A Graphical Measure of Aggregate Flexibility for Energy-Constrained Distributed Resources

Michael P. Evans, Simon H. Tindemans, *Member, IEEE* and David Angeli, *Fellow, IEEE*

Abstract—We consider the problem of dispatching a fleet of heterogeneous energy storage units to provide grid support. Under the restriction that recharging is not possible during the time frame of interest, we develop an aggregate measure of fleet flexibility with an intuitive graphical interpretation. This analytical expression summarises the full set of demand traces that the fleet can satisfy, and can be used for immediate and straightforward determination of the feasibility of any service request. This representation therefore facilitates a wide range of capability assessments, such as flexibility comparisons between fleets or the determination of a fleet's ability to deliver ancillary services. Examples are shown of applications to fleet flexibility comparisons, signal feasibility assessment and the optimisation of ancillary service provision.

Index Terms—Distributed resources, energy storage systems, optimal control, aggregation, ancillary service

NOMENCLATURE

n	Total number of devices
D_i	ith device
$\mathcal{N}$	Set of all devices
$e_i(t)$	Energy of the ith device at time t
$u_i(t)$	Power output of the ith device at time t
$u(t)$	Vector of power outputs at time t, across all devices
$\bar{p}_i$	Maximum power rating of the ith device
$\bar{p}$	Vector of maximum power ratings, across all devices
P	Diagonal matrix of maximum power ratings, across all devices
$\mathcal{U}_{\bar{p}}$	Product set of power constraints, with vector of maximum powers $\bar{p}$
$x_i(t),\ z_i(t)$	Time-to-go of the ith device at time t
$x(t),\ z(t)$	Vector of time-to-go values at time t, across all devices
$\mathcal{X}$	State-space
$X(t)$	Vector of distinct time-to-go values at time t
$\mathcal{N}_j$	jth subset of devices
$x_{\mathcal{N}_j}(t)$	Equal time-to-go value of all members of $\mathcal{N}_j$, at time t
$\bar{U}_j$	Vector of maximum power outputs, across the jth subset of devices

$P^r(t)$	Reference at time t
$P^r_{[t_0,t)}$	Reference signal truncated to the interval $[t_0,t)$
$\mathcal{F}_{\bar{p},x}$	Feasible set from state x, with vector of maximum powers $\bar{p}$
$E_{P^r}(p)$	E-p transform of reference $P^r(\cdot)$, evaluated at power level p
$R(t)$	Worst-case reference at time t
$\Omega_{\bar{p},x}(p)$	Capacity curve from state x, with vector of maximum powers $\bar{p}$, evaluated at power level p

I. INTRODUCTION

A. Background

We consider the problem of operating a fleet of storage devices, i.e. a collection of these which are dispatched in a coordinated manner, in order to provide system support in supply-shortfall conditions. Recent years have seen the proliferation of multiple forms of storage onto electricity grids and, crucially, many of these resources are energy-constrained, for example due to physical capacity limits or operational limits set by users. These storage technologies can provide a range of valuable services to the system, with one advantage being their ability to shift consumption in time, thereby compensating for fluctuations in the output of intermittent generation. This offers significant potential as a means to replace conventional generation as electricity networks are decarbonised, according to the World Energy Council [1]. Increasingly, system operators are offering frameworks for convenient delivery of grid-support services; to date these include the California Independent System Operator (CAISO) [2] and the Pennsylvania-New Jersey-Maryland Interconnection RTO (PJM) [3] in the US. The latter offers the control framework within which we base our research: a system-wide regulation signal, representing the mismatch between electricity supply and demand, is broadcast and service participants endeavour to track this signal. Thus distributed resources are controlled centrally via their aggregate response. For the purposes of description we assume that an aggregator, defined as the entity responsible for the provision of system services through dispatch of the fleet, is contracted by the network operator as follows. The amount of power requested is updated at regular intervals, and the aggregator must make a decision as to which of its resources to deploy to meet the request. This aggregation of small participants into composite entities occurs in many ancillary service markets with minimum size requirements. Example fleets include uninterruptible power supplies with storage headroom and electric vehicle (EV) batteries returned

M.P. Evans and D. Angeli are with the Department of Electrical and Electronic Engineering, Imperial College London, UK (email: m.evans16@imperial.ac.uk; d.angeli@imperial.ac.uk). D. Angeli is in addition with the Department of Information Engineering, University of Florence, Italy. S.H. Tindemans is with the Department of Electrical Sustainable Energy, TU Delft, Netherlands (email: s.h.tindemans@tudelft.nl). This work was supported by EPSRC studentship 1688672.

to the manufacturer for recycling; grid support offers a convenient end-of-life deployment for the latter once they have deteriorated beyond usability in the EV. Note that if significant loss of load were at stake the network operator might instead take on the role of central dispatcher.

We focus on the decision making of the aggregator, and consider how to best dispatch devices in the presence of uncertain demand. Previous literature has addressed the coordinated dispatch of a fleet of distributed devices in comparable problem frameworks. Prior approaches can be broadly clustered into three categories: optimal control, mean-field control [4]–[6] and transactive energy [7]–[10]. Our approach lies in the first category and we directly optimise power flows between storage and the grid via explicit feedback policies. In contrast to mean field approaches, this allows us to achieve guaranteed optimal behaviour, as opposed to in expectation.

Within applications of optimal control to the dispatch of distributed devices, a common approach has been the composition of an optimisation problem of standard form for the entire network and across the full time horizon. The dispatch of each device at each time instant is then assigned as a decision variable in the optimisation problem. Objectives chosen have included minimisation of charging costs [11] or more generalised operational costs [11]–[16], or flattening of load profiles [17]. Varying based upon the problem framework, solution techniques applied to these problems have included linear solvers [12], mixed-integer linear solvers [13], combined mixed-integer linear/nonlinear solvers [11], robust optimisation [14], [15], particle swarm algorithms [16] and genetic algorithms [17].

We are interested in achieving optimality in a simpler way, negating the need to perform the large computations of this prior work. Moreover, we are interested in applying our techniques to settings with a complete lack of knowledge about future request signals, including probability distributions or forecasts. In these cases one would be unable to compose an optimisation problem covering all future time instants. For example, [18] composed a simpler optimisation problem and presented a decentralised solution, but still required a foreknown demand profile.

Specifically, we consider approaches where the optimal dispatch problem is divided into two coupled sub-problems. A real-time *control* algorithm dispatches the fleet of devices according to a common control signal, without knowledge of the future (i.e. a greedy algorithm). This control algorithm is paired with a matching *scheduling* problem that determines the capability of the fleet to meet future requests. It uses an aggregate representation of the flexibility limits of the device fleet to ensure that only feasible responses are scheduled. An example of this approach, applied to mean field control of thermostatically controlled loads, can be found in [4] and [19], for the control and scheduling problems respectively.

A number of previous proposals to control single storage devices can be interpreted through the lens of a greedy control strategy plus a separate scheduling component. [20]–[22] all implemented a greedy buffer policy in which the device charges as quickly as possible under conditions of excess supply and discharges as quickly as possible under excess demand, but chose different reference levels for defining this mismatch. In [20], thermal plant was dispatched so that it was forecast to maintain the storage unit at a predefined set level, and buffering was done in real-time to counter forecast errors. In [21] the buffer comparison level was instead the forecast net demand plus either a fixed offset or an offset based upon the forecast storage level. Cruise *et al.* [22] were concerned with arbitrage applications, where the authors showed that the choice to maximally charge or discharge should be based upon the notional value per unit of energy stored as follows: charge if this is above the buy price, discharge if this is below the sell price, otherwise do nothing.

Results with strict optimality proofs can be obtained by restricting the operational regime to discharging only. This corresponds to an assumption that the ability to satisfy the power request takes precedence over other objectives, allowing one to study properties of the system in a general sense, without considering price dynamics in detail. A natural objective for the aggregator within this setting is to postpone the time to failure, i.e. the time at which the aggregator is first unable to meet the requested power supply. In general, this approach is applicable whenever an event occurs which threatens security of supply, because it would give the network operator the most time in which to attempt recovery of normal operating conditions. A straightforward example of such an event would be the loss of a generator or transformer, or an islanding fault. Other relevant events might occur within nominally ordinary conditions, for example an unexpected shortfall in wind output or peak in demand. In these cases, purely surviving until the natural termination of the event would be sufficient to recover normal operation. In [23], the authors showed that, for a single energy storage device, the naive greedy discharge policy maximises inclusion in the set of future request signals that the device is able to meet. This policy therefore maximises time to failure. The corresponding scheduling problem is trivial for a single device, because the feasibility of a specific output signal is immediately determined from the power and energy limits of the device.

