view, it behaves like a single controllable plant. Chapter 4.

Aggregator
The company that enrolls customers' devices, runs the DERMS and sells the fleet's capability to a utility or market. In the lessons the aggregator is one large residential DER aggregator, which is not named.

Demand response (DR)
Asking customers' loads to reduce (or shift) consumption for a period, usually at a time of system stress. Chapter 2.

Flexibility
The ability to change net power exchange with the grid on request, while respecting the device's own limits and the customer's requirements. The central quantity of this book.

Baseline
The counterfactual consumption that would have happened without the control request. Needed to measure what a response delivered.

Dispatch
The act of deciding and sending control actions to a set of devices for a period.

Setpoint
The target value a controller tries to hold, for example the thermostat temperature.

State of charge (SOC)
The fraction of a battery's usable energy currently stored.

Nameplate rating
The maximum power printed on the device label. Rarely the amount actually available.

Point of common coupling (PCC)
The physical point where the customer's premises connect to the utility network; P_{grid} is measured there.

Pre-cooling / pre-heating
Deliberately moving a building's temperature before an event so that more flexibility is available during the event.

Event
A bounded period during which an aggregator asks enrolled devices to change their consumption; the unit of traditional demand response.

Rebound
The increase in load above baseline after an event ends, as devices recover the state they gave up during it.

Load shifting
Moving consumption from one period to another with the same total energy delivered.

Load shedding
Turning load off, sometimes involuntarily, to protect the system.

Demand-side management (DSM)
The historical umbrella for everything done on the consumption side: efficiency, tariffs, conservation and demand response.

Measurement and verification (M&V)
The procedure that estimates what an event delivered, by comparing metered load with a baseline.

Chance constraint
A constraint required to hold with a stated probability, for example $\Pr(P^{\text{del}} \geq 100\,\text{MW}) \geq 0.95$.

Telemetry
The stream of measurements a device reports to the platform (temperatures, state of charge, power, status).

API (application programming interface)
The documented way one piece of software talks to another; each device vendor exposes its own.

State estimation
Inferring the full internal state of a system from partial, noisy measurements.

MPC (model predictive control)
A control scheme that repeatedly estimates the state, optimizes actions over a finite horizon, applies the first action, observes, and repeats.

EMS (energy management system)
The control-center software of the bulk transmission grid: SCADA, state estimation, contingency analysis, economic dispatch, automatic generation control.

ADMS (advanced distribution management system)
The utility's software for operating the distribution network itself: feeder topology, switching, voltage, outage management.

SCADA (supervisory control and data acquisition)
The measurement-and-command infrastructure of a control center.

Volt/VAR optimization
Coordinated control of voltage and reactive power on a distribution feeder; an ADMS function.

Nameplate, available, firm capacity
Three successively smaller numbers: what is installed; what the current states permit; what can be promised with a stated confidence.

Derating
Reducing a resource's credited capacity below nameplate to account for unavailability; standard for conventional generators.

Coverage ratio
Enrolled nominal flexibility divided by committed capacity; the overbooking margin.

CVaR (conditional value at risk)
The expected shortfall in the worst fraction of outcomes, for example the mean delivery in the worst 5 % of days; a tail-risk measure.

Model reduction
Replacing a high-dimensional state by a small set of quantities that suffice for a decision.

Deadband (hysteresis)
The temperature band around the setpoint inside which a thermostat leaves the compressor in its current state.

Duty cycle
Fraction of time a cycling device is on; its average power is the duty cycle times its on-power.

Thermal time constant
The product RC of a building; the time for the indoor temperature to cover about 63 % of a step change in the outdoor temperature with the HVAC off.

Synchronization
The loss of diversity that occurs when many cycling devices receive the same command at the same time and thereafter switch together.

Staggered recovery
Ending an event for groups of devices at different times to avoid a synchronized rebound.

BESS (battery energy storage system)
A stationary battery with its inverter and controls; residential (a few kWh to tens of kWh) or utility scale.

Round-trip efficiency
Fraction of the energy put into a battery that comes back out; the product of charging and discharging efficiencies.

Degradation
Loss of battery capacity and life with use; cycling has a cost even when energy is free.

Arbitrage
Buying (charging) when the price is low and selling (discharging) when it is high.

Managed charging
Scheduling when an EV charges, within its plugged-in window, so that it is ready by departure.

V2G (vehicle-to-grid)
Discharging an EV battery into the grid; bidirectional charging.

Slack Time plugged in minus time needed to charge at full power; the scheduling freedom of one EV.

Urgency
Charging time needed divided by time left before departure; above one the deadline cannot be met.

Earliest-deadline-first (EDF)
A scheduling rule that serves the job with the nearest deadline first.

Valley filling
Placing flexible load in the hours where system load is lowest.

Counterfactual
What would have happened without the intervention; never observed, always inferred.

Control group
Enrolled devices deliberately left uncontrolled during an event so their load estimates the counterfactual.

Histogram state
A fleet representation that counts devices per state bin instead of listing every device.

Virtual battery
A fleet of flexible loads approximated by one power limit, one energy state and its bounds.

Common-mode uncertainty
Error shared by every device in a fleet (for example a weather forecast error); aggregation does not average it away.

Quantile forecast
A forecast reported as values exceeded with stated probabilities rather than as one number.

Scenario optimization
Optimizing expected cost over a set of sampled futures.

Robust optimization
Optimizing against the worst case inside an uncertainty set.

Receding horizon
Re-solving a finite-horizon plan at every step and applying only its first action; the mechanism of MPC.

Reference signal
The trajectory a controller is asked to track.

Feeder
A distribution circuit leaving a substation and serving a group of transformers and customers.

Hosting capacity
How much additional DER a network can accommodate before a constraint is violated.

Optimal power flow (OPF)
Dispatching resources to minimize a cost subject to the network's power-flow equations and limits.

Regulation
A fast balancing service that follows a frequently updated signal around a baseline.

