Chance-Constrained Ancillary Service Specification for Heterogeneous Storage Devices

Evans, Michael; Angeli, David; Strbac, Goran; Tindemans, Simon H.

DOI

[10.1109/ISGTEurope.2019.8905723](https://doi.org/10.1109/ISGTEurope.2019.8905723)

Publication date

2019

Document Version

Final published version

Published in

Proceedings of 2019 IEEE PES Innovative Smart Grid Technologies Europe, ISGT-Europe 2019

Citation (APA)

Evans, M., Angeli, D., Strbac, G., & Tindemans, S. H. (2019). Chance-Constrained Ancillary Service Specification for Heterogeneous Storage Devices. In *Proceedings of 2019 IEEE PES Innovative Smart Grid Technologies Europe, ISGT-Europe 2019* (pp. 1-5). [8905723] (Proceedings of 2019 IEEE PES Innovative Smart Grid Technologies Europe, ISGT-Europe 2019). IEEE .
<https://doi.org/10.1109/ISGTEurope.2019.8905723>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Chance-Constrained Ancillary Service Specification for Heterogeneous Storage Devices

Michael P. Evans, David Angeli* and Goran Strbac
Department of Electrical and Electronic Engineering
Imperial College London, UK
{m.evans16, d.angeli, g.strbac}@imperial.ac.uk

Simon H. Tindemans
Department of Electrical Sustainable Energy
TU Delft, Netherlands
s.h.tindemans@tudelft.nl

Abstract—We present a method to find the maximum magnitude of any supply-shortfall service that an aggregator of energy storage devices is able to sell to a grid operator. This is first demonstrated in deterministic settings, then applied to scenarios in which device availabilities are stochastic. In this case we implement chance constraints on the inability to deliver as promised. We show a significant computational improvement in using our method in place of straightforward scenario simulation. As an extension, we present an approximation to this method which allows the determined fleet capability to be applied to any chosen service, rather than having to re-solve the chance-constrained optimisation each time.

Index Terms—Energy storage systems, chance-constrained optimisation, optimal control, ancillary service

I. INTRODUCTION

Recent years have seen a significant increase in the proliferation of energy storage devices onto electricity networks. This is in large part due to their potential to replace conventional generation as electricity grids are decarbonised [1]. As part of this transition, increasing numbers of system operators offer tenders for the provision of ancillary services; in this paper we focus on balancing services, although our analysis would be applicable to other active power services as well. The majority of these services are specified in terms of time durations over which a ramp in power or a constant power output should be delivered. When considered for provision by storage, they are therefore often taken to be shaped delivery requirements, as in [2]. We place our analysis within such a framework, and in particular consider request profiles in the shape of pulses or trapezoids. We point out here, however, that our analysis is general and not limited to any particular shape.

Many prior studies have considered the application of storage to the provision of ancillary services. A particularly promising field is that of vehicle-to-grid (V2G), in which electric vehicles (EVs) are able to provide system support when connected to the network [3]–[6]. This has been investigated for a range of services, including spinning reserve [3], [4], regulation [4], [5] and reactive power compensation [6]. Crucially, the ability for an EV to participate is subject to its availability, which tends to be assumed stochastic. We incorporate this form of stochasticity into our analysis.

*Dr David Angeli is also with the Department of Information Engineering, University of Florence, Italy

Moreover, in these settings a common requirement is that the provision of an ancillary service must not reduce the quality of the primary service (e.g. transportation) below some user-defined threshold. A similar framework might apply to uninterruptible power supplies (UPSs) that are kept as a backup in case of power outage. Other device configurations instead include system service provision within their core tasks, as is common among home energy management systems including storage [7]. Such a system might, for example, be designed so as to trade off supporting the grid against maximising self-consumption of electricity generated from rooftop solar panels [8]. Wu *et al.* [9] combine the strands of V2G and home energy management by considering a system in which the only storage device is the household EV.

It is clear from these applications that a means to determine *a priori* the maximum service provision capability of a fleet of heterogeneous storage devices would be a valuable asset to an aggregator. This would allow them to determine the magnitude of a given shape of delivery requirement that they could offer to the system operator. When considering a single storage device, it is trivial to deduce its maximum ancillary service capability, simply by considering the power and energy limits of the device. The same cannot be said of a heterogeneous fleet of devices, however. In our prior work [10], we presented a transform through which one can represent the aggregate flexibility of a device pool, termed the *E-p* transform. An alternative aggregate fleet representation was developed independently by Cruise and Zachary [11]. For an aggregate representation of this form to exist requires that the problem be restricted to one of discharge only, as was discussed in detail in [10].