In [24], we proposed a greedy control policy that is optimal in this same sense for a fleet of arbitrary size. A feedback policy (which we will here refer to as the *optimal* policy) was presented and shown to maximise the time to failure, instantaneous maximum power and flexibility of a system supplied by energy-constrained distributed resources. The flexibility maximisation, in particular, was performed in a set-theoretic sense; which is to say that the set of future feasible scenarios is always as large as possible, regardless of the current fleet output. This can be interpreted to mean that, with no information about future signals, implementation of the policy is guaranteed to result in the latest failure time of the system. Moreover, by forming a general integrator model of storage, we are able to apply our analysis across a wide range of device types that have previously been considered within the paradigm of demand response [25]. These include EVs [26], [27], home storage devices [28], [29] and diesel generators [30], [31]. It should be noted that a comparable policy to that which we propose, applied to a continuum of devices, is also relevant in price-determined settings [32] and

enables the coordination of devices within areas of flat prices.

B. Contribution of this paper

We argue that the policy of [24] solves, in an optimal sense, the *control problem* for a fleet of energy-constrained distributed resources. While this policy can be used to determine fleet capability in a procedural way, i.e. by running the algorithm, this paper proposes an analytical method for undertaking the same assessment We present a transform that returns an aggregate representation of the fleet; this therefore enables longer-timescale planning and so can be used to solve the associated *scheduling problem*. Moreover, this transform can also be applied to a received request for immediate feasibility determination, which leads to greater insight into fleet capabilities. The contributions of this paper can be summarised as follows:

- We present a functional means to determine request feasibility, thereby improving upon the procedural application of the optimal policy of [24].
- We explore the range of capability assessments that the transform enables: specification of maximum system service provision; immediate feasibility determination for a received request and flexibility comparison between different fleets.
- We discuss how this transform has multiple useful properties that enable intuitive fleet characterisation.
- We show how this transform can be implemented graphically, thereby enabling feasibility determination by eye.
- We show how a graphical comparison can be made to the homogeneous fleet of the same total ratings.
- We derive additional analytical results relating to feasibility of requests, which lead to a greater understanding of the problem in general.

In deriving the transform, we utilise the result of [24] that there is a one-to-one mapping between feasible signals and those that will be met by the optimal policy. Use of the developed framework can be seen in Figure 1.
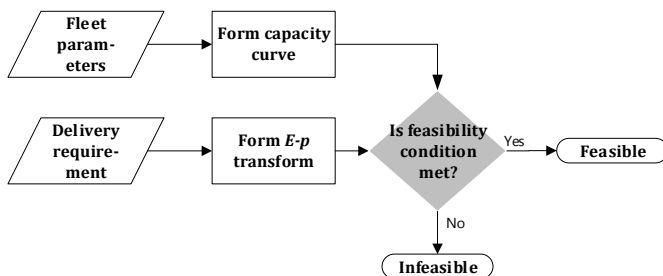

Fig. 1: Use of the presented framework to determine the feasibility of a given delivery requirement.

C. Motivating example

The technique that we present is able to compare fleets of devices utilised for system services, when often such a comparison appears far from trivial. To this end, the contribution of this paper is perhaps best motivated by means of an example. Consider the microgrid of Figure 2, onto which the system operator plans to connect a storage setup with one

of the following three device configurations (energy, power): A, consisting of (108 kWh, 4 kW) and (36 kWh, 18 kW); B, consisting of (104 kWh, 13 kW) and C, consisting of (90 kWh, 8 kW) and (54 kWh, 14 kW). Note that configurations A and C have the same total energy and power but distributed differently between devices. Configuration B has smaller values of each but is comprised of a single device. Hence, prior to the analysis that we present, it is unclear which configuration would be the best choice.
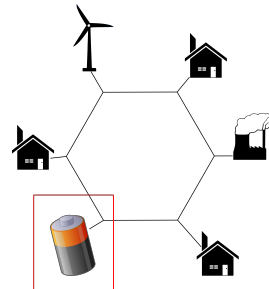

Fig. 2: An example microgrid, onto which the network operator plans to connect storage at the location highlighted by the red box.

D. Organisation of the paper

The remainder of this paper is organised as follows. Section II covers the mathematical formulation of the problem. This is followed by a presentation of analytical results in Section III, and the application of these results to numerical examples in Section IV. Section V then discusses the implications of these results. Finally, Section VI concludes the report and discusses relevant future work.

II. MATHEMATICAL FORMULATION

A. Problem description

We denote by n the number of energy-constrained storage units available to the aggregator, and define the set of all devices as $\mathcal{N} \doteq \{D_1, D_2, ..., D_n\}$. We do not impose any restrictions on homogeneity of devices and allow each device to have a unique discharging efficiency. For convenience we incorporate this into the model implicitly by considering the extractable energy of each device, $e_i(t)$. We choose the power delivered by each device to be the control input $u_i(t)$, and assume that this is measured externally so that efficiency is once again accounted for. This leads to integrator dynamics on the energy of each device,

$$\dot{e}_i(t) = -u_i(t), \tag{1}$$

subject to the assumed physical constraint $e_i(t) \geq 0$. We neglect network effects, in particular due to the high likelihood that storage devices are situated near to electricity-consuming devices, so that their deployment is unlikely to be constrained by network congestion or losses. We also assume the absence of cross-charging between devices, which corresponds to a regime of energy scarcity. We therefore restrict our devices to discharging operation only, so that the power of each device is constrained as $u_i(t) \in [0, \bar{p}_i]$, in which $\bar{p}_i$ denotes the maximum discharge rate of device D_i, and with the convention that discharging rates are positive. Note, once again, that there

is no homogeneity imposed on energy or power constraints across the ensemble of devices. We define the *time-to-go* of device D_i to be the remaining time for which it can run at its maximum power, i.e.

$$x_i(t) \doteq \frac{e_i(t)}{\bar{p}_i}, \tag{2}$$

and represent the state of each device by its time-to-go. We then form state, input and maximum power vectors as

$$x(t) \doteq \begin{bmatrix} x_1(t) & \dots & x_n(t) \end{bmatrix}^T, \tag{3}$$

$$u(t) \doteq \begin{bmatrix} u_1(t) & \dots & u_n(t) \end{bmatrix}^T, \tag{4}$$

$$\bar{p} \doteq \begin{bmatrix} \bar{p}_1 & \dots & \bar{p}_n \end{bmatrix}^T \tag{5}$$

respectively, so that we can write our dynamics in matrix form as $\dot{x}(t) = -P^{-1}u(t)$, in which $P \doteq \operatorname{diag}(\bar{p})$. We define the state space as $\mathcal{X} \doteq [0, +\infty)^n$, and form the product set of our constraints on all the inputs,

$$\mathcal{U}_{\bar{p}} \doteq [0, \bar{p}_1] \times [0, \bar{p}_2] \times \dots \times [0, \bar{p}_n], \tag{6}$$

allowing us to write our input constraints as $u(t) \in \mathcal{U}_{\bar{p}}$.

Now, as in [24], we partition the devices considered at a chosen time instant into collections of descending equal state value, without loss of generality leading to

$$x_1 = \dots = x_{s_1} > \dots > x_{s_{q-1}+1} = \dots = x_{s_q}, \tag{7}$$

in which q is the number of unique x_i-values. Note that we have dropped the explicit dependence on t as we are considering a single time instant. We then denote the subsets formed as $\mathcal{N}_i \doteq \{ D_j : x_j = x_{s_i} \}$, $i = 1, 2, \dots, q$, and similarly denote by $x_{\mathcal{N}_i} \doteq x_{s_i}$, $i = 1, 2, \dots, q$ the (single) time-to-go value of the devices in each subset. We finally form a condensed state vector of distinct values only as

$$X \doteq \begin{bmatrix} x_{\mathcal{N}_1} & \dots & x_{\mathcal{N}_q} \end{bmatrix}^T, \tag{8}$$

and denote the maximum power vectors corresponding to each subset as

$$\bar{U}_i \doteq \begin{bmatrix} \bar{p}_{s_{i-1}+1} \dots \bar{p}_{s_i} \end{bmatrix}^T, \quad i = 1, \dots, q, \tag{9}$$

with the convention that $s_0 = 0$.