Operating reserve
Capacity held back, unused, to respond to an unexpected event.

Value stacking
: Using one device for several services at different times; forbidden to count the same capability twice for services that can call at once.

Event budget
: The number of events a program allows per season; a finite resource to allocate over time.

Vendor adapter
: The software that translates between a device vendor's interface and the platform's common model.

Cadence
: How often a module of the system runs.

Edge DERMS
: A DERMS operated by an aggregator on customer-owned devices at the grid edge, as opposed to a utility-owned system.

Enrollment
: The customer-program step of signing a customer and connecting their device; a VPP's capacity is bounded by it.

Snapback
: An operator's term for the post-event rebound.

Firm multi-hour dispatch
: Grouping thermostats and scheduling the groups so the fleet's multi-hour reduction is approximately flat.

Packetized energy management
: A coordination technique in which devices request short "packets" of energy that a coordinator grants or defers.

Receding horizon
: See model predictive control.

Chapter 1

What is a DER, and what makes it a grid resource?

1.1 Devices and resources

The first distinction of the subject is simple but is often lost.

Key idea

A *distributed energy resource* (DER) is a physical device. It is not useful to the grid merely because it exists. It becomes useful when it can *change its net power exchange with the grid when asked.*

Batteries, electric vehicles (EVs), thermostats with the heating and cooling plant behind them, rooftop solar photovoltaics (PV), electric water heaters and controllable commercial loads are all DERs. Each one sits on the customer's side of the meter, each one moves electrical energy, and each one was bought for a reason that has nothing to do with the grid: the owner wants a comfortable house, a charged car, backup power, or a lower bill.

For an aggregator the only question that matters at first is:

Can this device change its net power exchange with the grid when requested?

The rest of this chapter makes precise what "can" means in that sentence. It depends on the device's internal state, on how long the change must last, on what the customer will tolerate, and on quantities that are not known in advance. Those four dependencies, state, time, customer, and uncertainty, organize the chapter.

Jargon: DER, aggregator, flexible capacity

DER is any customer-side device that generates, stores or consumes electricity. An *aggregator* is the company that enrolls many such devices and operates them together (a commercial aggregation platform in these lessons). *Flexible capacity* is the amount of power the aggregator can add to or remove from the grid on request. This chapter makes "flexible capacity" precise.

1.2 The device as a dynamical system

This section sets up the vocabulary used for every device in the book: a power balance at the meter, and a state, a control and a disturbance behind each term.

1.2.1 Power exchange at the point of common coupling

Consider one house connected to the grid, shown in Figure 1.1. Define $P_{\text{grid}}(t)$ as the instantaneous power drawn from the grid at the point of common coupling, the physical point where the house meets the network. Positive values mean the house is importing, negative values mean it is exporting.

A simple balance at that point is

$$P_{\text{grid}}(t) = P_{\text{load}}(t) + P_{\text{EV}}(t) + P_{\text{bat}}(t) - P_{\text{PV}}(t). \tag{1.1}$$

Here P_{load} is ordinary household consumption (lights, cooking, appliances), P_{EV} is EV charging power, P_{bat} is battery power with the convention that $P_{\text{bat}} > 0$ means charging and $P_{\text{bat}} < 0$ means discharging, and P_{PV} is solar generation, which always enters with a minus sign because it offsets import.

Equation (1.1) is conservation of energy at one node: whatever the house consumes or stores must come from the grid or from its own solar generation. The same statement holds at every bus of a power network, and Section 1.8 uses it to connect the chapter to the TERRA framework.

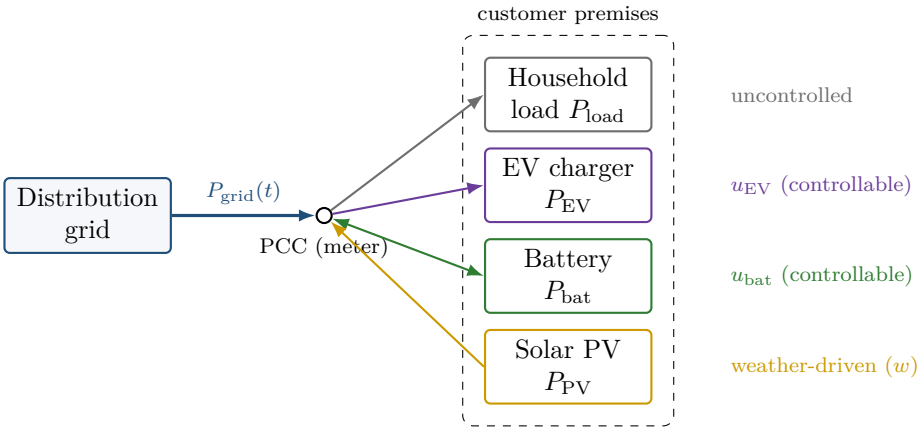


Figure 1.1: One house as a power-balance node. The net exchange P_{grid} at the point of common coupling (PCC) is the sum of the device powers, equation (1.1). Only some of the terms are controllable.

Now suppose some of the terms can be commanded. A control signal u_{EV} might set EV charging power; u_{bat} might set battery charging or discharging. Then the net exchange can be written schematically as

$$P_{\text{grid}} = g(x, u, w), \tag{1.2}$$

with three kinds of argument:

- x, the **internal device states**: what the devices "remember" from the past (how full the battery is, how warm the house is);
- u, the **control actions**: what the aggregator is allowed to choose;
- w, the **external disturbances or conditions**: what nobody chooses (outdoor temperature, sunshine, whether the car is home).

This triple (x, u, w) is the standard language of control theory, and it is worth adopting from the start because every later lesson (DERMS, dispatch, load shaping) is a statement about how to pick u given x and a forecast of w. Table 1.1 lists the three arguments for the common device types.

1.2.2 State dependence of available power

At this point a DER starts to differ from a conventional generator. Consider the question

Table 1.1: The state, control and disturbance of the common DER types.