This paper explores in greater depth how the *E-p* transform can be used to specify ancillary service capability. In particular, we move beyond deterministic settings. Chance constraints are commonly applied to the dispatch of storage under uncertainty [12]–[14], and we utilise these to ensure that the probability of unsuccessful delivery of the procured service lies below a user-defined threshold. This involves a Monte Carlo simulation in which the *E-p* transform of [10] is used as a binary feasibility check on each drawn trace. We point out here that an alternative method for such a check would be a procedural simulation in time spanning the total duration of the delivery requirement. We therefore investigate the speed-

up achieved in using the E - p transform for such a check. In prior work we presented a policy, in both continuous [15] and discrete [16] time, with a one-to-one mapping between feasible requests and those that will be met under the implementation of the policy. It will therefore be this that we will use to perform a procedural check.

II. MATHEMATICAL FORMULATION

A. Problem description

We denote the constituent units of the heterogeneous fleet as D_i , $i = 1, \dots, n$, each with extractable energy $e_i(t)$ subject to the assumed physical constraint $e_i(t) \geq 0$. We choose as our control input $u_i(t)$ the power extracted from each device, measured externally so as to take into account any inefficiencies present, leading to the integrator dynamics $\dot{e}_i(t) = -u_i(t)$. For the reasons mentioned in the Introduction, we restrict ourselves to discharging operation, with device-specific limit $\bar{p}_i$, so that $u_i(t) \in [0, \bar{p}_i]$. We choose as our state variable the *time-to-go*, defined for each device as $x_i(t) \doteq e_i(t)/\bar{p}_i$. We then stack state, input and maximum power values across devices as follows:

$$x(t) \doteq [x_1(t) \ \dots \ x_n(t)]^T, \quad (1)$$

$$u(t) \doteq [u_1(t) \ \dots \ u_n(t)]^T, \quad (2)$$

$$\bar{p} \doteq [\bar{p}_1 \ \dots \ \bar{p}_n]^T, \quad (3)$$

allowing us to rewrite the dynamics in matrix form as $\dot{x}(t) = -P^{-1}u(t)$, in which $P \doteq \text{diag}(\bar{p})$. We also form the product set of the power constraints, $\mathcal{U}_{\bar{p}} \doteq [0, \bar{p}_1] \times [0, \bar{p}_2] \times \dots \times [0, \bar{p}_n]$, so that our constraints can be compactly written as $u(t) \in \mathcal{U}_{\bar{p}}$ and $x(t) \geq 0$. We will be interested in the ability of the fleet to meet a request profile, which we denote by $P^r: [0, +\infty) \mapsto [0, +\infty)$. We say that any profile which the fleet can satisfy without violating any constraints is *feasible*, and we define the set of such signals as follows:

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

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

This then allows us to transform a determination of request feasibility into an evaluation of membership of this set.

B. E - p transform

The main tool that we will use in our analysis is the E - p transform of [10]. We therefore reproduce the following definitions here:

Definition II.2. *Given a request profile $P^r: [0, +\infty) \rightarrow [0, +\infty)$, we define its E - p transform as the following function:*

$$E_{P^r}(p) \doteq \int_0^\infty \max\{P^r(t) - p, 0\} dt, \quad (4)$$

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

Definition II.3. *We define the **capacity** of a system to be the E - p transform of the worst-case request profile that it can meet, i.e. $\Omega_{\bar{p},x}(p) \doteq E_R(p)$, where $R(\cdot)$ is defined as*

$$R(t) \doteq \sum_{i=1}^n \bar{p}_i [H(t) - H(t - x_i)], \quad (5)$$

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

Property II.4. *A request profile $P^r(\cdot)$ is feasible if, and only if, its E - p transform is dominated by the capacity curve 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 property can be interpreted to mean that the capacity curve represents the full set of request profiles that a fleet can satisfy. As discussed in detail in [10], one can therefore use this curve to determine the maximum ancillary service capability of the fleet.

In this paper, we use the following bisection method to find the maximum feasible service magnitude for a given capacity curve. Start with the interval $[0, \sum_{i=1}^n \bar{p}_i]$, which is known to contain the maximum feasible magnitude. Then, continually bisect and update this interval until its width is less than a chosen tolerance. Finally, as we are interested in the largest feasible magnitude, take as the solution the lower bound.