We denote by $P^r : [0, +\infty) \mapsto [0, +\infty)$ a power reference signal received by the aggregator, and in addition denote a truncation of such a signal as

$$P^r_{[t_0, t)} \doteq \begin{cases} P^r(\tau), & \text{if } \tau \in [t_0, t) \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

We utilise equivalent notation for the truncation of any other signal also. For a reference to be *feasible*, there must exist a control signal able to satisfy it for all time without violating any constraints. These are both the power constraints on the control signal itself and the energy constraints that apply to the resulting state trajectory. We therefore define the set of feasible reference signals as follows:

Definition II.1. *The set of feasible power reference signals, for a system with maximum power vector $\bar{p}$ and initial state $x = x(0)$, is defined as*

$$\mathcal{F}_{\bar{p},x} \doteq \big\{ P^r(\cdot) \colon \exists u(\cdot), \ z(\cdot) \colon \forall t \geq 0, \ 1^T u(t) = P^r(t),$$
$$u(t) \in \mathcal{U}_{\bar{p}}, \ \dot{z}(t) = -P^{-1}u(t), \ z(0) = x, \ z(t) \geq 0 \big\}.$$

We interpret the inclusion of a given feasible set by another as an increase in *flexibility*.

B. Optimal Feedback Policy

Due to its maximisation of the feasible set, we will continually refer to the optimal policy of [24]. We therefore reproduce it here as follows, in which it should be noted that explicit time dependencies have been dropped as we are again considering instantaneous values. Without loss of generality, form subsets of devices with equal value in descending order, of which there will be q, so that $x_{\mathcal{N}_1} > x_{\mathcal{N}_2} > \dots > x_{\mathcal{N}_q}$. The explicit feedback law is then calculated as a fraction r_i of the maximum power $\bar{U}_i$ according to:

$$r_i = \begin{cases} 1, & \text{if } \sum_{j \leq i} 1^T \bar{U}_j \leq P^r \\ 0, & \text{if } \sum_{j < i} 1^T \bar{U}_j \geq P^r \\ \frac{P^r - \sum_{j \leq i} 1^T \bar{U}_j}{1^T \bar{U}_i}, & \text{otherwise,} \end{cases} \tag{11a}$$

$$u^*(x, P^r) \doteq \begin{bmatrix} r_1 \bar{U}_1^T & \dots & r_q \bar{U}_q^T \end{bmatrix}^T. \tag{11b}$$

In words, this policy allocates devices in descending order of time-to-go, at maximum power (with up to one subset of devices run at a fraction of maximum power to exactly meet the reference). This then depletes as few devices as possible, thereby maximising the number of available devices and so giving the fleet the best chance to meet an unknown future reference.

Denoting by $z^*(\cdot)$ the state trajectory under the application of (11), the closed-loop dynamics are

$$\dot{z}^*(t) = -P^{-1} u^* \big(z^*(t), P^r(t) \big). \tag{12}$$

III. RESULTS ON FEASIBILITY

In this section we present theoretical results which hold in general. These results are of greater importance to our arguments than their derivations and so, for clarity of argument, their proofs have been omitted from the main body of this paper and can be found in the Appendix.

A. Interesting system properties

The existence of a feasible set which includes all others (maintained via implementation of the optimal policy (11) - see [24] for more details) allows us to derive additional useful properties of the system considered. Moreover, the optimality of the presented policy allows us to utilise it as a proxy for this largest feasible set. We undertake such derivations here, implicitly utilising this policy in each case. The proofs of Lemmas III.1 and III.2 can be found in the Appendix.

Lemma III.1. *Piecewise constant reference signals are permutable whilst maintaining feasibility.*

One can interpret this result to mean that the ability of a system to satisfy a given reference signal is independent of the variation of that signal in time; rather it is dependent on the total time spent at each power level.

Lemma III.2. *Given two feasible piecewise constant reference signals which are permuted versions of one another, the final state under both references is the same when the optimal policy is used.*

B. E-p transform

The ability to utilise the policy (11) as a proxy for the feasible set also enables the derivation of a graphical representation of this set. In this section we present a transform, followed by a discussion of its uses to a grid operator or aggregator. The proofs of Lemma III.5 and Theorem III.8 can be found in the Appendix.

Definition III.3. *Given a power reference $P^r : [0, +\infty) \mapsto [0, +\infty)$, we define its E-p transform as the following function:*

$$E_{P^r}(p) \doteq \int_0^\infty \max\{P^r(t) - p, 0\}dt, \qquad (13)$$

interpretable as the energy supplied above any given power rating, p.

The following properties of this transform are of particular relevance to our analysis:

Property III.4. *The E-p transform intersects the $E-$ and $p-$axes at $\int_0^\infty P^r(t)dt$ (total energy supplied) and $\sup_t P^r(t)$ (maximum power supplied) respectively.*

Lemma III.5. *The E-p transform is convex and monotone.*

We additionally define the *worst-case* reference signal, and form its E-p transform as follows:

Definition III.6. *The worst-case reference signal, $R(\cdot)$, that can be fulfilled by a given system runs all devices at full power until they deplete; which will take place in order of ascending time-to-go. This signal can be calculated as follows:*

$$R(t) \doteq \sum_{i=1}^n \bar{p}_i[H(t) - H(t - x_i)], \qquad (14)$$

in which $H(\cdot)$ denotes the Heaviside step function.

Definition III.7. *We define the **capacity** of a system, $\Omega_{\bar{p},x}(p)$, to be the E-p transform of the worst-case reference signal that can be met by the system,*

$$\Omega_{\bar{p},x}(p) \doteq E_R(p), \qquad (15)$$

in which $R(\cdot)$ is defined as in (14).

We now consider comparing the E-p transform of an arbitrary reference to the capacity as follows.

Theorem III.8. *A reference signal $P^r(\cdot)$ is feasible if, and only if, its E-p transform is dominated by the capacity of the system, i.e. $E_{P^r}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p \iff P^r(\cdot) \in \mathcal{F}_{\bar{p},x}$.*

This result justifies our use of the terminology *worst-case* reference signal for $R(\cdot)$ and *capacity* for its E-p transform. Definitions III.6 and III.7 allow us to transform the complex object that is the set of feasible reference signals into a simple one-dimensional curve: that corresponding to the capacity of the system. This then enables a wide range of fleet capability assessments to be undertaken. For a given fleet, one could ascertain, for example, the longest ramp input of set gradient or highest-gradient ramp of a fixed duration that could be satisfied. Likewise, this could apply to pulse inputs, as we

will demonstrate in Section IV, or system services such as primary or secondary response.

If, in addition to the capacity curve, we form the E-p transform of a received reference signal, Theorem III.8 then allows for immediate and straightforward determination of the reference feasibility; simply by testing whether the capacity curve dominates the reference curve. This operation might provide valuable insight to a grid operator or aggregator when contemplating the feasibility of a request profile. Moreover, since the state fully defines the capacity of the system, the grid operator or aggregator would be able to make such deductions based solely upon the current state of their system.

We are also able to use Theorem III.8 to compare different device fleets as follows:

Corollary III.9. *The feasible set of system-state pair $(\bar{p}^a, x^a)$ includes that of $(\bar{p}^b, x^b)$ if, and only if, the capacity curve of $(\bar{p}^a, x^a)$ dominates that of $(\bar{p}^b, x^b)$, i.e.*

$$\Omega_{\bar{p}^a,x^a}(p) \geq \Omega_{\bar{p}^b,x^b}(p) \ \forall p \iff \mathcal{F}_{\bar{p}^a,x^a} \supseteq \mathcal{F}_{\bar{p}^b,x^b}. \quad (16)$$

An alternative use of the E-p transform would therefore be the comparison among fleets exemplified by the motivating example of the Introduction. We also point out here that, as a result of the convexity and monotonicity of the transform, forming a capacity curve based upon lower bounds to the charge level of each device would allow any of these assessments to be implemented in a robust manner. Coupled with Lemma III.5, Corollary III.9 also allows us to deduce the following:

Property III.10. *Given total energy and total power ratings E_{max} and p_{max} respectively, the most flexible distribution of devices has a capacity curve that is a straight line between $(0, E_{max})$ and $(p_{max}, 0)$. This is equivalent to a single device of the same ratings.*

Property III.11. *The deviation of a system's capacity curve from the maximum-flexibility curve of the same total ratings represents a **flexibility gap** resulting from its heterogeneous nature.*

IV. NUMERICAL RESULTS

A. Comparison between fleets

We here present the capacities of the three device configurations discussed as a motivating example in the Introduction. The three capacity curves can be seen in Figure 3, in which it can be seen that C is the unambiguous best choice as its E-p curve dominates the others. Clearly, the more even distribution of C as compared to A offers more flexibility. When comparing A and B, if the network operator expects high-power, short-duration or low-power, long-duration request signals then A is preferred. If, however, signals with intermediate values of power and energy are expected, then B is preferred instead.