Device	State x	Control u	Disturbance w
Battery	SOC; temperature; cycle count	charge/discharge power P_{bat}	household load P_{load}, PV output
Electric vehicle	SOC; plugged-in or not; departure time t_{dep}	charging power P_{EV}	arrival time, trips, driver overrides
Thermostat / HVAC	indoor temperature T_{in}; compressor on/off; current setpoint T_{set}	setpoint T_{set} (or direct on/off)	outdoor temperature T_{out}, solar gain, occupancy
Water heater	tank temperature; stratification	element on/off	hot-water draws
Solar PV	(none controllable in most installs)	curtailment, if any	irradiance, clouds

Can this battery give me 5 kW right now?

The answer depends on the state. If the state of charge is SOC = 90% the answer is yes. If SOC = 8% and the owner insists on keeping a 10% emergency reserve, the answer is no, even though the label on the inverter says 5 kW.

So the available power is not simply the nameplate rating, $P^{\max} = 5\,\text{kW}$. It is a function of the state,

$$P^{\text{avail}} = P^{\text{avail}}(x). \tag{1.3}$$

For the battery in the example a reasonable form is

$$P^{\text{avail}}(\text{SOC}) = \begin{cases} P^{\max}, & \text{SOC} \geq \text{SOC}^{\text{res}} + \delta, \\ P^{\max} \cdot \dfrac{\text{SOC} - \text{SOC}^{\text{res}}}{\delta}, & \text{SOC}^{\text{res}} \leq \text{SOC} < \text{SOC}^{\text{res}} + \delta, \\ 0, & \text{SOC} < \text{SOC}^{\text{res}}, \end{cases} \tag{1.4}$$

where δ is a small band over which the inverter tapers power as it approaches the reserve. Figure 1.2 draws this. The exact taper is a detail; the point is the shape: a flat region where the device is fully available, a reserve region where it is not available at all, and a transition between them.

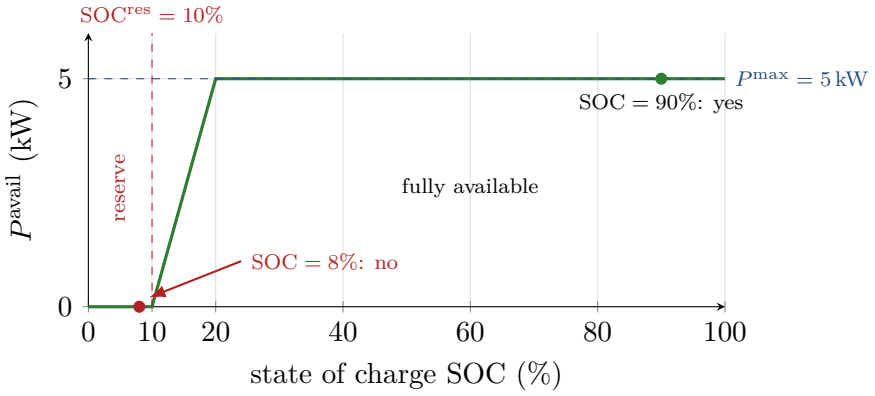


Figure 1.2: Available discharge power of a 5 kW battery as a function of state of charge, equation (1.4), with a 10 % owner reserve and a 10 % taper band. The same device answers "yes" or "no" to the same request depending only on its state.

The state that matters is different for each device type:

$$\begin{aligned} \text{battery:} &\quad x \supset \{\text{SOC}\}, \\ \text{EV:} &\quad x \supset \{\text{SOC},\ \text{plugged-in status},\ t_{\text{dep}}\}, \\ \text{thermostat:} &\quad x \supset \{T_{\text{in}},\ \text{HVAC on/off},\ T_{\text{set}}\}. \end{aligned}$$

The symbol $\supset$ reads "contains at least"; a real model may carry more (battery temperature, EV charger rating, building thermal mass), but these are the states without which no sensible answer can be given.

Principle 1.1. *DER flexibility is* state-dependent*: the power a device can deliver on request is a function $P^{\text{avail}}(x)$ of its internal state, not a fixed nameplate number.*

1.3 Flexibility as a quantity

With the vocabulary in place, flexibility can be defined as a quantity. Four properties follow in turn: it is the change caused by control, it has a duration, it is bounded by power and by energy, and it is a state-dependent feasible set.

1.3.1 Controllability versus flexibility

Suppose a house has a smart thermostat. Technically the aggregator can change the setpoint,

$$u_t = T_{\text{set}}.$$

Now imagine the house is already warm, the thermostat is satisfied, and the compressor is off. Raising the setpoint by another degree produces essentially no change in power:

$$\Delta P \approx 0.$$

The device is fully controllable and yet, at that moment, provides no flexibility. Controllability is a property of the interface; flexibility is a property of the physics at the current operating point.

A useful conceptual definition is

Flexibility = the change in feasible grid power caused by allowable control.

or, as a quantity at a single time step,

$$\Delta P_t = P_t^{\text{base}} - P_t^{\text{ctrl}}, \tag{1.5}$$

where P_t^{base} is the *baseline*, the power the premises would have drawn with no control, and P_t^{ctrl} is the power actually drawn under control. With the sign in (1.5), a positive ΔP_t is a load *reduction*, which is what a grid under stress usually asks for.

Example 1.1 (One thermostat at 5 PM)**.** If the baseline would have been $P_t^{\text{base}} = 3.5\,\text{kW}$ and control reduces consumption to $P_t^{\text{ctrl}} = 1.0\,\text{kW}$, then

$$\Delta P_t = 3.5 - 1.0 = 2.5\,\text{kW}.$$

That 2.5 kW is the quantity a virtual power plant is interested in, not the 3.5 kW of load and not the air conditioner's 4 kW nameplate.

