Efficient Estimation of Upper Bounds on Arbitrage Values for Energy Storage Devices

Robert J. Barthorpe[1]

[1]*Department of Mechanical Engineering, The University of Sheffield, Sheffield, United Kingdom*
Correspondence author. Email: r.j.barthorpe@sheffield.ac.uk

ABSTRACT

Achieving flexibility through energy storage will be a key element in the transition to a low carbon energy grid. Decision making around investment in energy storage assets is a challenging task, and is typically reliant upon evaluating the expected performance of the planned asset on the basis of a set of economic metrics. The levelised cost of storage (LCOS) is one such metric, although significant limitations have been recognised, not least that temporal characteristics are not included. This means that the arbitrage value that may be achieved by the store - in the general case that charging and discharging prices vary over time - cannot be represented, with standing losses also not straightforwardly included. A potential solution is presented by the net present value (NPV) metric. The NPV describes the discounted present value of an asset considering both costs and revenue streams. As such, it offers a means of incorporating temporal behaviour including arbitrage value into the value calculation. This paper focuses on the development of a highly efficient method for estimating an upper bound on the arbitrage value that may be achieved by a given storage device, such that it may be included with the NPV metric.

Keywords: *Storage Metrics, Energy Arbitrage, Optimisation Tools*

1. INTRODUCTION

It is widely accepted that energy storage will be key to unlocking the benefits of future energy systems driven by renewable generation. It is expected that the storage options to be implemented will span a huge range of scales and technologies, and having access to efficient and reliable evaluation metrics will be of great importance in arriving at the most appropriate option for a given application. This paper focuses on one such metric, *net present value* (NPV), and in particular the development of an efficient method for evaluating the revenue term that is central to its calculation.

1.1. Storage Metrics

A number of metrics have been proposed for evaluating the economic benefit of energy storage devices. Among these, the levelised cost of storage (LCOS) has found particular traction, and is broadly defined as the total lifetime cost of the storage device divided by the total energy delivered over that time. The LCOS metric is derived from the widely used levelised cost of energy (LCOE) metric developed for comparison of energy generators. The commonly adopted definition in recent literature is,

$$LCOS = \frac{CAPEX + \sum_{n=1}^{N} \frac{CC_n + OPEX_n}{(1+r)^n} + \frac{EoL}{(1+r)^{N+1}}}{\sum_{n=1}^{N} \frac{Eout_n}{(1+r)^n}} \tag{1}$$

where $CAPEX$ is capital expenditure, $OPEX$ is operational expenditure, EoL is the end of life cost, CC is the charging cost, $Eout$ is the energy discharged from the system, n denotes the year of operation of the system up to the end of its lifetime in year N and r is the discount rate.

However, LCOS as commonly applied presents several limitations, the most significant of which is that temporal characteristics are not naturally included. It remains common practice for a constant energy price to be assumed, with the energy output by the system $Eout$ and the cost of charging the system CC being based on an assumed number of charge cycles per annum.

One alternative is the NPV metric, which describes the discounted present value of an asset considering both costs (one-off and recurring) and revenue streams. As such, it offers a means of incorporating temporal behaviour including arbitrage value into the value calculation via the revenue term. The NPV of a storage

P. Droege and L. Quint (eds.), *Proceedings of the International Renewable Energy Storage and Systems Conference (IRES 2023)*, Atlantis Highlights in Engineering 32,
https://doi.org/10.2991/978-94-6463-455-6_3

asset may be defined as,

$$NPV = \sum_{n=1}^{N} \frac{R_n - OPEX_n}{(1+r)^n} - \frac{EoL}{(1+r)^{N+1}} - CAPEX \quad (2)$$

where R_n denotes the revenue returned within year n.