III. APPLICATION TO STOCHASTIC SETTINGS

A. Problem of interest

We take the point of view of an aggregator that wants to sell grid support services to a system operator. More specifically, we restrict these to balancing services under supply-shortfall conditions, so that our delivery requirements are strictly positive. We consider the setting in which devices have stochastic availability, which results in a random initial state vector and therefore a random capacity curve. We investigate the scenario in which an aggregator offers a service to a grid operator of a predefined shape and duration, but where they are free to specify the magnitude, m , of that service; and they are incentivised to offer as large a value as they can. We therefore aim to find the maximum such magnitude that the aggregator can sell, subject to a chance constraint on feasibility. We model their risk aversion via a level c , encoding the maximum probability of being unable to supply the contracted service that they are willing to accept. Their task, for a given $P^r(\cdot)$, is then to solve the following chance-constrained optimisation problem:

$$\begin{aligned} \max_m \quad & mP^r(\cdot) \\ \text{s.t.} \quad & \Pr[mP^r(\cdot) \in \mathcal{F}_{\bar{p},x}] \geq 1 - c, \end{aligned} \quad (*)$$

in which $\Pr[\cdot]$ denotes the probability operator. We consider scenarios in which a probability distribution of device availabilities is known. Despite this, however, we do not have an analytical expression for the probability distribution of interest. Instead, we are able to sample from this distribution, and we implement Monte Carlo sampling to find an estimate of the solution to $(*)$.

There are many scenarios to which this might apply in reality. For example, availability of EVs to participate in V2G

services is clearly limited by whether or not they happen to be plugged in when the request is received. Similarly, UPSs might be allowed to offer regulation services, but this would have to be curtailed in the event that they were called upon to deliver backup power. Under any such scenario with stochastic device availability, a chance-constrained capability assessment would tell the aggregator what magnitude of service provision to offer the system operator. We apply the E - p transform to this problem for the following reasons. Not only does it offer a speed-up over procedural simulation, as we will demonstrate, but it also offers intuition and insights into the factors affecting fleet capability. This enables us to find routes to computationally efficient approximations.

B. Monte Carlo procedure

As mentioned previously, for a given maximum power and state vector, we can find the maximum service capability of the fleet using its capacity curve. In our problem framework, the maximum power vector is predefined, but the state vector is allowed to vary based on device availability. We denote by $a_i \in \{0,1\}$ the availability of device D_i and stack availabilities across devices as $a \doteq [a_1, \dots, a_n]^T$. The sample-state corresponding to a given realised availability is then $x^s = a \circ x$, the Hadamard product of the availability and the (full-availability) state. For a given full-availability state, our Monte Carlo procedure utilises this as follows. For each sample, we draw an availability vector a and compute $\Omega_{\bar{p}, x^s}$. Using this, we then find the maximum feasible magnitude of the chosen service. We repeat this N times, then take the maximum magnitude which would have resulted in feasibility in $(1-c)N$ of the samples. This procedure can be seen diagrammatically in Figure 1.

Fig. 1: The Monte Carlo procedure

C. Quantile approximation

We here present an approximation to the above procedure, as follows. We draw N samples and each time form the corresponding capacity curve. We then determine the single curve such that $(1-c)N$ of the observed capacity curves lie

below it, point-wise in p . The single curve formed is then used as a deterministic capacity. This procedure can be seen in Figure 2. We point out here that this is an optimistic approximation, because pointwise dominance of $(1-c)N$ capacity curves is a less stringent condition than dominance over the whole domain of p . If this approximation is observed to be reasonably accurate, however, it might still be worthwhile. By producing a single curve, it would enable the determination of the maximum magnitude of *any* ancillary service, without needing to compose new samples each time. In addition, optimisation with respect to a quantile curve should be much simpler than optimisation with respect to an ensemble.

Fig. 2: The quantile approximation procedure

IV. NUMERICAL RESULTS

A. Scenario studied

A case study illustrates the procedures described above as follows. We choose the fleet size to be 500 devices. We draw maximum power and energy ratings for each device from device-independent log-normal distributions, based on the assumption that relative variation in these values across the fleet is more important than absolute variation. Inspired by typical EV parameters, we assign 450 (90%) of the devices to have maximum power values with a mean of 3.3 kW and a standard deviation of 1 kW, each of the log-normal distribution, and the remaining 50 devices to have a rating of 50 kW. This represents the partial access to rapid chargers, under the assumption that all chargers have V2G capability equal to their rated power. We then choose the current energy values of the (entire) fleet to have a mean of 40 kWh and a standard deviation of 10 kWh, again of the log-normal distribution.

B. Ancillary service specification procedure

Initially, we consider the capability of the fleet when all devices are fully available. In this case all parameter values are known explicitly and therefore the problem is deterministic. For this initial example we consider a service to deliver a simple pulse with a duration of 4 h, and find the maximum magnitude pulse of this form that the fleet can provide. The E - p curve corresponding to a signal of this type will be linear, with gradient equal to the negative duration and p -intercept

equal to the pulse magnitude. Therefore this task involves finding the straight line of gradient -4, with the maximum p -intercept out of those that are contained within the feasible region. This procedure can be seen in Figure 3a, with the corresponding pulses shown in Figure 3b.