Jargon: baseline

The baseline is a *counterfactual*: what would have happened without the request. It cannot be measured, only estimated, because once control is applied the uncontrolled outcome can no longer be observed. Baseline estimation is one of the oldest arguments in demand response and gets its own treatment in Chapter 2.

1.3.2 Duration dependence of flexibility

A second property separates a DER from a generator. Suppose the thermostat above gives 2.5 kW of load reduction. Can it sustain that for five minutes? Almost certainly. For four hours? Probably not, because the house warms up and eventually the compressor has to run again to keep the customer within their comfort band.

So flexibility is not a number; it has a time dimension. The better object is a function of two arguments,

$$F(t, \tau), \tag{1.6}$$

where t is when the service begins and τ is how long the response must be sustained. For the thermostat, numbers like

$$F(17{:}00,\ 15\text{ min}) = 2.5\,\text{kW} \qquad \text{but} \qquad F(17{:}00,\ 4\text{ h}) = 0.7\,\text{kW}$$

are entirely plausible. The second number is the *average* reduction the device can hold over four hours without violating comfort: it may deliver 2.5 kW at the start, then cycle back on, so that over the whole window the sustainable reduction is much smaller. Figure 1.3 plots this dependence for the two devices used in this chapter: the battery of Section 1.3.3 and the air-conditioned house of Section 1.5.2. The horizontal axis is the required duration τ; the vertical axis is the average reduction that can be sustained over that duration. The battery curve is flat until its energy is exhausted and then falls as $1/\tau$; the thermostat curve is flat only until the compressor must restart and then falls toward a small sustained value.

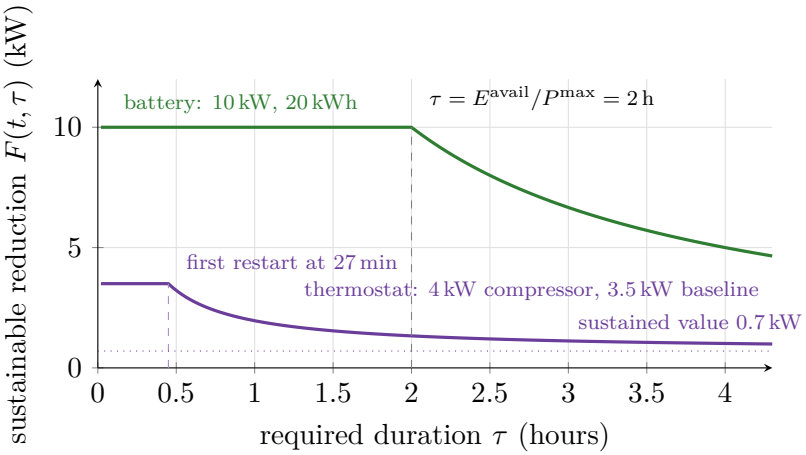


Figure 1.3: Sustainable average reduction versus required duration for the two devices of this chapter. The battery (green) holds its full 10 kW until its 20 kWh is exhausted at two hours, then the average it can sustain falls as E^{avail}/τ. The thermostat (purple) is the air-conditioned house of Section 1.5.2: a 4 kW compressor with a 3.5 kW baseline average, switched off at t. The full 3.5 kW reduction relative to baseline lasts until the compressor must restart at 27 minutes; after that the average reduction over the window, $[3.5\,t^* + 0.7\,(\tau - t^*)]/\tau$, settles toward the 0.7 kW that holding 24°C instead of 22°C sustains.

A conventional generator, by contrast, has essentially $F(t, \tau)$ independent of τ as long as there is fuel. The dependence on τ is what makes a DER a fundamentally different kind of resource.

Principle 1.2. *DER flexibility is* time-dependent*: it is a function $F(t, \tau)$ of when the response starts and how long it must last, and for most devices it decays with τ.*

1.3.3 Power and energy limits

The duration dependence has a simple closed form for a battery, which makes the battery the clearest case. Suppose a battery has a rated power of $P^{\max} = 10\,\text{kW}$ and an available energy of $E^{\text{avail}} = 20\,\text{kWh}$. It can deliver 10 kW for approximately

$$\frac{E^{\text{avail}}}{P^{\max}} = \frac{20\,\text{kWh}}{10\,\text{kW}} = 2 \text{ hours},$$

but it cannot provide 10 kW for four hours. Over four hours the most it can average is $20/4 = 5\,\text{kW}$.

So there are at least two independent quantities that describe a DER:

P	**how fast** the device can move power (kW)
E	**how long** it can keep doing so (kWh)

Figure 1.4 draws the feasible region in the (P, τ) plane: any request with $P \le P^{\max}$ and $P\,\tau \le E^{\text{avail}}$ is feasible.

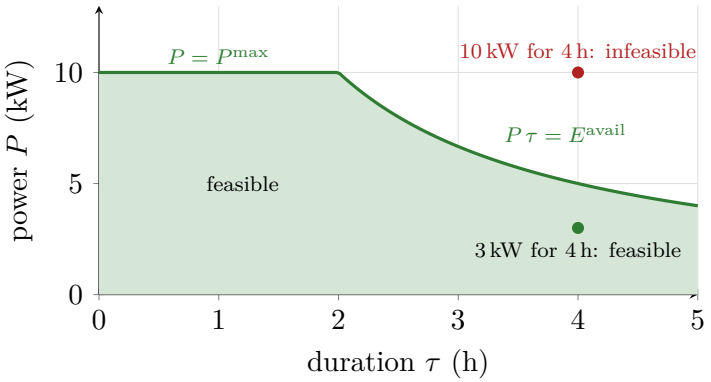


Figure 1.4: The power–duration capability of a 10 kW, 20 kWh battery. The power limit caps the height, the energy limit caps the area under a request.

For the other device types the "energy" is not stored as electricity but the same two-quantity picture applies:

- For a **thermostat**, the analog of stored energy is the building's *thermal inertia*: how much heat the walls, floors and air can absorb before the temperature leaves the comfort band.
- For an **EV**, it is the combination of the battery's SOC and how long the car remains plugged in before t_{dep}.
- For a **water heater**, it is the tank's thermal mass between the minimum acceptable and maximum safe water temperature.