B. Feasibility of a received request

1) Alternative policy choices: As in [24], we now compare indicative behaviour under the optimal policy to the implementation of two greedy heuristic alternatives. Note that [23] discusses sub-optimal policies such as these in more detail. The comparative policies chosen are as follows:
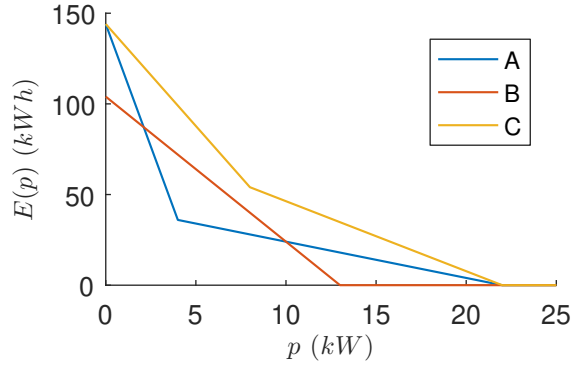
Fig. 3: Capacity curves for the motivating example of Section I-C

i) *Lowest Power First*. Order the devices by maximum power, without loss of generality leading to $\bar{p}_1 \leq \bar{p}_2 \leq \ldots \leq \bar{p}_n$, and allocate devices in order from D_1 to D_n as

$$u_i = \mathbb{1}[x_i > 0] \cdot \min\left\{\bar{p}_i, \ P^r - \sum_{j<i} u_j\right\}, \qquad (17)$$

in which $\mathbb{1}[\cdot]$ denotes the indicator function. Note that the choice of allocation between devices of equal maximum powers is made arbitrarily.

ii) *Proportion of Power*. In this case, no ordering is required and each device is run according to

$$u_i = \mathbb{1}[x_i > 0] \cdot \min\left\{\bar{p}_i, \ \frac{\bar{p}_i P^r}{\sum_{i:\ x_i > 0} \bar{p}_i}\right\}. \qquad (18)$$

2) Simulation results: We compose a scenario in which there are 10,000 devices with initial time-to-go and maximum power values generated from uniform distributions, as $x_i \sim U(0, 10)$ h and $\bar{p}_i \sim U(0, 1.5)$ kW respectively. In an attempt to model a realistic dispatch, we choose a stepwise reference signal that is updated hourly, and draw each value $P^r[k]$ from the normal distribution $P^r[k] \sim N(2, 0.8)$ MW. In this way all devices will be depleted by the end of a single day according to the optimal policy, and we set our simulation horizon to 1 day. We point out here that distributions are used solely for setting parameter values and therefore that this is a deterministic setup. The results of this case study can be seen in Figure 4, and can be interpreted as follows. The optimal policy provides the highest feasible reference up to its time to failure, as it postpones emptying devices until absolutely necessary, resulting in the latest time to failure. If this scenario represents some failure mode requiring less than 12.8 h for resolution, any of the three policies are capable of maintaining full functionality. If, however, resolution requires between 13.2 and 16.3 h, the policy that we present is the only one out of the three that is capable of avoiding lost load. Beyond 16.3 h we are able to say that there exists no policy capable of avoiding lost load.

3) E-p analysis: For the scenario described above, we here demonstrate the use of the E-p transform in determining the feasibility of the full reference over the 24 h period (known to be infeasible) and in addition the same signal truncated to the time to failure under the optimal policy (feasible by construction), as can be seen in Figure 5. We also plot the
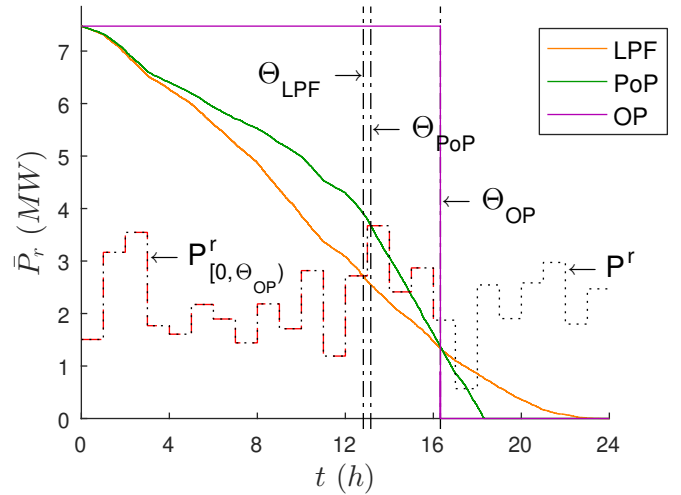

Fig. 4: The maximum available power, $\bar{P}^r(t)$, under the implementation of the optimal (OP), lowest power first (LPF) and proportion of power (PoP) policies. The time to failure under each policy is represented by a black dash-dotted line and denoted Θ. The corresponding reference signal is shown for comparison, both in full (the dotted black line) and truncated to the time to failure under the optimal policy (the dashed red line).

maximum-flexibility curve for comparison, and highlight the flexibility gap. As it can be seen, the truncation of the reference signal corresponds to bringing its E-p transform to just below the capacity curve, i.e. into the feasible region. In addition, the E-p and capacity curves intersect at $p = 0$, which corresponds to the use of all available energy and the resulting depletion of all devices by the time to failure. Moreover, in this case there would be no advantage to homogenising the fleet, as the same truncation of the reference would be required to achieve feasibility.
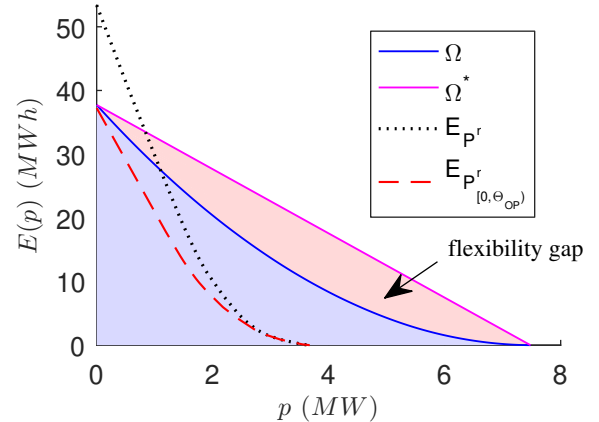

Fig. 5: Use of the E-p transform to determine reference feasibility, for the case study of Figure 4. Capacity curves are represented by solid lines and transformed reference curves by dashed or dotted lines (consistent with Figure 4). Ω denotes the capacity curve, and the feasible region defined by this is shaded in blue. Ω^* denotes the maximum-flexibility capacity curve; the red shaded area between these two curves is the flexibility gap.

C. Fleet capability assessment

We here determine the capability of a fleet to provide a user-defined system service. We choose as an example service

a pulse input of fixed duration, and demonstrate how the E-p transform can be used to find the largest-magnitude feasible pulse of 15 s duration. We choose the fleet to be 10 devices with equal maximum discharge rating $\bar{p}_i = 1.5$ kW. We generate initial state values for half the fleet uniformly, as $x_i \sim U(0, 20)$ s, and choose the other half to start full, at $x_i = 20$ s. This setup might, for example, represent uninterruptible power supplies that are allowed to take part in short-duration frequency regulation. A pulse signal will have a corresponding E-p curve that is linear, of gradient equal to the negative of the duration. The pulse magnitude is then equal to the p-intercept of this curve. Our task is to find the curve of this form with the largest p-intercept out of those which lie in the feasible region. Figure 6a demonstrates this approach across three pulses of varying magnitude, which are then shown in Figure 6b, and it can be seen that an 11.2 kW pulse is the largest that the fleet can feasibly meet. Note that this example is composed at a drastically smaller device scale than was previously explored, demonstrating how the presented technique is applicable regardless of the scale of the problem.
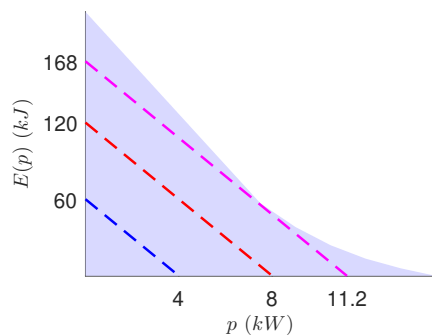
(a) Three example E-p curves corresponding to pulse inputs of 15 s duration. The magnitude of each pulse is equal to the curve's p-intercept. The blue shaded area represents the feasible region.
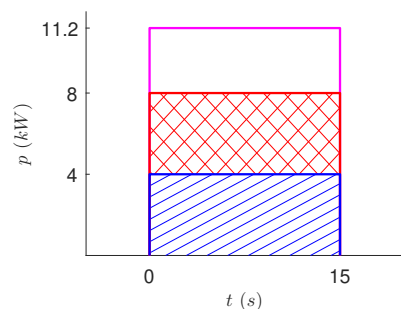
(b) The three pulses of 15 s duration considered. To enable a direct comparison, the unfilled magenta pulse lies behind the cross-hatched red pulse, which in turn lies behind the hatched blue pulse.

