

Coalition Formation with Limited Information Sharing for Local Energy Management*

Luke Rickard¹, Paola Falugi¹ and Eric C. Kerrigan²

Abstract—Distributed energy systems with prosumers require new methods for coordinating energy exchange among agents. Coalitional control provides a framework in which agents form groups to cooperatively reduce costs; however, existing bottom-up coalition-formation methods typically require full information sharing, raising privacy concerns and imposing significant computational overhead.

In this work, we propose a limited information coalition-formation algorithm that requires only limited aggregate information exchange among agents. By constructing an upper bound on the value of candidate coalitions, we eliminate the need to solve optimisation problems for each potential merge, significantly reducing computational complexity while limiting information exchange. We prove that the proposed method guarantees cost no greater than that of decentralised operation.

Coalition strategies are optimised using a distributed approach based on the Alternating Direction Method of Multipliers (ADMM), further limiting information sharing within coalitions. We embed the framework within a model predictive control scheme and evaluate it on real-world data, demonstrating improved economic performance over decentralised control with substantially lower computational cost than full-information approaches.

I. INTRODUCTION

Energy production around the world is currently undergoing a significant paradigm shift from fossil fuels to renewable energy sources. One part of this transition is the change from relying on a few large, high-capacity power plants to depending on many smaller, widely distributed generation units [1]. In particular, we are interested in the rise of “prosumers”, participants in the energy market who both consume and produce electricity [2], [3]. This motivates exploring novel control techniques better suited to the new energy generation landscape.

One promising area of research that leverages this shift in generation is *coalitional control* [4], [5]. Coalitional control is an approach that considers grouping entities involved in a large problem into separate coalitions that can be solved separately. We refer to this initial grouping of agents, before the subproblems are solved, as coalition formation. For the grid scenario, coalitional control can involve buildings agreeing to trade energy with one another, aiming to pay less than they would by interacting only with the grid.

Forming coalitions allows agents to reduce their individual costs without incurring large computational or communication overheads, and without sharing full local information with agents outside the coalition (once coalitions are

formed). Revealing private information (such as usage data) in the case of the smart grid, can lead to vulnerabilities, allowing attackers to infer various household traits [6], [7].

However, while recent work that has explored bottom-up coalition formation [5] addresses computational complexity, it still requires full information sharing during coalition formation, negating these privacy benefits. In particular, the technique in [5] required solving the complete optimisation problem for every pair of agents, requiring each agent to broadcast all optimisation variables and constraints to all other agents. Not only does this introduce privacy concerns, but it also represents a computational bottleneck. We seek a “limited information sharing” scheme that avoids sharing full local demand, generation, storage, and constraint data. Instead, agents only share tentative grid trade plans during coalition formation, and coalition trade plans afterwards. This reduces information shared and computational costs.

We make use of Model Predictive Control (MPC) techniques to optimise energy trades and battery charging profiles, with forecasted energy usage and production for a day ahead. MPC has been studied in this context in [8], but it is limited to a single building, whereas we consider a group of buildings that may form into coalitions. Our work extends upon current extensions of MPC to coalitions [9], [10] by reducing the amount of information shared (and computational requirements) for each possible coalition formation step, and dynamically updates coalitions at every time step to account for changing conditions.

Coalitional control aligns with techniques in distributed control, such as the Alternating Direction Method of Multipliers [11], explored for energy grids in [12], [13]. While distributed control provides a powerful framework for optimising systems composed of multiple agents, it typically assumes a fixed interaction structure and does not address how agents should be grouped. In many applications, forming such groups is desirable to limit information sharing, reduce computational burden, or manage communication costs. Coalitional control addresses this by determining groups of agents (coalitions) that can then optimise their decisions jointly, e.g., using distributed optimisation techniques.