This is why the state-space picture of Section 1.2.1 maps so naturally onto DERs: "how fast" is a limit on u and "how long" is a limit on how far x may travel.

1.3.4 The feasible control set and the state transition

Sections 1.2.2 to 1.3.3 together give the appropriate description of a DER. Instead of saying "this battery is a 10 kW resource" one says

$$u_t \in \mathcal{U}(x_t), \tag{1.7}$$

read as: *the allowable control actions depend on the current state.* For the battery, $\mathcal{U}(\text{SOC})$ is the interval $[-P^{\text{dis}}(\text{SOC}),\ P^{\text{chg}}(\text{SOC})]$, where the discharge limit P^{dis} is the curve of Figure 1.2

and the charge limit P^{chg} has the mirror-image shape near full charge; for the thermostat it is the set of setpoints that keep T_{in} inside the comfort band over the next step.

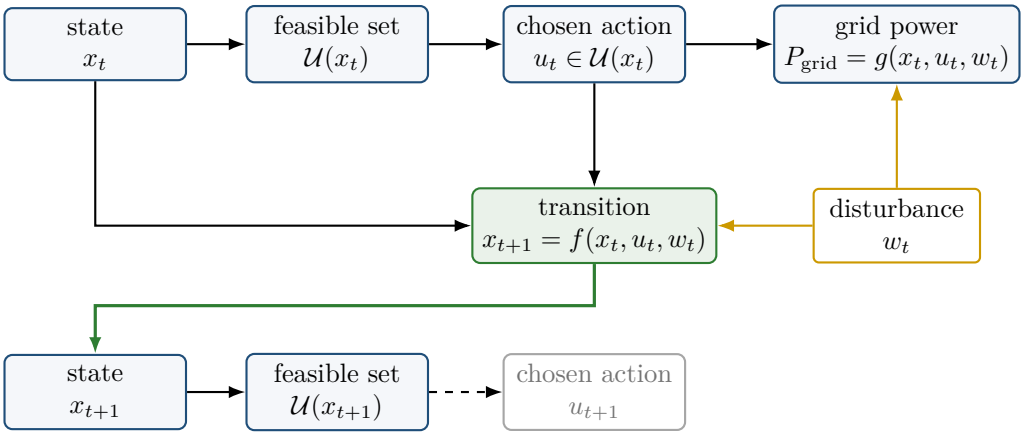
Then the device evolves:

$$\boxed{x_{t+1} = f(x_t, u_t, w_t)} \tag{1.8}$$

This equation accounts for a large part of the DERMS problem, because it states that *today's dispatch changes tomorrow's flexibility*:

- If I heavily discharge a battery now, $\text{SOC}_{t+1} \downarrow$ and less is available later.
- If I pre-cool a home now, $T_{\text{in},t+1} \downarrow$ and I may have *more* flexibility later.
- If I delay EV charging now, SOC_{t+1} may become too low to meet the driver's departure requirement.

DER control is therefore intrinsically *intertemporal*: a decision at t cannot be evaluated without its effect on $\mathcal{U}(x_{t+1})$, $\mathcal{U}(x_{t+2})$, and so on. Figure 1.5 shows the loop.



Intertemporal coupling. The action chosen at t sets the state at $t+1$, and the state at $t+1$ sets which actions are allowed at $t+1$.

Figure 1.5: The state–control loop of a single DER. The disturbance w_t (weather, occupancy, driving) enters both the output and the transition. The green arrow is the coupling that makes dispatch a planning problem rather than a switch.

Principle 1.3. *Because $x_{t+1} = f(x_t, u_t, w_t)$, every dispatch decision spends or stores future flexibility. DER control is intertemporal.*

1.4 Constraints and uncertainty

Two further properties come from outside the device: the customer's requirements and the aggregator's ignorance of the future.

1.4.1 Customer constraints

Conventional generation does not have this in the same form. A thermostat is installed because somebody wants to be comfortable. An EV exists because somebody needs to drive. A home battery may exist mainly for backup during outages. So the control problem carries constraints

that express the owner's purpose:

$$\text{thermostat:} \quad T^{\min} \le T_{\text{in}}(t) \le T^{\max} \quad \text{for all } t, \tag{1.9}$$

$$\text{EV:} \quad \text{SOC}(t_{\text{dep}}) \ge \text{SOC}^{\text{req}}, \tag{1.10}$$

$$\text{battery:} \quad \text{SOC}(t) \ge \text{SOC}^{\text{res}} \quad \text{for all } t. \tag{1.11}$$

These are not secondary conditions. They define whether a grid service is feasible at all. A 4 kW reduction that leaves the house at 27°C is not a service; it produces a complaint and, typically, withdrawal from the program. Formally, the customer constraints are part of what shapes $\mathcal{U}(x_t)$ in (1.7): the feasible set is the set of actions that respect both the device's physics *and* the owner's requirements over the horizon.

Two features make customer constraints harder than device constraints:

- They are stated over a *horizon*, not an instant. Equation (1.10) is a constraint at a future time t_{dep} that restricts actions now.
- They are *soft in practice*. A customer who is slightly uncomfortable may tolerate it once; one who is uncomfortable repeatedly will override the thermostat or leave the program. The override is itself part of the disturbance w.

Principle 1.4. *DER flexibility is* customer-constrained*: comfort bands, departure requirements and reserves are hard constraints on the control problem, and violating them removes the resource.*

1.4.2 Uncertainty

The central forecasting problem can now be stated. Suppose 10 000 EVs are enrolled, each with a 7 kW charger. Does the aggregator have

$$10\,000 \times 7\,\text{kW} = 70\,\text{MW}$$

available? It does not. At 6 PM on a given day:

- some cars are not plugged in;
- some are already fully charged, so there is no charging to defer;
- some need charging urgently to meet SOC^{req} by t_{dep};
- some customers will override the control;
- chargers have different ratings and some are on single-phase circuits that cannot reach 7 kW;
- some communications will fail, so the command never arrives or the confirmation never returns.