Fig. 6: Use of the E-p transform to determine the largest feasible pulse magnitude of 15 s duration. The colour of each E-p curve in Figure 6a matches that of the corresponding pulse signal in Figure 6b.

V. DISCUSSION

We here discuss the consequences of the above results. Firstly, as we have shown, the feasibility of a received or derived reference (for example that corresponding to a particular ancillary service) can be immediately determined. This analysis could also be extended into stochastic settings, i.e. those with random state, maximum power or reference values, by sampling scenario traces and rapidly checking the feasibility of each. This could, for example, allow a network operator to devise a day-ahead plan based upon a forecast as follows. If the probability, found through a Monte Carlo simulation, that the storage will be unable to meet the net demand exceeds some user-defined limit, the network operator could preventively take measures to reduce the high risk of loss of load. Alternatively, robust estimates could be composed from lower bounds on state or maximum power values, or from upper bounds on reference values. Similarly, the transform could be embedded in a range of algorithms, either as a rapid binary feasibility check or as inequality constraints in an optimisation problem. It is here where the computational improvements of the proposed transform over straightforward simulation might become crucial.

Ongoing work involves extending these results into scenarios in which infeasibility is found to occur. The one-to-one mapping between feasible references and those that will be met by the optimal policy enables one to identify routes to feasibility in these cases. Moreover, the convenient mathematical properties (convexity, monotonicity) of the E-p transform offer significant advantage in the construction of extended control strategies that minimise energy shortfalls.

We now briefly mention the scalability of the optimal policy in practical implementations. One need only broadcast the following three values: 1) the state above which devices should become active at full speed, 2) the state above which devices should become active at a fraction of full speed and 3) that fraction. The complexity of the broadcast will therefore remain unchanged as the number of devices increases. The calculation of these three values, however, requires up to one pass through the devices and will therefore scale as $\mathcal{O}(n)$. Similarly, it is worth mentioning the scalability of the proposed transform. This is formed via a single pass through the device fleet, hence scales as $\mathcal{O}(n)$. Where necessary, in order to decrease the computational complexity below this, it would also be possible to cluster devices by time-to-go and approximate from below the capacity curve for each cluster (using its smallest time-to-go value). The aggregated capacity curve could then be used to obtain a robust determination of feasibility.

VI. CONCLUSIONS AND FUTURE WORK

This paper has extended the results of [24], in the form of additional analytic properties of a system in which storage is used for grid support. We have conceived a transform that compactly represents the capacity of the system and can be used to characterise fleet capability; this then complements the optimal policy when making decisions at longer timescales. The transform can be used to calculate the maximum provision by the fleet of a range of system services, as well as directly

compare flexibility across different fleets. In addition the E-p transform allows for simple and immediate determination of the feasibility of any received request profile, through a comparison of E-p curves. We have discussed the ways in which these assessments could be extended to stochastic settings. The optimal policy can of course also be used for device scheduling across a range of applications.

This work has predominantly considered settings in which references are feasible, only performing our analysis up until the time to failure under the optimal policy. In future work the authors intend to extend this into scenarios in which a reference cannot be met for all time, in which case they plan to investigate the policy that would result in the least amount of lost load. Additional future extensions include the incorporation of charging requests and cross-charging among devices.

APPENDIX

A. Proof of Lemma III.1

Given a feasible reference signal $P^r(\cdot)$, consider the class of modified signals defined by:

1) An arbitrary partition of $[t_1, T)$,

$$P^r(\cdot) = P^r_{[t_1,t_2)} \otimes P^r_{[t_2,t_3)} \otimes ... \otimes P^r_{[t_v,t_{v+1}=T)}, \quad (A.1)$$

in which v is the number of parts and the $\otimes$ operator denotes the concatenation of two signals.

2) An arbitrary permutation of the part indices $1, 2, ..., v$. Letting ψ denote the vector dictating the order of these parts, the partition induced by ψ is then

$$\tilde{P}^r(\cdot) \doteq P^r_{[t_{\psi_1},t_{\psi_1+1})} \otimes P^r_{[t_{\psi_2},t_{\psi_2+1})} \otimes ... \otimes P^r_{[t_{\psi_v},t_{\psi_v+1})}. \quad (A.2)$$

$P^r(\cdot)$ is feasible, and so let $x(\cdot)$ and $u(\cdot)$ satisfy the following conditions for all non-negative t:

$$u(t) \in \mathcal{U}_{\bar{p}}, \; 1^T u(t) = P^r(t),$$
$$\dot{x}(t) = -P^{-1}u(t), \; x_i(t) \geq 0, \quad i = 1, 2, ..., n. \quad (A.3)$$

Moreover, due to the inability of devices to charge, we can utilise the fact that $x(t) \geq 0 \;\; \forall t \iff x(T) \geq 0$ for an expression equivalent to the last condition. Now, consider permuting the input $u(\cdot)$ in time, corresponding to the permutation of the reference, to produce $\tilde{u}(\cdot)$, i.e.

$$\tilde{u}(\cdot) \doteq u_{[t_{\psi_1},t_{\psi_1+1})} \otimes u_{[t_{\psi_2},t_{\psi_2+1})} \otimes ... \otimes u_{[t_{\psi_v},t_{\psi_v+1})}, \quad (A.4)$$

in which $u_{[t_1,t_2)}$ denotes the signal $u(\cdot)$ truncated to the half-open interval $[t_1, t_2)$. Consider each condition of the feasibility of $\tilde{P}^r(\cdot)$ individually as follows. By construction, for all non-negative t,

$$\tilde{u}(t) \in \mathcal{U}_{\bar{p}},$$
$$1^T \tilde{u}(t) = \tilde{P}^r(t), \quad (A.5)$$
$$\dot{\tilde{x}}(t) = -P^{-1}\tilde{u}(t),$$

in which $\tilde{x}(\cdot)$ denotes the trajectory that fulfils the permuted reference and is initialised as $\tilde{x}(0) = x(0)$. In addition,

$$\tilde{x}(T) = \tilde{x}(0) - \int_0^T P^{-1}\tilde{u}(\tau)d\tau$$
$$= x(0) - \int_0^T P^{-1}u(\tau)d\tau = x(T) \geq 0. \quad (A.6)$$

Therefore $\tilde{P}^r(\cdot)$ satisfies all the necessary conditions and we are able to deduce that $P^r(\cdot) \in \mathcal{F}_{\bar{p},x} \iff \tilde{P}^r(\cdot) \in \mathcal{F}_{\bar{p},x}$, in which the condition is necessary as well as sufficient as a result of its symmetry. $\blacksquare$

B. Proof of Lemma III.2

Let the system under consideration have initial state x and maximum power vector $\bar{p}$. Then consider partitioning a feasible piecewise constant reference signal $P^r(\cdot)$, of length T, into the partitions before and after time t, i.e. $P^r(\cdot) = P^r_{[0,t)} \otimes P^r_{[t,T)}$. Now form a new signal as $\hat{P}^r(\cdot) \doteq \tilde{P}^r_{[0,t)} \otimes P^r_{[t,T)}$, in which $\tilde{P}^r_{[0,t)}$ is a permuted version of $P^r_{[0,t)}$. Using the result of Lemma III.1, we know that if $P^r(\cdot)$ is feasible then so too must be $\hat{P}^r(\cdot)$.