(a) Use of the E - p transform to determine the maximum-magnitude 4 h pulse that the fleet can provide. The feasible region is shaded in green and the requests are represented by dashed lines.

(b) The three 4 h pulses 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. 3: Use of the E - p transform to determine the largest feasible pulse magnitude of 4 h duration. The colour of each E - p curve in Figure 3a matches that of the corresponding pulse signal in Figure 3b.

C. Trapezoidal delivery requirements

Having demonstrated how one can use the E - p transform to specify maximum ancillary service provision, we now perform this task on a less straightforward service. We consider a trapezoidal signal in which the total duration is split equally among ramping up at a fixed rate, maintaining constant power and ramping down to zero at the (negative) same fixed rate. We find the largest magnitude of such a signal that the fleet can provide, across a range of total durations. Note that, as in the previous example, the magnitude of this delivery requirement is given by the p -intercept of each curve. The results of this can be seen in Figure 4a, with the corresponding request profiles shown in Figure 4b.

D. Chance-constrained optimisation

We now consider the case in which device availability is stochastic, varying according to a known probability distribution. We assume an independent Bernoulli distribution for each device, with a 60% probability that the device is available. We implement a Monte Carlo simulation to determine the maximum magnitude 2 h trapezoidal signal that the fleet can now provide, across chance constraint probability levels, c ,

(a) Use of the E - p transform to determine the maximum-magnitude trapezoid that the fleet can provide, across a range of total durations. The feasible region is shaded in green and the requests are represented by dashed lines, with their total durations given in the figure legend.

(b) The trapezoidal requests considered.

Fig. 4: Use of the E - p transform to determine the largest-magnitude feasible trapezoid, across a range of total durations. The colour of each E - p curve in Figure 4a matches that of the corresponding pulse signal in Figure 4b.

of 50%, 10% and 1%. We apply the procedure of Figure 1, with the number of samples, N , set to 10^4 . The results of this experiment are shown in Table I and plotted as dashed E - p curves in Figure 5. 95% bootstrap confidence intervals were observed to be smaller than the significant digits shown.

E. Quantile approximation

We now investigate the approximation to the chance-constrained result described in Section III. We use the same fleet as in the previous experiment, including the same distribution on availabilities of devices, and we again set $N = 10^4$. The results of this can be seen in Table I, and the single-curve approximations are plotted as solid lines in Figure 5. Again, 95% bootstrap confidence intervals were observed to be smaller than the significant digits shown. Interestingly, this experiment returns a magnitude of provision that differs from the accurate value by less than 1%, suggesting that this quantile technique may be a worthwhile approximation to make. This is corroborated by Figure 5, where for each probability level there is a very good match (i.e. a small gap) between the accurate request E - p curves and the approximate capacity curves. We reiterate, however, that this estimate is not robust; we know from the preceding results that each approximated magnitude would lead to a greater probability of failure than the aggregator is willing to accept.

c value	Accurate magnitude (MW)	Approximated magnitude (MW)	Relative Error
50%	1.79	1.80	0.98%
10%	1.66	1.67	0.35%
1%	1.55	1.56	0.75%

TABLE I: The magnitude values determined by each procedure, and the approximation errors.

Fig. 5: Approximating the above results using a single quantile curve, across a range of probability levels, c (shown in the figure legend). In each case, the accurate maximum E - p curve is represented by a dashed line, with the approximate capacity curve represented by a solid line of the same colour.

F. Timing comparison

We now investigate the speed-up achieved in using the E - p transform to determine the maximum service magnitude for a given sample-state, in place of straightforward time-sequential simulation using the discrete time policy of [16]. We assume that the trapezoidal request is implemented as a piecewise constant signal of period 1 min, and for increased efficiency configure the routine so that it terminates as soon as the current value of the request cannot be met. A binary feasibility check is called once per iteration of the bisection procedure, but the number of iterations required depends on the sampled state. For a direct comparison, therefore, we find the average time taken across the bisection procedure, over the same 10^4 sampled states in each case. For the E - p case we include formation of the capacity curve within this, since it is not required during the straightforward simulation. The results of this comparison can be seen in Table II, where it can be seen that the E - p transform enables a 2.6x improvement over straightforward simulation. This experiment was performed in MATLAB on a workstation with an Intel Xeon processor and 64 GB of memory.