Figure 1.6 shows the nameplate total shrinking through these filters to the amount actually deliverable.

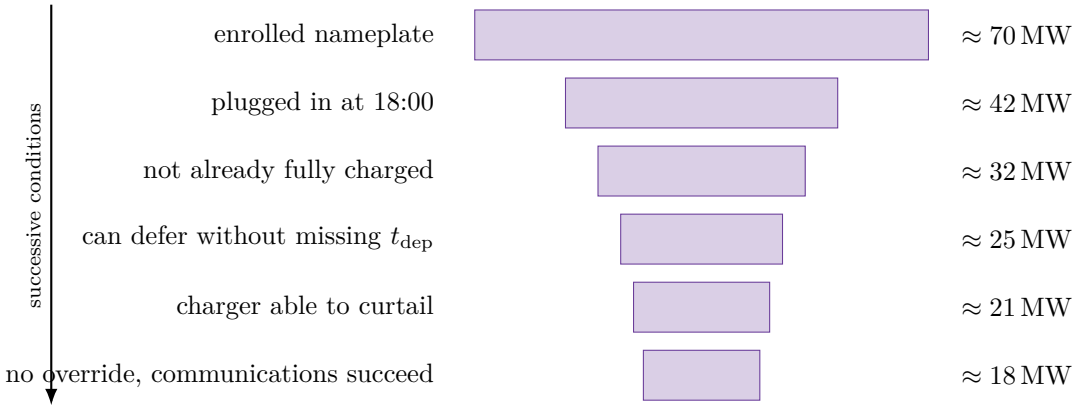
The consequence is that P^{avail} is not a number to be computed but a random variable to be forecast. What the aggregator really wants is a probability distribution,

$$p\Big(P^{\text{avail}} \,\Big|\, \text{information available at time } t\Big), \tag{1.12}$$

conditioned on everything observed so far: telemetry from the devices, weather, calendar, historical participation rates. This is the point where the two tracks of the lesson sequence meet:

Track A asks	*How much flexibility will be available?*
Track B asks	*Given that flexibility, how should we use it?*

Principle 1.5. *DER flexibility is* uncertain*: the aggregator never knows P^{avail} exactly, only a distribution conditioned on current information. Forecasting and control cannot be separated.*



The numbers are illustrative. The shape is not: the deliverable amount is a small, uncertain fraction of nameplate, and each condition is itself a random quantity on any given day.

Figure 1.6: From nameplate to deliverable: 10 000 enrolled EVs at 7 kW each are a 70 MW label but an 18 MW resource on this evening, and a different number tomorrow.

1.5 From one device to a fleet

The chapter closes its argument by placing the single device in the hierarchy the rest of the book climbs, and by working one thermostat through an event.

1.5.1 The hierarchy from device to fleet

Figure 1.7 organizes the domain. Reading left to right on the first row: a *device* has a *state*; the state determines the *feasible control*; applying a control produces a *power response*; the response, net of the baseline and over a duration, is the *flexibility*. The second row is the transition that the later lessons are about: individual flexibilities are aggregated into a *fleet* flexibility, and a fleet operated as a single controllable plant is a *virtual power plant*.

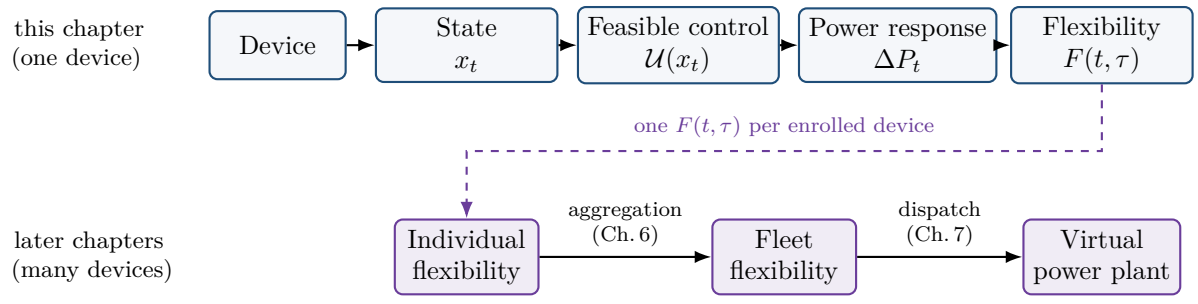


Figure 1.7: The two-level hierarchy of the subject. The top row is one device; the bottom row is what the rest of the book builds on top of it.

The step from the first row to the second, aggregation, carries most of the mathematical difficulty (fleet feasible sets are not simply sums of individual ones), and it is also where the uncertainty of Section 1.4.2 becomes manageable: ten thousand uncertain devices are far more predictable in total than any one of them.

1.5.2 Worked example: thermostat pre-cooling

All of the properties above appear in a single air-conditioned home. The parameters are listed in Table 1.2.

Table 1.2: Parameters of the thermostat example.