Denote by z and $\tilde{z}$ the states reached at time t under the application of the optimal policy to meet $P^r(\cdot)$ and $\hat{P}^r(\cdot)$ respectively. Without loss of generality, let the vector x be ordered in decreasing time-to-go, and let the same device ordering be used for z and $\tilde{z}$. Then, because ordering is preserved under the optimal policy, z and $\tilde{z}$ will also be in descending order. From the above result, we are able to say that $P^r_{[t,T)} \in \mathcal{F}_{\bar{p},z}$ and $P^r_{[t,T)} \in \mathcal{F}_{\bar{p},\hat{z}}$. As this holds for any feasible choice of reference $P^r(\cdot)$, we are then able to say that $\mathcal{F}_{\bar{p},\hat{z}} = \mathcal{F}_{\bar{p},z}$. We use this to prove the required result by contradiction as follows. Consider firstly the case in which $z_n > \tilde{z}_n$, in which n indexes the device with the smallest time-to-go value. We are able to construct the following reference signal:

$$P^r(t) = \begin{cases} \sum_{i=1}^n \bar{p}_i, & \text{if } t \in [0, z_n) \\ 0, & \text{otherwise}, \end{cases} \quad (A.7)$$

for which we can trivially see that $P^r(\cdot) \in \mathcal{F}_{\bar{p},z}$ and $P^r(\cdot) \notin \mathcal{F}_{\bar{p},\tilde{z}}$, hence the feasible sets defined by the two states must be distinct.

Now, consider an arbitrary device index j and assume that

$$z_i = \tilde{z}_i, \; i = n, n-1, ..., j+1, \quad (A.8a)$$
$$z_j > \tilde{z}_j. \quad (A.8b)$$

We are able to construct the following reference signal:

$$P^r(t) = \begin{cases} \sum_{i=1}^n \bar{p}_i, & \text{if } t \in [0, z_n) \\ \sum_{i=1}^{n-1} \bar{p}_i, & \text{if } t \in [z_n, z_{n-1}) \\ \vdots \\ \sum_{i=1}^j \bar{p}_i, & \text{if } t \in [z_{j+1}, z_j) \\ 0, & \text{otherwise}, \end{cases} \quad (A.9)$$

for which we can trivially see that $P^r(\cdot) \in \mathcal{F}_{\bar{p},z}$ and $P^r(\cdot) \notin \mathcal{F}_{\bar{p},\tilde{z}}$, hence the feasible sets defined by the two states must be distinct. Induction of this argument from device index n to 1 returns the result that the feasible sets must be distinct whenever $\exists k: z_k > \tilde{z}_k$. In addition, due to the arbitrary allocation of z and $\tilde{z}$, this result must hold whenever $z \neq \tilde{z}$. Hence the permutation must not alter the terminal state. $\blacksquare$

C. Proof of Lemma III.5

Consider two arbitrary power levels, p and q for which $p \geq q$. These must satisfy

$$\max\{P^r(\tau) - p, 0\} \leq \max\{P^r(\tau) - q, 0\} \;\; \forall \tau, \quad (A.10)$$

and so $E_{P^r}(p) \leq E_{P^r}(q)$, i.e. the E-p curve is monotone. In addition, the definition of the E-p transform leads to the following left- and right-derivatives respectively:

$$\frac{dE(p)}{dp}^{-} = -\mu\big(\{\tau\colon P^r(\tau) \geq p\}\big), \qquad \text{(A.11a)}$$

$$\frac{dE(p)}{dp}^{+} = -\mu\big(\{\tau\colon P^r(\tau) > p\}\big), \qquad \text{(A.11b)}$$

in which $\mu(\cdot)$ denotes the Lebesgue measure operator. Hence

$$0 \geq \frac{dE(p)}{dp}\bigg|_{p} \geq \frac{dE(p)}{dp}\bigg|_{q}, \qquad \text{(A.12)}$$

thus the gradient is negative and monotonically increasing, and so the curve must be convex. ∎

D. Proof of Theorem III.8

For this proof we initially form the following supporting Lemmas:

Lemma VII.1. *If a **piecewise constant** reference signal $P^r(\cdot)$ is feasible, its E-p transform is dominated by the capacity of the system, i.e. $P^r(\cdot) \in \mathcal{F}_{\bar{p},x} \implies E_{P^r}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p$.*

Proof. Firstly, the result of Lemma III.1, combined with the fact that $P^r(\cdot)$ is both piecewise constant and feasible, allows us to form a non-increasing equivalent to $P^r(\cdot)$, which we denote $\tilde{P}^r(\cdot)$ and take to be the reference of interest for the remainder of this proof.

We then compose a framework as follows. Consider the partition induced at the initial time by the optimal policy, and let $\mathcal{N}_i$, $i = 1, 2, ..., l$ be the corresponding subsets of devices. Making use of these subsets, then form the worst-case reference of the system, $R(\cdot)$ as in (14), so that this can be compared to $\tilde{P}^r(\cdot)$. We note that $R(\cdot)$ is bang-bang in all components of the input (i.e. it runs them at full power from time $t = 0$ up to depletion). Denoting by $\tilde{E}(\cdot)$ and $E^*(\cdot)$ the energy vectors resulting from implementation of the optimal policy to meet references $\tilde{P}^r(\cdot)$ and $R(\cdot)$ respectively, we are therefore able to say that $\tilde{E}(t) \geq E^*(t) \ \forall t \geq 0$. Moreover, since $\tilde{P}^r(\cdot)$ and $R(\cdot)$ are both feasible,

$$1^T \dot{\tilde{E}}(t) = -\tilde{P}^r(t) \quad \forall t \geq 0, \qquad \text{(A.13)}$$

$$1^T \dot{E}^*(t) = -R(t) \quad \forall t \geq 0. \qquad \text{(A.14)}$$

Now, consider any power level p, and define the following:

$$t_P \doteq \inf\{t \geq 0\colon \tilde{P}^r(t) \leq p\}, \qquad \text{(A.15)}$$

$$t_R \doteq \inf\{t \geq 0\colon R(t) \leq p\}, \qquad \text{(A.16)}$$

which exist finite. We are then directly able to make the following comparison:

$$
\begin{aligned}
E_{\tilde{P}^r}(p) &= \int_0^{t_P} [-1^T \dot{\tilde{E}}(t) - p]dt \\
&= 1^T \tilde{E}(0) - 1^T \tilde{E}(t_P) - pt_P \\
&\leq 1^T E^*(0) - 1^T E^*(t_P) - pt_P \\
&= \int_0^{t_P} [-1^T \dot{E}^*(t) - p]dt \\
&\leq \int_0^{t_R} [R(t) - p]dt = \Omega_{\bar{p},x}(p),
\end{aligned}
\qquad \text{(A.17)}
$$

where the final inequality follows from consideration of the cases $t_P < t_R$, $t_P = t_R$ and $t_P > t_R$. ∎

Lemma VII.2. *Given two reference signals $P^r(\cdot)$ and $\tilde{P}^r(\cdot)$, if*

$$P^r(t) \geq \tilde{P}^r(t) \geq 0 \quad \forall t \qquad \text{(A.18)}$$

and $P^r(\cdot)$ is feasible, then $\tilde{P}^r(\cdot)$ is also feasible.

Proof. Since $P^r(\cdot)$ is feasible, there exists a $u(\cdot)$ induced by this reference which satisfies all necessary conditions. Consider now constructing an input $\tilde{u}(\cdot)$ simply as

$$\tilde{u}(t) = \begin{cases} \frac{\tilde{P}^r(t)}{P^r(t)} u(t), & \text{if } P^r(t) > 0 \\ 0, & \text{otherwise.} \end{cases} \qquad \text{(A.19)}$$

Noting firstly that (A.18) leads to $P^r(t) = 0 \implies \tilde{P}^r(t) = 0$, we are able to see by construction that $1^T \tilde{u}(t) = \tilde{P}^r(t) \ \forall t$. In addition the feasibility of $P^r(\cdot)$ combined with $0 \leq \tilde{u}(t) \leq u(t) \ \forall t$ leads to the result that $\tilde{u}(\cdot)$ must meet all the other necessary conditions for feasibility as in Definition II.1. ∎

Lemma VII.3. *If a **piecewise continuous** signal $P^r(\cdot)$ is feasible, its E-p transform is dominated by the capacity of the system, i.e. $P^r(\cdot) \in \mathcal{F}_{\bar{p},x} \implies E_{P^r}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p$.*

Proof. Given $P^r(\cdot)$, consider constructing k piecewise constant signals, $\tilde{P}^{r,i}(\cdot)$, $i = 1, ..., k$, the sum of which approximates from below the original signal, i.e.

$$P^{r,k}(t) \doteq \sum_{i=1}^{k} \tilde{P}^{r,i}(t) \leq P^r(t) \quad \forall t, \qquad \text{(A.20a)}$$

$$\lim_{k \to +\infty} \left\| \sum_{i=1}^{k} \tilde{P}^{r,i} - P^r(t) \right\|_{\infty} = 0. \qquad \text{(A.20b)}$$

We know from the result of Lemma VII.2 that $P^{r,k}(\cdot)$ is feasible, and in particular that $E_{P^{r,k}}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p, k$. Letting $k \to \infty$, we then see that $E_{P^r}(p) = \lim_{k \to \infty} E_{P^{r,k}}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p$. ∎

Lemma VII.4. *If the E-p transform of a reference signal $P^r(\cdot)$ is dominated by the capacity of the system, the signal is feasible, i.e.*

$$E_{P^r}(p) \leq \Omega_{\bar{p},x}(p) \ \forall p \implies P^r(\cdot) \in \mathcal{F}_{\bar{p},x}. \qquad \text{(A.21)}$$

Proof. For a reference signal to be feasible there must exist an input that meets all power and energy constraints as in Definition II.1; to prove the claim it is sufficient to find such an input signal. We apply the optimal feedback (11) and construct accordingly our chosen input, which meets all power constraints by construction. All that remains is to show that it does not result in negative energy stored in any device, which we do as follows.