The research question addressed in this paper is how to go about evaluating revenue for a given period based upon the characteristics of the storage device and the available energy price records, focusing on arbitrage as a key revenue stream. An approach is developed that makes use of graphical modelling methods to efficiently arrive at an estimate of the upper bound on the arbitrage value that may be achieved for a given set of storage parameters, facilitating evaluation of storage assets using the NPV metric. This enables a range of analyses, from comparison of competing storage technologies for a given application to decisions on whether increased CAPEX costs to achieve improved system performance can be justified. These analyses may be further extended by considering the *carbon* arbitrage benefits that may be achieved, using the same graphical modelling approach.

The layout of the paper is as follows. The proposed methodology is presented in §2, with demonstration for a representative set of storage devices in §3. Discussion and brief conclusions are presented in §4.

2. METHODOLOGY

The approach adopted in this paper employs graphical modelling techniques, and proceeds by:

1) Ingesting and preprocessing historical (or forecast) time series data on available energy prices and associated carbon intensities.
2) Identifying peaks and troughs in the time series that represent candidate 'sell' and 'buy' points, respectively.
3) Setting up a directed acyclic graph (DAG) to represent the buy/sell/hold actions ('edges') connecting the identified buy/sell points ('vertices').
4) Searching for the path through this graph that optimises carbon/cost benefits, efficiently returning an estimated upper bound on cost and/or carbon savings.

The approach is motivated by methods employed with financial analysis which seek to identify the maximum profit that may be realised over a period of time through optimisation of trading decisions. This application is illustrated in §2.1, followed by extension to the energy arbitrage case in §2.2.

2.1. *Stock Trading*

The approach proceeds by constructing a weighted, bipartite DAG denoted $G(U, V, E, R)$ to represent the trading problem. A bipartite graph is one whereby the vertices may be partitioned into two disjoint, independent, nonempty subsets U and V, with all edges in E connecting a vertex in U to a vertex in V. In a weighted graph, the set R contains values assigned to each edge in E.

In the stock trading problem, we may consider U to contain the set of all candidate "buy" points and V to contain the candidate "sell" points. The set E consists of two types of directed edge: *trading edges* which start at a vertex in U and end at a vertex in V, and *forward edges* which start at a vertex in V and end at a vertex in U. It should be noted that this approach essentially represents an all-in, "buy low, sell high" approach to maximising profit. It is also an *ex-post* analysis, so assumes that data is perfectly known for the period considered.

Graph construction begins with identification of the subsets U and V (e.g. through application of a peak picking method to identify local maxima and minima within the price record) and by setting $E = R = \emptyset$.

A new trading edge (i, j) between a vertex $i \in U$ and $j \in V$ is created and assigned a weight $R_{i,j}$ if and only if:

- $t_j \geq t_i$ to ensure temporal consistency, where t_i and t_j are the timestamps associated with i and j respectively.
- The trade is profitable, i.e.,

  $$R_{i,j} = p_j(1-c) - p_i(1+c) \geq 0 \quad (3)$$

  where c is the rate of transaction costs (here treated as equal for buy and sell actions) and p_i and p_j are the prices associated with i and j respectively.

A new forward edge between vertices $j \in V$ and $i \in U$ is created and assigned a weight $R_{j,i} = 0$ if and only if:

- $t_i \geq t_j$ to ensure temporal consistency.

Once the graph is constructed a range of analyses may be performed. Most pertinently, the length of the longest admissible path through the graph represents the optimal solution to the problem i.e. the maximum profit that may be made during the period considered for the stock trading case. Note that in practice it is commonplace to replace searching for the longest path through $G(U, V, E, R)$ with a search for the shortest path through $G(U, V, E, -R)$ (commonly referred to as $-G$); these approaches are equivalent for DAGs and allow the use of efficient shortest path search algorithms. The

Table 1. Stock trading data: buy points

t	1	3	5	7	9	11
p_t	100	90	126	140	110	168

Table 2. Stock trading data: sell points

t	2	4	6	8	10	12
p_t	120	160	150	160	170	185

shortest path algorithm is of complexity $\mathcal{O}(|E|)$, solving in linear time. Retaining only those trading edges that are profitable promotes sparsity within the graph and thus increases the efficiency of solution. A second analysis of less relevance to the stock trading case but of value is to the energy storage scenario is to consider the returns that may be realised as the number of trades within the period is constrained. This extension is omitted from the current study but will be returned to in future work.