Air conditioner electrical power when running	$P_{\mathrm{AC}} = 4\,\mathrm{kW}$
Indoor temperature at 5 PM	$T_{\mathrm{in}} = 22°\mathrm{C}$
Customer comfort band	$21°\mathrm{C} \le T_{\mathrm{in}} \le 24°\mathrm{C}$
Outdoor temperature	$T_{\mathrm{out}} = 32°\mathrm{C}$
Current setpoint	$T_{\mathrm{set}} = 22°\mathrm{C}$, compressor running

The event. At 5 PM the aggregator raises the setpoint by one degree, $T_{\mathrm{set}} : 22°\mathrm{C} \to 23°\mathrm{C}$. The thermostat is now satisfied, the compressor turns off, and the grid load changes approximately

$$4\,\mathrm{kW} \to 0, \qquad \text{so initially } \Delta P \approx 4\,\mathrm{kW}.$$

The decay. But the house is 10 degrees cooler than outdoors, so heat leaks in and the indoor temperature starts rising,

$$\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} > 0.$$

A simple and adequate model of a house for this purpose is a single thermal mass C (kWh per °C) behind a thermal resistance R (°C per kW) to the outdoors, with the air conditioner removing heat at rate Q (kW thermal) when it runs:

$$C\,\frac{\mathrm{d}T_{\mathrm{in}}}{\mathrm{d}t} = \frac{T_{\mathrm{out}} - T_{\mathrm{in}}}{R} - Q\,s(t), \qquad s(t) \in \{0, 1\}. \tag{1.13}$$

With the compressor off ($s = 0$) the solution is an exponential approach to T_{out} with time constant RC:

$$T_{\mathrm{in}}(t) = T_{\mathrm{out}} - \left(T_{\mathrm{out}} - T_{\mathrm{in}}(0)\right) e^{-t/RC}. \tag{1.14}$$

Eventually T_{in} reaches $T^{\max} = 24°\mathrm{C}$ and the compressor must restart to avoid violating the comfort band. With the numbers of Table 1.2 and $RC = 2\,\mathrm{h}$, equation (1.14) gives the first restart at

$$t^{*} = RC \ln \frac{T_{\mathrm{out}} - T_{\mathrm{in}}(0)}{T_{\mathrm{out}} - T^{\max}} = 120\,\mathrm{min} \times \ln \frac{10}{8} \approx 27\,\mathrm{min}.$$

From then on the compressor cycles to hold 24°C. The heat it must remove per unit time is now proportional to $T_{\mathrm{out}} - 24 = 8°\mathrm{C}$ instead of $T_{\mathrm{out}} - 22 = 10°\mathrm{C}$, so its average electrical power is 8/10 of the baseline average. If the baseline average is 3.5 kW (the 4 kW compressor running 87.5% of the time at 22°C), the sustained load is 2.8 kW and the sustained reduction is 0.7 kW.

Two reductions therefore describe the same event: 4 kW at the instant the compressor stops, measured against the instantaneous load, and 3.5 kW measured against the baseline *average*, which is the quantity a settlement uses. The distinction is the subject of Chapter 2. Either way the full reduction lasts only until t^{*} and then settles toward 0.7 kW; in the language of Section 1.3.2, $F(17{:}00, 15\,\mathrm{min})$ is the full value while $F(17{:}00, 4\,\mathrm{h})$ is about 1 kW, and Figure 1.3 plotted exactly this history.

Key idea

Available flexibility decays as the thermal state evolves, *unless you manipulate the state beforehand.*

Pre-cooling. The transition equation (1.8) says the state at the start of the event is something the aggregator can influence. Before 5 PM, run the air conditioner a little harder and bring the house down,

$$T_{\text{in}} : 22°\text{C} \to 21.5°\text{C},$$

which is still inside the comfort band. The building's thermal mass now holds extra "cold", and when the compressor is switched off at 5 PM it takes longer for T_{in} to reach 24°C. Future flexibility has been increased by spending a little energy earlier. Figure 1.8 plots both trajectories using equation (1.14) with a 2-hour time constant.

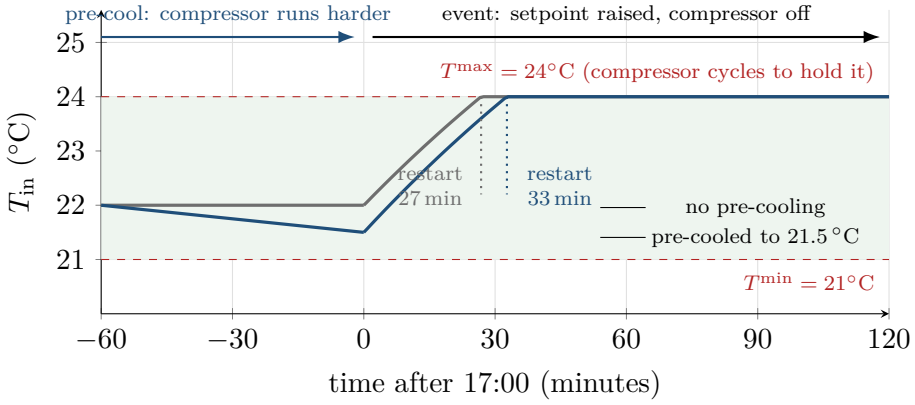


Figure 1.8: Indoor temperature around a 5 PM event. The shaded band is the customer's comfort range; the horizontal axis is minutes relative to the start of the event. Before the event the gray trajectory holds 22°C and the blue one is cooled to 21.5°C. After the event starts both rise according to equation (1.14) with $RC = 2$ h and $T_{\text{out}} = 32°\text{C}$ until they reach $T^{\max}$, where the compressor restarts and cycles. Pre-cooling by half a degree moves the first restart from 27 to 33 minutes; a deeper pre-cool, or a slower building, moves it more.

The gain in this example is modest because half a degree is a small investment in a fast building. With a 3-hour time constant and a pre-cool to 21°C, the first restart moves from 40 minutes to 57. The qualitative lesson does not depend on the numbers: pre-cooling is a genuine control problem, with a cost (the extra energy and the slightly higher load before the event) and a benefit (more flexibility during it), and the right amount depends on the forecast of T_{out}, on the building's RC, and on when the event will be called. That is exactly the shape of the dispatch problems in the later chapters.

1.6 The working definition

The chapter closes with the definition to retain, followed by Table 1.3, which lists the four properties of DER flexibility established above and the mathematical object that captures each.

Definition 1.1 (DER as a grid resource)**.** A DER becomes a grid resource when its *feasible, state-dependent power response* can be *predicted* and *controlled* while respecting *device and customer constraints*.