Consider the partition induced by the optimal policy at a state $x = x(t)$. Let $\mathcal{N}_1(x), \mathcal{N}_2(x), ..., \mathcal{N}_{q(x)}(x)$, in which $q(x)$ is the number of unique values in x, be the corresponding subsets ordered by decreasing time-to-go. Additionally denote the maximum power vectors corresponding to each subset as

$\bar{U}_1(x), \bar{U}_2(x), ..., \bar{U}_{q(x)}(x)$. Now, consider the power level s defined according to the final state as follows:

$$s \doteq \sum_{i=1}^{q(x(\infty))-1} 1^T \bar{U}_i(x(\infty)). \qquad (A.22)$$

This power level corresponds to a summation across all devices except members of the subset with (equal) smallest time-to-go value at the final time. We choose this level because the last subset of devices would be the first to be pushed negative under the optimal policy; if this subset has non-negative time-to-go then all devices must have non-negative time-to-go.

Now, we know that the subsets of devices are monotonically expanding in time. Hence the final subsets are well-defined and in addition are composed of subsets formed at any previous time instant. Thus we know that the power level s corresponds to a summation across one or more subsets at any time along the considered solution. We define

$$l\colon x(t) \mapsto k, \text{ such that } \sum_{i=1}^{k} 1^T \bar{U}_i(x(t)) = s. \qquad (A.23)$$

Consider now the worst-case reference from an arbitrary starting state x. We are able to write this and the corresponding capacity explicitly as a function of the condensed state vector X as follows:

$$R(t, X) = \sum_{j=1}^{q(x)} \left[H(t) - H(t - x_{\mathcal{N}_j}) \right] 1^T \bar{U}_j(x), \qquad (A.24a)$$

$$\Omega(p, X) = \int_0^{\infty} \max\{R(t, X) - p, 0\} dt, \qquad (A.24b)$$

in which $H(\cdot)$ denotes the Heaviside step function. This expression is only valid for x non-negative, however we extend it to arbitrary x as follows. We invert the correspondence of the worst-case reference to give time as a function of power:

$$t(p, X) = \begin{cases} \infty, & \text{if } p = 0 \\ x_{\mathcal{N}_1}, & \text{if } 0 < p \leq 1^T \bar{U}_1(x) \\ x_{\mathcal{N}_2}, & \text{if } 1^T \bar{U}_1(x) < p \leq \sum_{i=1}^{2} 1^T \bar{U}_i(x) \\ \vdots & \\ x_{\mathcal{N}_{q(x)}}, & \text{if } \sum_{i=1}^{q(x)-1} 1^T \bar{U}_i(x) < p \leq p_{max} \\ 0, & \text{otherwise,} \end{cases} \qquad (A.25)$$

in which $p_{max} \doteq \sum_{i=1}^{q(x)} 1^T \bar{U}_i(x) = \sum_{j=1}^{n} \bar{p}_j$, and we then redefine the capacity as an integration along the p-axis:

$$\Omega(p, X) \doteq \int_p^{p_{max}} t(p', X) dp'. \qquad (A.26)$$

Given this generalised definition of the capacity, consider taking partial derivatives with respect to each distinct state, evaluated at the power level s (recall that this is independent of the current state). This leads to

$$\frac{\partial \Omega}{\partial x_{\mathcal{N}_i}}\bigg|_s = \begin{cases} 1^T \bar{U}_i(x), & \text{if } i > l(x) \\ 0, & \text{otherwise.} \end{cases} \qquad (A.27)$$

Now, denoting by $m(x)$ the highest group index of devices running at positive power, we know that the dynamics corresponding to the optimal policy (11) are

$$\dot{x}_{\mathcal{N}_i} = \begin{cases} -1, & \text{if } i < m(x) \\ -r_i, & \text{if } i = m(x) \\ 0, & \text{if } i > m(x), \end{cases} \qquad (A.28a)$$

in which

$$r_i = \frac{P^r - \sum_{j=1}^{i-1} 1^T \bar{U}_j(x)}{1^T \bar{U}_i(x)}. \qquad (A.28b)$$

We are then able to consider the 6 possible cases that arise for the following product for each subset $\mathcal{N}_i$ and at the power level s:

$$\dot{x}_{\mathcal{N}_i} \cdot \frac{\partial \Omega}{\partial x_{\mathcal{N}_i}}\bigg|_s, \qquad (A.29)$$

as seen in Table I.

	$i \leq l(x)$	$i > l(x)$
$i < m(x)$	0	$-1^T \bar{U}_i(x)$
$i = m(x)$	0	$-r_i 1^T \bar{U}_i(x)$
$i > m(x)$	0	0

TABLE I: The 6 possible cases for the expression of (A.29).

As it can be seen from the table, there are only 2 cases in which the expression of (A.29) is non-zero: when $l(x) < i \leq m(x)$. These can be converted back to power conditions as follows:

$$\dot{x}_{\mathcal{N}_i} \cdot \frac{\partial \Omega}{\partial x_{\mathcal{N}_i}}\bigg|_s$$
$$= \begin{cases} -1^T \bar{U}_i(x), & \text{if } s < \sum_{j=1}^{i} 1^T \bar{U}_j(x) \leq P^r \\ \sum_{j=1}^{i-1} 1^T \bar{U}_j(x) - P^r, & \text{if } s < P^r < \sum_{j=1}^{i} 1^T \bar{U}_j(x) \\ 0, & \text{otherwise.} \end{cases} \qquad (A.30)$$

Utilising this expression, we can then use the chain rule to form the partial derivative with respect to time of the capacity, evaluated at power level s and state x, by summing across subsets. This operation is valid because each subset of devices can in effect be aggregated into a single virtual device, and leads to

$$\frac{\partial \Omega}{\partial t}\bigg|_{s,x} = \sum_{i=1}^{q(x)} \dot{x}_{\mathcal{N}_i} \cdot \frac{\partial \Omega}{\partial x_{\mathcal{N}_i}}\bigg|_s$$
$$\stackrel{(\dagger)}{=} -\sum_{i=l(x)+1}^{q(x)} \max\left\{0, \min\left\{P^r - \sum_{j=1}^{i-1} 1^T \bar{U}_j(x), 1^T \bar{U}_i(x)\right\}\right\}$$
$$= \sum_{i=1}^{l(x)} \max\left\{0, \min\left\{P^r - \sum_{j=1}^{i-1} 1^T \bar{U}_j(x), 1^T \bar{U}_i(x)\right\}\right\}$$
$$- \sum_{i=1}^{q(x)} \max\left\{0, \min\left\{P^r - \sum_{j=1}^{i-1} 1^T \bar{U}_j(x), 1^T \bar{U}_i(x)\right\}\right\}$$
$$\stackrel{(\ddagger)}{=} \min\{P^r, s\} - P^r = -\max\{P^r - s, 0\}, \qquad (A.31)$$

in which the equality $(\dagger)$ results from the removal of the cases for which $l(x) \geq i$ leading to a product of 0, and the equality $(\ddagger)$ results from the following arguments. The negative summation can be interpreted as the negative of the

aggregate power across all devices, which must equal P^r as a result of the construction of the input according to the optimal policy. The positive summation is the same series curtailed at the index corresponding to the level s; if $P^r > s$ then all devices up to this level will be run at full power and their aggregate will equal s, otherwise the aggregate will once again be the total power P^r. Integration of (A.31) gives

$$\Omega\big(s, x(\infty)\big) = \Omega\big(s, x(0)\big) - \int_0^\infty \max\big\{P^r(t) - s, 0\big\}dt, \tag{A.32}$$

which coupled with the condition (A.21) leads directly to $\Omega\big(s, x(\infty)\big) \geq 0$. Moreover,

$$\Omega\big(s, x(\infty)\big) = x_{\mathcal{N}_{q(x(\infty))}}(\infty) \cdot 1^T \bar{U}_{q(x(\infty))}\big(x(\infty)\big), \tag{A.33}$$

therefore $x_{\mathcal{N}_{q(x(\infty))}}(\infty) \geq 0$. As the time-to-go values cannot increase over time the reference must therefore be feasible. ∎

Theorem III.8 then follows directly from Lemmas VII.3 and VII.4. ∎

References

[1] "E-storage: Shifting from Cost to Value," World Energy Council, Tech. Rep., 2016.