In this work, we propose a coalition formation framework that reduces both computational complexity and information sharing. The approach enables agents to evaluate potential coalition merges using only limited, aggregate information, without requiring full knowledge of local models or solving a full optimisation problem for each candidate merge. The main contributions of this paper are as follows:

- We propose a coalition formation algorithm that avoids

*This work was supported by EPSRC grant EP/Z536106/1

¹University of East London

²Imperial College London

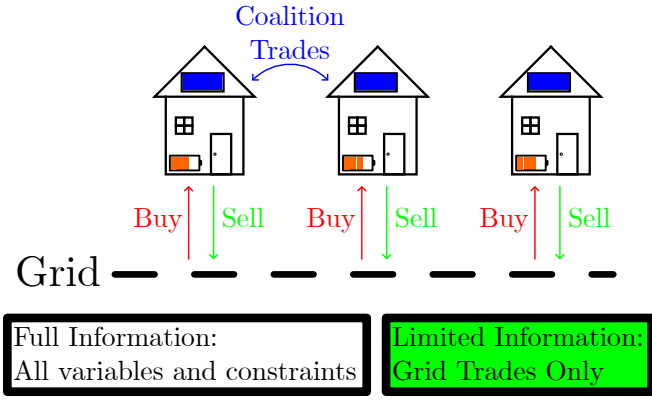


Fig. 1: Overview of the energy system with coalitions. Buildings may trade with the grid, or with other buildings within their coalition (but not those outside the coalition). Our limited-information approach shares only planned grid trades, whereas full-information methods require sharing all optimisation variables and constraints.

solving optimisation problems for candidate merges by using an efficiently computable upper bound on coalition value.

- The proposed method requires only aggregate grid trade information, thereby reducing information sharing and preserving agent privacy.
- We prove the resulting coalition structure guarantees a total cost no greater than that of decentralised operation.
- We integrate the proposed coalition formation scheme with distributed optimisation via ADMM and a model predictive control framework.
- Numerical results on real-world data demonstrate improved performance over decentralised control and lower computational cost than full-information methods.

While we focus on energy markets in this paper, our technique is generally applicable to all trading games, and we discuss how it may be extended to allow a trade-off between optimality and information sharing.

A. Notation

Where we have a sequence $\{G(t)\}_{t=1}^T$, we use G to refer to the entire sequence collectively (i.e. $G = \{G(t)\}_{t=1}^T$).

II. BACKGROUND

A. Building Model

We consider a collection of N buildings $\mathcal{M}_1, \dots, \mathcal{M}_N$, each with some fixed demand and generation profiles $d_i(t), g_i(t) \in \mathbb{R}_{\geq 0}$, corresponding to the energy consumed, and produced, at each time step t across the time horizon $\mathcal{T} = \{1, \dots, T\}$. Each building has a battery with initial state $\text{SoC}_i(1) = \text{SoC}_i^{\text{init}}$, whose state $\text{SoC}_i(t) \in [0, \overline{\text{SoC}}_i]$ evolves according to the following dynamics

$$\text{SoC}_i(t+1) = \text{SoC}_i(t) + \rho_i^c u_i^c(t) - \frac{1}{\rho_i^d} u_i^d(t), \quad (1)$$

where $\rho_i^c, \rho_i^d \in (0, 1)$ are charge and discharge efficiencies respectively, and $u_i = (u_i^c, u_i^d) \in [0, \bar{u}_i]^2$ are charge and discharge signals. The building may trade energy with the grid by purchasing or selling energy to the grid at each time step $G_i(t) = (G_i^{\text{buy}}(t), G_i^{\text{sell}}(t)) \in \mathbb{R}_{\geq 0}^2$, and we refer to the total trade with the grid as $G_i^{\text{tot}}(t) = G_i^{\text{buy}}(t) - G_i^{\text{sell}}(t)$. At each time step, the building i must satisfy a power balance

$$d_i(t) + u_i^c(t) = g_i(t) + u_i^d(t) + G_i^{\text{tot}}(t). \quad (2)$$

We then define a linear cost function for each building as

$$L_i(G_i) = \sum_{t=1}^T \left(P^{\text{buy}}(t) G_i^{\text{buy}}(t) - P^{\text{sell}}(t) G_i^{\text{sell}}(t) \right), \quad (3)$$

where $P^{\text{buy}}(t), P^{\text{sell}}(t)$ are the prices offered by the grid for buying and selling electricity. We assume $P^{\text{buy}}(t) > P^{\text{sell}}(t) > 0$ for all t . Under this assumption, and with $\rho_c, \rho_d \in (0, 1)$, simultaneous charging and discharging is suboptimal [14]. Finally, we denote the feasible set of variables $\{G_i, u_i\}$ by