Each italicized phrase is one of the sections above: feasible (Section 1.3.4), state-dependent (Section 1.2.2), predicted (Section 1.4.2), controlled (Sections 1.3.1 and 1.3.2), device and customer constraints (Section 1.4.1). The definition is more useful than the expansion "distributed energy resource" alone, and it already explains why an aggregator needs forecasting, optimization and control, not merely an interface that turns devices on and off.

Table 1.3: The four properties of DER flexibility, and the object that captures each.

Property	Object	Consequence
State-dependent	$P^{\text{avail}}(x)$, $\mathcal{U}(x)$	nameplate is not availability
Time-dependent	$F(t,\tau)$, $x_{t+1} = f(x_t, u_t, w_t)$	flexibility decays, and today's use changes tomorrow's
Customer-constrained	$T^{\min} \le T_{\text{in}} \le T^{\max}$, $\text{SOC}(t_{\text{dep}}) \ge \text{SOC}^{\text{req}}$	the service is feasible only inside the owner's band
Uncertain	$p(P^{\text{avail}} \mid \text{info}_t)$	forecast a distribution, not a number

1.7 Exercises

Exercise 1.1. A home battery has $P^{\max} = 5\,\text{kW}$, a usable capacity of 13.5 kWh, is at SOC = 60%, and the owner keeps a 20% reserve. For how long can it discharge at full power? What is $F(t, 3\,\text{h})$?

Exercise 1.2. Using equation (1.14), show that the time until the compressor must restart after being switched off at $T_{\text{in}}(0)$ is $t^* = RC \ln\!\big[(T_{\text{out}} - T_{\text{in}}(0))/(T_{\text{out}} - T^{\max})\big]$. Evaluate it for $RC = 2\,\text{h}$, $T_{\text{out}} = 32°\text{C}$, $T^{\max} = 24°\text{C}$ and $T_{\text{in}}(0) \in \{22, 21.5, 21\}°\text{C}$.

Exercise 1.3. An EV arrives at 18:00 with SOC = 30% on a 60 kWh pack, charges at 7 kW, and must leave at 07:00 with SOC $\ge$ 80%. How many hours of charging are needed, and how many hours of *deferral* does the aggregator have? What happens to that deferral if the driver changes the departure to 22:00?

Exercise 1.4. Write the feasible set $\mathcal{U}(x_t)$ for the EV of the previous exercise at a time t before departure, in terms of SOC_t, $t_{\text{dep}} - t$, SOC^{req} and the charger rating.

1.8 Connection to the TERRA framework

TERRA models a physical system as a graph: *nodes* are places where a conserved quantity is balanced, *edges* carry flows of that quantity between nodes, and *components* attached to the graph contribute equations (residuals) that must all be satisfied at once. Every quantity the engine solves for is collected in one state vector, and the whole model is the statement "all residuals are zero". On top of that, the framework treats unknown parameters and unmeasured flows as quantities to be *inferred* from measurements, with a probability distribution rather than a point value. The lesson of this chapter maps onto that picture almost line by line.

- **The house is a conservation node.** Equation (1.1) is a power balance at the point of common coupling: inflow from the grid equals the sum of the device flows. In TERRA that is one node with five edges (grid, load, EV, battery, PV) and one residual, $P_{\text{grid}} - P_{\text{load}} - P_{\text{EV}} - P_{\text{bat}} + P_{\text{PV}} = 0$. A feeder with a thousand houses is a thousand such nodes joined by the network's own balance equations.
- **Each device is a component with a state.** The battery contributes its state-of-charge dynamics (an accumulator), the house contributes the thermal mass of equation (1.13) (a second conservation law, for heat, in the same graph), the EV contributes a battery plus a plugged-in schedule. The transition map f of equation (1.8) is what the engine integrates in time; the output map g is the set of edge flows the solved state implies. Multi-conservation, electrical and thermal in one model, is precisely what a thermostat DER needs.

- **Feasible sets and customer constraints are bounds on the state.** $T^{\min} \le T_{\text{in}} \le T^{\max}$ and $\text{SOC} \ge \text{SOC}^{\text{res}}$ are constraints on components of the state vector. Checking whether a dispatch is feasible is checking whether the integrated trajectory stays inside those bounds; finding the best dispatch is an optimization over the control inputs with the graph as the model.

- **Flexibility $F(t, \tau)$ is a query on the model.** "Switch this compressor off at t and integrate forward: how long until a bound is hit, and what was the average change in the grid edge flow?" is a forward simulation of the graph. The chapter's $F(t, \tau)$ is that query's answer, and the pre-cooling experiment of Section 1.5.2 is the same query run from a different initial state.

- **Uncertainty is inference, not a fudge factor.** The distribution $p(P^{\text{avail}} \mid \text{info}_t)$ of equation (1.12) is what the framework produces when device states (an EV's SOC, a building's RC) are only partially observed: measurements at the meter constrain, but do not pin down, the hidden states, and the posterior over those states propagates into a posterior over the available power. The funnel of Figure 1.6 is a chain of such partially observed conditions.

- **Aggregation is the same graph, larger.** The step from one device to a fleet (Figure 1.7) adds nodes and edges but no new kind of equation. This is the practical reason to adopt a conservation-graph view from the first lesson: the single-house model of Figure 1.1 and the virtual-power-plant model of a later chapter are the same object at different scales.

In summary, a DER is a component on a conservation graph whose state is partially known and whose control is constrained; a grid resource is such a component whose edge flow can be predicted and steered inside those constraints. The rest of the book is about doing that for many components at once.

END OF THE PREVIEW

Distributed Energy Resources as Grid Resources

These are the opening pages. The complete book runs 308 pages.

The full PDF opens with a token. Write to **rmarquez123@gmail.com**, name the book, and I will send you one.

One token opens every book on the site.

Ricardo Marquez, Ph.D.

ricardo-marquez.com