In order to illustrate the method, the example presented in [1] is recreated. The example is based on the artificial stock price data shown in Fig. 1, with the data is sliced into candidate buy and sell points as shown in Tables 1 and 2. The resulting graph for a transaction cost of $c = 10\%$ is shown in Fig. 2, with the optimal path through the graph shown in red. Note that dummy start and end points (labelled 13 and 14) are included to aid illustration; the start node has all $i \in U$ as successor vertices, with the end node being a successor to all $j \in V$. In both cases the edges from/to these nodes are zero-weighted. The optimal profit returned by the sequence of trades illustrated in Fig. 2 is 91.4, the sum of the weights associated with the trading edges on the optimal path.

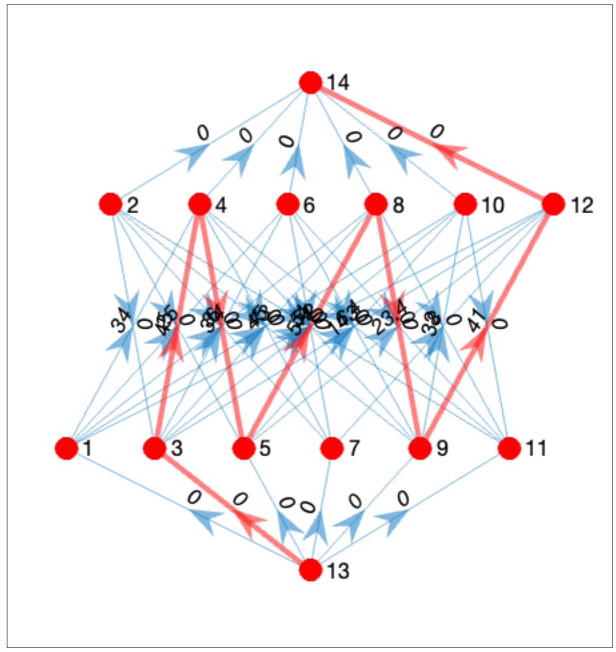

Figure 2 Directed acyclic graph (blue) and optimal path (red) for the stock trading problem.

2.2. Energy Arbitrage

Adapting the approach presented in §2.1 to the case of energy arbitrage requires modification of the trading edge weight calculation, and consideration of the price record used. These steps allow the incorporation of key storage parameters (e.g. overall capacity, charge and discharge rates and efficiencies, standing losses) within the optimal revenue calculation. Graph construction proceeds along similar lines, with subsets of candidate charge and discharge points identified from the appropriate price record, and with $E = R = \emptyset$.

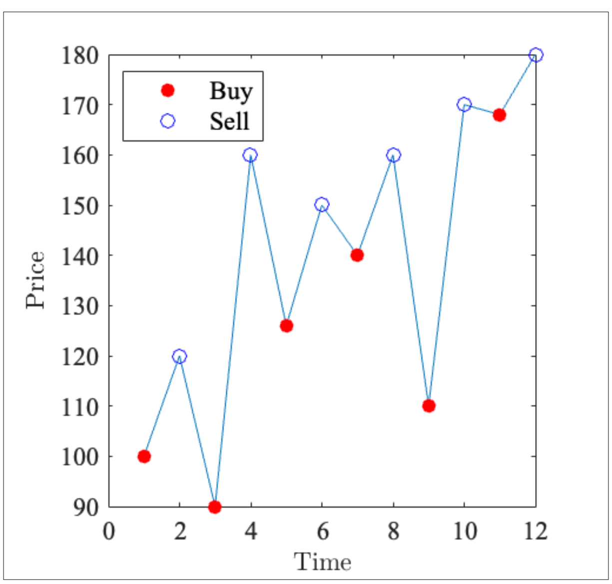

Figure 1 Exemplar stock trading data.