$$\begin{aligned} \mathfrak{X}_i &= \{G_i \in \mathbb{R}_{\geq 0}^{2 \times T}, u_i \in [0, \bar{u}_i]^{2 \times T} : \\ & d_i(t) + u_i^c(t) = g_i(t) + u_i^d(t) + G_i^{\text{tot}}(t) \quad \forall t \in \mathcal{T}, \\ & \text{SoC}_i(t+1) = \text{SoC}_i(t) + \rho_i^c u_i^c(t) - \frac{1}{\rho_i^d} u_i^d(t) \quad \forall t \in \mathcal{T}, \\ & \text{SoC}_i(t) \in [0, \overline{\text{SoC}}_i] \quad \forall t \in \mathcal{T} \cap \{T+1\}, \\ & \text{SoC}_i(1) = \text{SoC}_i^{\text{init}} \}. \end{aligned} \quad (4)$$

B. Open-loop Optimisation

We first seek to minimise the cumulative cost for all buildings $\sum_{i=1}^N L_i(G_i)$ in an open-loop fashion, we later consider how to extend to model predictive control in Section IV. We therefore define an optimisation problem

$$\min_{\{G_i, u_i\} \in \mathfrak{X}_i, i=1, \dots, N} \sum_{i=1}^N L_i(G_i) \quad (5)$$

1) *Decentralised*: We can naturally separate this problem into N linear programs

$$\sum_{i=1}^N \min_{\{G_i, u_i\} \in \mathfrak{X}_i} L_i(G_i) \quad (6)$$

which can then be solved directly using off-the-shelf optimisation packages. However, we can improve upon this optimum by allowing for buildings to form into groups, which we call coalitions, in which energy is freely traded to benefit all members of the coalition.

C. Coalitions

We consider the problem of forming a coalitional structure, defined in the sequel, to reduce the shared cost of all agents.

Definition 1 (Coalition Structure): A coalition structure $\mathcal{C} = \{\mathcal{C}_1, \dots, \mathcal{C}_M\}$, denotes a disjoint partitioning of the complete set of building indices

$$\bigcup_{j=1}^M \mathcal{C}_j = \{i\}_{i=1}^N, \quad \mathcal{C}_i \cap \mathcal{C}_j = \emptyset, \quad i, j = 1, \dots, M, i \neq j. \quad (7)$$

Note that a coalition \mathcal{C}_j may be a collection of buildings, or a single building.

We now introduce the additional optimisation variables $C_i^{\text{buy}}(t), C_i^{\text{sell}}(t) \in \mathbb{R}_{\geq 0}$ for each building, which allow that building to buy from, and sell to, its coalition. Again, we use $C_i^{\text{tot}}(t) = C_i^{\text{buy}}(t) - C_i^{\text{sell}}(t)$ to refer to the total traded with the coalition, and $C_i(t) = (C_i^{\text{buy}}(t), C_i^{\text{sell}}(t))$ to refer to the pair. The power balance in (2) is then amended to

$$d_i(t) + u_i^e(t) = g_i(t) + u_i^d(t) + G_i^{\text{tot}}(t) + C_i^{\text{tot}}(t), \quad (8)$$

and the constraint set \mathfrak{X}_i is amended similarly (with other terms unaffected).

For a given coalition $\mathcal{C}_j \in \mathcal{C}$, we can then define the optimisation problem as

$$\begin{aligned} V(\mathcal{C}_j) = & \min_{\{G_i, u_i, C_i\} \in \mathfrak{X}_i, i \in \mathcal{C}_j} \sum_{i \in \mathcal{C}_j} L_i(G_i) \\ & \text{subject to } \sum_{i \in \mathcal{C}_j} C_i^{\text{tot}}(t) = 0, \forall t \in \mathcal{T}, \end{aligned} \quad (9)$$

with a new constraint to maintain a power balance for the coalitional trades. For the entire coalition structure, we then seek to minimise the cumulative value $\sum_{j=0}^M V(\mathcal{C}_j)$.