[2] "Load Participation and Demand Response," CAISO, 2017, http://www.caiso.com/participate/Pages/Load/Default.aspx.

[3] "Ancillary Services," PJM, 2017, https://www.pjm.com/markets-and-operations/ancillary-services.aspx.

[4] S. H. Tindemans, V. Trovato, and G. Strbac, "Decentralized Control of Thermostatic Loads for Flexible Demand Response," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 5, pp. 1685–1700, 2015.

[5] A. Busic and S. Meyn, "Distributed randomized control for demand dispatch," *2016 IEEE 55th Conference on Decision and Control, CDC 2016*, pp. 6964–6971, 2016.

[6] D. Métivier and M. Chertkov, "Mean Field Control for Efficient Mixing of Energy Loads," *Preprint arXiv:1810.00450*, pp. 1–17, 2018.

[7] S. CHEN and C.-C. LIU, "From demand response to transactive energy: state of the art," *Journal of Modern Power Systems and Clean Energy*, vol. 5, no. 1, pp. 10–19, 2017.

[8] H. S. Nunna and D. Srinivasan, "Multiagent-Based Transactive Energy Framework for Distribution Systems with Smart Microgrids," *IEEE Transactions on Industrial Informatics*, vol. 13, no. 5, pp. 2241–2250, 2017.

[9] J. Hu, G. Yang, H. W. Bindner, and Y. Xue, "Application of Network-Constrained Transactive Control to Electric Vehicle Charging for Secure Grid Operation," *IEEE Transactions on Sustainable Energy*, vol. 8, no. 2, pp. 505–515, 2017.

[10] Z. Liu, Q. Wu, K. Ma, M. Shahidehpour, Y. Xue, and S. Huang, "Two-Stage Optimal Scheduling of Electric Vehicle Charging based on Transactive Control," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, pp. 1–11, 2018.

[11] T. Alharbi and K. Bhattacharya, "Optimal scheduling of energy resources and management of loads in isolated / islanded microgrids," *Canadian Journal of Electrical and Computer Engineering*, vol. 40, no. 4, pp. 284–295, 2017.

[12] N. G. Paterakis, O. Erdinc, I. N. Pappi, A. G. Bakirtzis, and J. P. Catalao, "Coordinated Operation of a Neighborhood of Smart Households Comprising Electric Vehicles, Energy Storage and Distributed Generation," *IEEE Transactions on Smart Grid*, vol. 7, no. 6, pp. 2736–2747, 2016.

[13] W. Alharbi and K. Raahemifar, "Probabilistic coordination of microgrid energy resources operation considering uncertainties," *Electric Power Systems Research*, vol. 128, pp. 1–10, 2015.

[14] C. Battistelli, L. Baringo, and A. J. Conejo, "Optimal energy management of small electric energy systems including V2G facilities and renewable energy sources," *Electric Power Systems Research*, vol. 92, pp. 50–59, 2012.

[15] X. Bai and W. Qiao, "Robust Optimization for Bidirectional Dispatch Coordination of Large-Scale V2G," *IEEE Transactions on Smart Grid*, vol. 6, no. 4, pp. 1944–1954, 2015.

[16] A. Y. Saber and G. K. Venayagamoorthy, "Resource scheduling under uncertainty in a smart grid with renewables and plug-in vehicles," *IEEE Systems Journal*, vol. 6, no. 1, pp. 103–109, 2012.

[17] M. Alonso, H. Amaris, J. G. Germain, and J. M. Galan, "Optimal charging scheduling of electric vehicles in smart grids by heuristic algorithms," *Energies*, vol. 7, no. 4, pp. 2449–2475, 2014.

[18] L. Gan, U. Topcu, and S. H. Low, "Optimal Decentralized Protocols for Electric Vehicle Charging," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 940–951, 2013.

[19] V. Trovato, S. H. Tindemans, and G. Strbac, "Leaky storage model for optimal multi-service allocation of thermostatic loads," *IET Generation, Transmission & Distribution*, vol. 10, no. 3, pp. 585–593, 2016.

[20] A. I. Bejan, R. J. Gibbens, and F. P. Kelly, "Statistical aspects of storage systems modelling in energy networks," *2012 46th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6, 2012.

[21] N. Gast, D. C. Tomozei, and J. Y. Le Boudec, "Optimal generation and storage scheduling in the presence of renewable forecast uncertainties," *IEEE Transactions on Smart Grid*, vol. 5, no. 3, pp. 1328–1339, 2014.

[22] J. R. Cruise, R. J. Gibbens, and S. Zachary, "Optimal control of storage for arbitrage, with applications to energy systems," *2014 48th Annual Conference on Information Sciences and Systems, CISS 2014*, pp. 1–6, 2014.

[23] G. Edwards, S. Sheehy, C. J. Dent, and M. C. Troffaes, "Assessing the contribution of nightly rechargeable grid-scale storage to generation capacity adequacy," *Sustainable Energy, Grids and Networks*, vol. 12, pp. 69–81, 2017.

[24] M. Evans, S. H. Tindemans, and D. Angeli, "Robustly Maximal Utilisation of Energy-Constrained Distributed Resources," *2018 Power Systems Computation Conference (PSCC)*, pp. 1–7, 2018.

[25] D. S. Callaway and I. A. Hiskens, "Achieving Controllability of Electric Loads," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, 2010.

[26] S. Shao, M. Pipattanasomporn, and S. Rahman, "Grid integration of electric vehicles and demand response with customer choice," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 543–550, 2012.

[27] G. Wenzel, M. Negrete-Pincetic, D. E. Olivares, J. MacDonald, and D. S. Callaway, "Real-Time Charging Strategies for an Electric Vehicle Aggregator to Provide Ancillary Services," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, pp. 1–11, 2017.

[28] K. M. Tsui and S. C. Chan, "Demand response optimization for smart home scheduling under real-time pricing," *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 1812–1821, 2012.

[29] Z. Wang, C. Gu, F. Li, P. Bale, and H. Sun, "Active demand response using shared energy storage for household energy management," *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 1888–1897, 2013.

[30] M. Mazidi, A. Zakariazadeh, S. Jadid, and P. Siano, "Integrated scheduling of renewable generation and demand response programs in a microgrid," *Energy Conversion and Management*, vol. 86, pp. 1118–1127, 2014.

[31] X. Wang, A. Palazoglu, and N. H. El-Farra, "Operational optimization and demand response of hybrid renewable energy systems," *Applied Energy*, vol. 143, pp. 324–335, 2015.

[32] A. De Paola, D. Angeli, and G. Strbac, "A Semi-Decentralized Scheme for Integration of Price-Responsive Appliances in the Electricity Market," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6729–6736, 2017.

Michael Evans received the B.A. and M.Eng. degrees in Engineering from the University of Cambridge, UK, in 2015. He is currently pursuing the Ph.D. degree in the Department of Electrical and Electronic Engineering at Imperial College London, UK.

Simon Tindemans (M'13) received the M.Sc. degree in physics from the University of Amsterdam in 2004, and the Ph.D. degree from Wageningen University / AMOLF in 2009. He was a Research Fellow with Imperial College London (2011-2017) and is currently an Assistant Professor with the Department of Electrical Sustainable Energy at Delft University of Technology, the Netherlands.

David Angeli (F'15) was born in Siena, Italy, in 1971. He received the B.S. degree in Computer Science Engineering and the Ph.D. in Control Theory from University of Florence, Italy, in 1996 and 2000, respectively. Since 2000 he was an Assistant and Associate Professor (2005) with the Department of Systems and Computer Science, University of Florence. In 2008 he joined, as a Senior Lecturer, the Department of Electrical and Electronic Engineering of Imperial College London, where he is currently a Reader in Nonlinear Control and the Director of the MSc in Control Systems. He has been a Fellow of the IEEE since 2015 and of the IET since 2018. Overall he has authored more than 100 journal papers in the areas of Stability of nonlinear systems, Control of constrained systems (MPC), Chemical Reaction Networks theory and Smart Grids.