The only difference in graph construction arises from definition of the trading edge weights. A new trading edge (i, j) between a vertex $i \in U$ and $j \in V$ is created and assigned a weight $R_{i,j}$ if and only if:

- $t_j \geq t_i$ to ensure temporal consistency, where t_i and t_j are the timestamps associated with i and j respectively.
- The trade is profitable, i.e.,

$$R_{i,j} = p_j \eta_{dch} f(t_i, t_j) - p_i \eta_{ch} \geq 0 \qquad (4)$$

  where and η_{ch} and η_{dch} are respectively the charge and discharge efficiencies and $f(t_i, t_j)$ is a standing loss function.

As before, a new forward edge between vertices $j \in V$ and $i \in U$ is created and assigned a weight $R_{j,i} = 0$ if and only if:

- $t_i \geq t_j$ to ensure temporal consistency.

The second element in tailoring the approach to the energy arbitrage case is through consideration of the price record from which the graph is constructed. In general, the cost of fully charging the store from empty at a given time t_i will be given by,

$$p_i = \int_{t=t_i}^{t_i+T^{ch}} p(t) P^{ch}(t) dt \qquad (5)$$

where $P_{ch}(t)$ describes the charging power profile, $p(t)$ represents the energy price record and T^{ch} is the time taken to complete a full charge. If the charge rate $P^{ch}(t)$ is assumed to be constant for the duration of the charging process, the charging time is given by,

$$T^{ch} = \frac{E}{P^{ch}} \tag{6}$$

with the time-dependency of P^{ch} omitted. Similarly, the energy cost associated with complete discharge of the store starting at time t_j may be expressed as,

$$p_j = \int_{t=t_j}^{t_j+T^{dch}} p(t)P^{dch}(t)dt \tag{7}$$

Note that there is no requirement for $P^{ch}(t)$ and $P^{dch}(t)$ to be either constant or to be linear functions of time, as would be required with a linear programming approach to the arbitrage estimation problem. This permits accurate representation of arbitrarily complex charge and discharge profiles without impacting computational efficiency. It also allows for the inclusion of other impacts on the charging price, as is the case in the demonstration example presented in §3. In the example, it is assumed that the (thermal) storage device is charged via a heat pump, conferring a coefficient of performance (CoP) advantage. The charge and discharge costs in Eqns. 5 and 7 may be be straightforwardly adjusted to consider the combined heat pump and thermal store system through addition of the time varying CoP, resulting in,

$$p_i = \int_{t=t_i}^{t_i+T^{ch}} \frac{p(t)P^{ch}(t)}{CoP(t)}dt \tag{8}$$

and,

$$p_i = \int_{t=t_i}^{t_i+T^{dch}} \frac{p(t)P^{dch}(t)}{CoP(t)}dt \tag{9}$$

Note that positive CoP values will have the effect of reducing the “spread” between charge and discharge costs, and thus the arbitrage value returned.

Finally, the approach shown in Eqns. 8 and 9 for energy prices may equivalently be applied to carbon intensity data through simple substitution of $p(t)$. This is demonstrated in § 3.3.

3. DEMONSTRATION

The approach introduced in §2.2 is demonstrated for a full year of weather, carbon and energy data for a location in Nottingham, UK. The period considered is 1st January to 31st December 2021. The parameters of the storage device used for illustration are presented in Table 3. These are chosen to be representative of a thermal storage device used to service intraday domestic hot water (DHW) demand within a residential property. The system configuration considered is that the thermal storage system is used in conjunction with an air source heat pump, modelled using the method proposed in [2] and assuming a charging temperature of $T_{ch} = 60°$C.