Method	Average run-time (ms)
Discrete time optimal policy	3.42
E - p transform	1.33

TABLE II: The total computation time using each method.

V. CONCLUSIONS AND FUTURE WORK

This paper has demonstrated use of the E - p transform to determine request feasibility, in both deterministic and stochastic settings. We have shown how this approach can be used to specify maximum service provision for a fleet of heterogeneous storage devices, and demonstrated that the speed-up achieved as compared to procedural simulation can

be significant. In deterministic settings the maximum service capability takes the form of a single curve, whereas in stochastic settings chance constraints can be utilised to ensure that the probability of successful delivery lies above a chosen limit. We have also investigated an approximation to this chance-constrained analysis. By forming a single capacity curve, this enables chance-constrained optimisation to be performed across any chosen service.

In future work the authors plan to develop an analytical basis for this approximation, investigate its accuracy when supplied with fewer samples and further explore its applicability. They also intend to embed these procedures into larger optimisation routines.

REFERENCES

- [1] "E-storage: Shifting from Cost to Value," World Energy Council, Tech. Rep., 2016.
- [2] D. M. Greenwood, K. Y. Lim, C. Patsios, P. F. Lyons, Y. S. Lim, and P. C. Taylor, "Frequency response services designed for energy storage," *Applied Energy*, vol. 203, pp. 115–127, 2017.
- [3] Y. Ota, H. Taniguchi, T. Nakajima, K. M. Liyanage, J. Baba, and A. Yokoyama, "Autonomous distributed V2G (vehicle-to-grid) satisfying scheduled charging," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 559–564, 2012.
- [4] E. Sortomme and M. A. El-Sharkawi, "Optimal scheduling of vehicle-to-grid energy and ancillary services," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 351–359, 2012.
- [5] S. Han, S. Han, and K. Sezaki, "Estimation of achievable power capacity from plug-in electric vehicles for V2G frequency regulation: Case studies for market participation," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 632–641, 2011.
- [6] M. C. Kisacikoglu, B. Ozpineci, and L. M. Tolbert, "Examination of a PHEV bidirectional charger system for V2G reactive power compensation," *Conference Proceedings - IEEE Applied Power Electronics Conference and Exposition - APEC*, pp. 458–465, 2010.
- [7] B. Zhou, W. Li, K. W. Chan, Y. Cao, Y. Kuang, X. Liu, and X. Wang, "Smart home energy management systems: Concept, configurations, and scheduling strategies," *Renewable and Sustainable Energy Reviews*, vol. 61, pp. 30–40, 2016.
- [8] J. Moshövel, K. P. Kairies, D. Magnor, M. Leuthold, M. Bost, S. Gähres, E. Szczechowicz, M. Cramer, and D. U. Sauer, "Analysis of the maximal possible grid relief from PV-peak-power impacts by using storage systems for increased self-consumption," *Applied Energy*, vol. 137, pp. 567–575, 2015.
- [9] X. Wu, X. Hu, X. Yin, and S. J. Moura, "Stochastic Optimal Energy Management of Smart Home With PEV Energy Storage," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 2065–2075, 2018.
- [10] M. P. Evans, S. H. Tindemans, and D. Angeli, "A Graphical Measure of Aggregate Flexibility for Energy-Constrained Distributed Resources," *IEEE Transactions on Smart Grid*, in press, pp. 1–12, 2019.
- [11] J. Cruise and S. Zachary, "Optimal scheduling of energy storage resources," *Preprint arXiv:1808.05901*, pp. 1–15, 2018.
- [12] H. Akhavan-Hejazi and H. Mohsenian-Rad, "Energy Storage Planning in Active Distribution Grids: A Chance-Constrained Optimization With Non-Parametric Probability Functions," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1972–1985, 2018.
- [13] T. Zhao, X. Pan, S. Yao, C. Ju, and L. Li, "Strategic Bidding of Hybrid AC/DC Microgrid Embedded Energy Hubs: A Two-stage Chance Constrained Stochastic Programming Approach," *IEEE Transactions on Sustainable Energy*, vol. PP, no. c, pp. 1–1, 2018.
- [14] Y. Xu, T. Zhao, S. Zhao, J. Zhang, and Y. Wang, "Multi-objective chance-constrained optimal day-ahead scheduling considering BESS degradation," *CSEE Journal of Power and Energy Systems*, vol. 4, no. 3, pp. 316–325, 2018.
- [15] 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.
- [16] M. P. Evans, S. H. Tindemans, and D. Angeli, "Minimising Unserved Energy Using Heterogeneous Storage Units," *IEEE Transactions on Power Systems*, in press, pp. 1–10, 2019.