3.1. Price Arbitrage

The effective price record for 2021 is shown in Fig. 3. The data used is taken from the Octopus Agile Incoming tariff for the Nottingham region. This tariff is broadly based upon the UK wholesale energy price, although weighted such that higher prices are charged during peak periods (5pm-10pm daily) with lower prices offered through the rest of the day. It should be noted that 2021 represented a somewhat atypical period for energy prices which rose markedly towards the end of the year. This resulted in a reduction in the effective spread between candidate buy and sell points, and thus a reduction in the density of profitable trading edges in this portion of the graph. This is apparent when considering the sequence of optimal trades presented in Fig. 4, with fewer trades made during December. Indeed, the optimal number of trades for the given year was 320, indicating that there were days within the year when it would not be profitable to operate the store. For the period shown, the optimal strategy would save £18.67 per kWh of storage, or £186.72 pa in total. By evaluating the carbon arbitrage outcome that would be returned by the price optimal approach, it was found that a comparatively small quantity (20.05 kgCO_2) would be saved. The following section considers the savings that may be achieved by adopting a carbon optimal approach.

3.2. Carbon Arbitrage

Information on operational carbon savings is of value for numerous purposes, from whole life carbon assessment (WLCA) to the evaluation of trade offs between cost and carbon optimal approaches. In this section, the analysis applied to arrive at a price-optimal strategy is repeated in order to arrive at a carbon-optimal alternative through substitution of the price record with appropriate grid carbon intensity (GCI) data. The effective carbon intensity for the period considered is presented in Fig. 6. The data is provided by National Grid ESO’s Carbon Intensity API, which provides both forecast and historical national and regional level UK GCI data at a half hourly time resolution.

For the period shown, the carbon optimal strategy involved 308 charge cycles and would save 14.0 kgCO_2/kWh of storage, or 140.42 kgCO_2pa in total. It would also deliver £25.38 of total price arbitrage benefit. It is informative to compare these values to the £186.72 pa and 20.05 kgCO_2pa savings returned by the price optimal strategy in §3.1. These represent the (approximate) extrema of the operational cost/carbon trade off that may

Table 3. Storage parameters used for demonstration

Cap [kWh]	P_{ch} [kW]	P_{dch} [kW]	η_{ch} [%]	η_{dch} [%]	ϵ_{loss} [%/Δt]
10	8	8	93	93	99.58

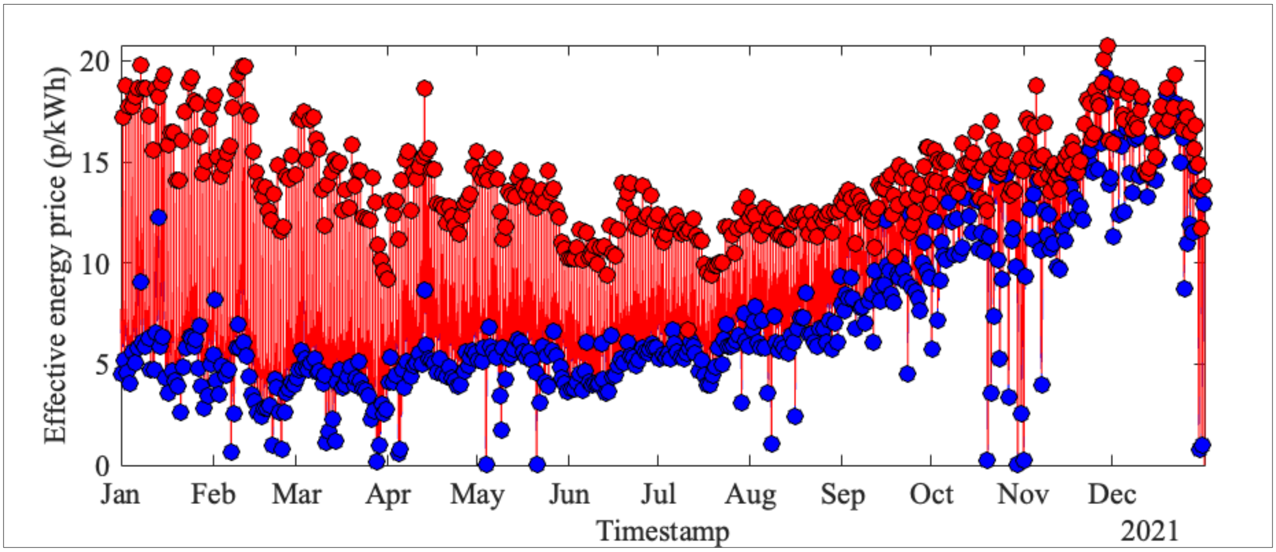

Figure 3 Effective energy price plot showing candidate buy (blue) and sell (red) points.

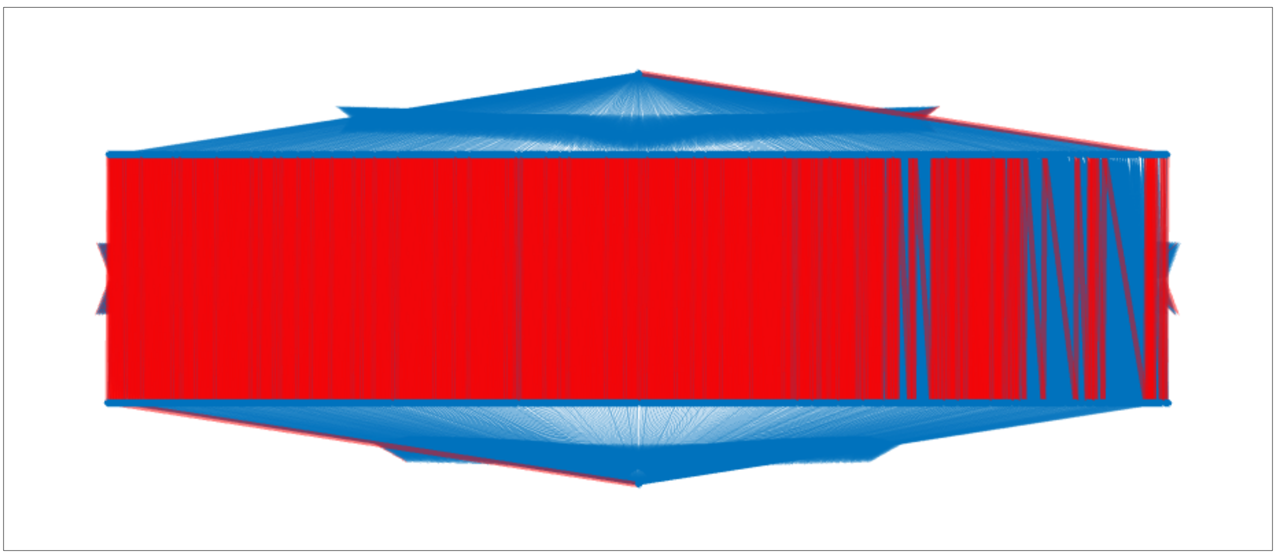

Figure 4 Directed acyclic graph (blue) and optimal path (red) for the price arbitrage problem.

be achieved for the given data and storage parameters. Extension to more formal consideration of this trade off is briefly discussed in §4.

3.3. *NPV and LCOS outcomes*

To enable an initial demonstration of NPV outcomes using the DAG-based arbitrage estimation approach, and to indicate how these outcomes could be used to inform development decisions, a simplified sensitivity study is presented. The adopted economic parameters for the storage device are presented in Table 4. Table 5 presents the arbitrage and subsequent NPV outcomes for variations on the exemplar storage device used for illustration in §3.1 and 3.3. With the NPV value to hand, it is straightforward to evaluate whether an additional CAPEX investment is justified

Table 4. Exemplar storage device parameters

Store	CAPEX [£]	OPEX [£/pa]	EOL [£]	r [%]	Life
1	3000	50	100	3	25

4. DISCUSSION

Extensive discussion is not warranted given the comparatively early stage of the work. However, a number of advantages of the approach are apparent from the initial exposition presented in this paper. Central to this is the efficiency of the method: once the graph is constructed, an optimal solution may be achieved in linear time ($\mathcal{O}(|E|)$). By way of illustration, for the

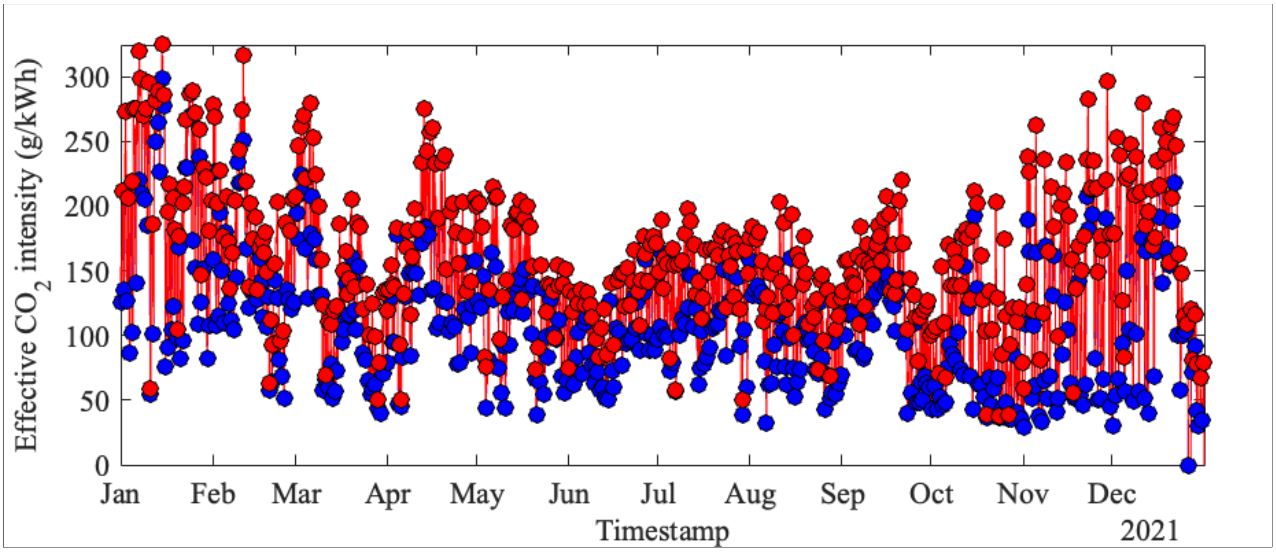

Figure 5 Effective carbon price plot showing candidate buy (blue) and sell (red) points.

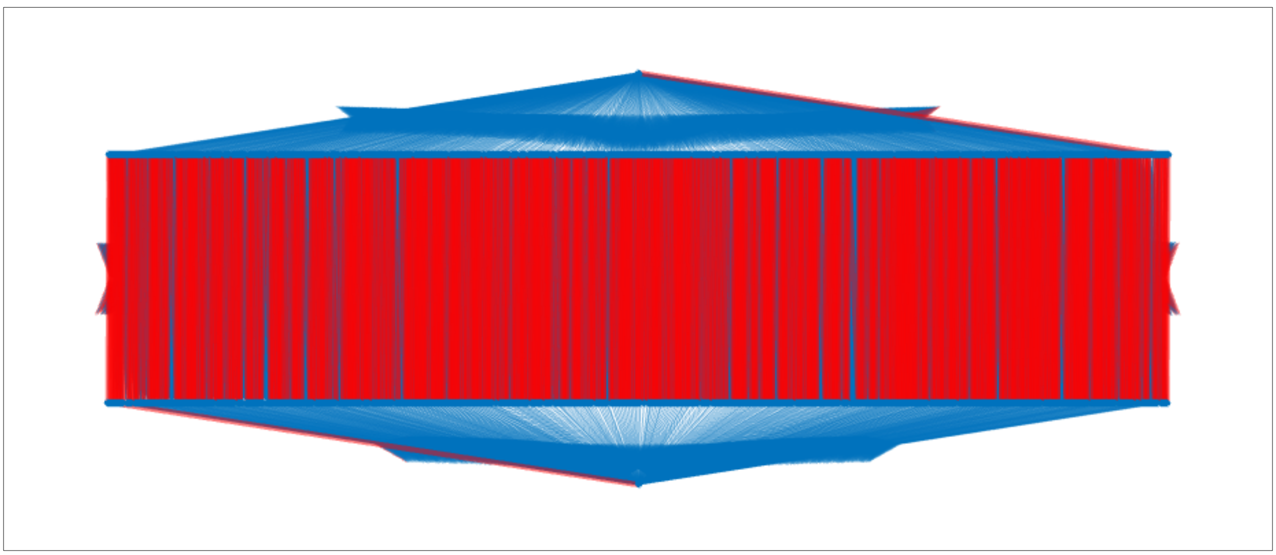

Figure 6 Effective carbon price plot showing candidate buy (blue) and sell (red) points.

Table 5. NPV results

Case	P_{ch} **[kW]**	P_{dch} **[kW]**	η_{ch} **[%]**	η_{dch} **[%]**	P_{loss} **[W]**	ϵ_{loss} **[%/Δt]**	R **[£ pa]**	**NPV [£]**
Ideal	11	11	100	100	0	100	271	2858
Midpoint	8	8	93	93	77	99.58	187	758
Loss$_{low}$	8	8	93	93	27	99.86	207	1258
Loss$_{high}$	8	8	93	93	127	99.25	168	283
Power$_{high}$	11	11	93	93	77	99.58	191	858
Power$_{low}$	5	5	93	93	77	99.58	177	508
Eff$_{high}$	8	8	98	98	77	99.58	219	1558
Eff$_{low}$	8	8	88	88	77	99.58	156	-17

examples shown in §3.1 and 3.3, graph construction took approximately 0.1s and shortest path optimisation approximately 0.02s using a modern laptop (Apple M1 Pro). A second major advantage of the approach is that it presents the possibility of including accurate nonlinear charge and discharge behaviour without any increase in computation cost. This sets the approach apart from linear programming approaches, which require linear (or linearised) behaviour.

The principal limitation of the approach as presented in the current study is the fact that it is restricted to representing an "all in" arbitrage strategy. Crucially, this assumes that demand must exist at all discharge points, which may not be the case in all applications (e.g. use of thermal energy for domestic space heating). It also omits consideration of part-charging; while it can be demonstrated that part-charging will play no role in the optimal solution for situations where there are no standing losses (see [1]), this is not the case for the majority of energy storage technologies. Implementing peak picking to identify candidate buy and sell points and thus limit the size of the graph introduces a further contribution potential contribution to sub-optimality. Finally, as stated at the outset this is an *ex-post* analysis; data must be available, whether from historical records or through forecasting. Achieving performance close to the identified optima in practice is the remit of effective control.

The next stage for the presented work is to conduct a full comparison to alternative approaches, principally linear programming (LP). This will enable both comparison of execution times and quantification of the deviation from the true optimum that arises due to the "all in" investment strategy. It will also be informative to more thoroughly explore use of the metric to evaluate the trade off between cost and carbon optimal approaches; the efficiency of graph construction and solution presents the possibility of exploring such trade offs in a detailed fashion without excessive computational expense.

ACKNOWLEDGEMENTST

This research was funded by the UK Department for Energy Security and Net Zero (DESNZ) through the Advanced Distributed Storage for grid Benefit (ADSorB) project.

REFERENCES

[1] Brandouy O, Mathieu P, Veryzhenko I, "Algorithmic determination of the maximum possible earnings for investment strategies" Decision Support Systems, 54:2, (2013) 816-82.

[2] Staffell I, Brett D, Brandon N, Hawkes A, "A review of domestic heat pumps" Energy & Environmental Science, 5, (2012) 9291-9306

[3] R. Floyd, "Algorithm 97, shortest path algorithms", Operations Research, 17 (1969) 395–412