We now have $M \leq N$ linear programs, with the new constraint preventing further separation. We assume that trades within a coalition are settled at 0 cost; this allows us to focus on the limited information coalition formation mechanism and derive a simple achievable upper bound. Extensions to unequal internal trading costs (e.g. to account for transmission costs) are left for future work.

III. COALITION FORMATION

In order to optimise the cumulative cost of all agents, it is also necessary to optimise over coalition structures. We now consider a number of techniques for performing this optimisation, with varying degrees of information sharing. We assume that there is an upper bound on the size of a coalition \mathcal{C}_{max} ; this constraint may be introduced to limit the number of agents with whom operational information is shared, or to lower communication/computational overhead.

A. Optimal (Size-Limited) Coalition

First, we consider the problem of directly obtaining the true optimal coalition, hence we seek to solve

$$\begin{aligned} & \min_{\mathcal{C}} \sum_{\mathcal{C}_j \in \mathcal{C}} V(\mathcal{C}_j) \\ & \text{subject to } \bigcup_{j=1}^M \mathcal{C}_j = \{i\}_{i=1}^N \\ & \quad \mathcal{C}_j \cap \mathcal{C}_k = \emptyset, j, k = 1, \dots, M, i \neq j \\ & \quad |\mathcal{C}_j| \leq \mathcal{C}_{\text{max}}, j = 1, \dots, M, \end{aligned} \quad (10)$$

which can be formulated as a mixed-integer linear program (due to the optimisation over coalition structures \mathcal{C}), since we now optimise also over each building's membership of coalitions. The brute-force approach to solving this problem is to consider all possible sets of coalitions, find the optimal cost for each, and select the best amongst these. However, this approach is combinatorial in the number of agents (requiring iteration over all possible partitions of a set).

Algorithm 1 Bottom-Up Coalition Formation

```

1:  $\mathcal{A}_{\text{new}} \leftarrow \{\mathcal{M}_1, \dots, \mathcal{M}_N\}$ 
2: repeat
3:    $\mathcal{A} \leftarrow \mathcal{A}_{\text{new}}$ 
4:   for all  $a \in \mathcal{A}$  do
5:      $\{G_i^*, u_i^*, C_i^*\}_{i \in a} \leftarrow \text{argmin}_{\{G_i, u_i, C_i\}_{i \in a}} \sum_{i \in a} L_i(G_i) \triangleright$ 
     Subject to  $\{G_i, u_i, C_i\} \in \mathfrak{X}_i$  and  $\sum_{i \in a} C_i^{\text{tot}}(t) = 0, t \in \mathcal{T}$ 
6:     for all  $\tilde{\mathcal{C}} \in \{(\mathcal{A}_j \cup \mathcal{A}_k) | j = 1, \dots, |\mathcal{A}| - 1, k = j +$ 
      $1, \dots, |\mathcal{A}|\}$  do
7:       if  $|\tilde{\mathcal{C}}| \leq \text{max coalition size}$  then
8:          $\text{val}(\tilde{\mathcal{C}}) \leftarrow V(\mathcal{A}_j) + V(\mathcal{A}_k) - V(\tilde{\mathcal{C}})$ 
9:          $\text{sortedCosts} \leftarrow \text{sort}(\text{val})$ 
10:         $\mathcal{A}_{\text{new}} \leftarrow \emptyset$ 
11:        for  $(\tilde{\mathcal{C}}, \text{val})$  in  $\text{sortedCosts}$  do
12:          if  $\text{val}(\tilde{\mathcal{C}}) > 0$  and  $\tilde{\mathcal{C}} \cap (\bigcup \mathcal{A}_{\text{new}}) = \emptyset$  then  $\triangleright$ 
          Coalition beneficial and neither agent accounted for
13:             $\mathcal{A}_{\text{new}} \leftarrow \mathcal{A}_{\text{new}} \cup \{\tilde{\mathcal{C}}\}$ 
14:             $\mathcal{A}_{\text{new}} \leftarrow \mathcal{A}_{\text{new}} \cup (\mathcal{A} \setminus (\bigcup \mathcal{A}_{\text{new}})) \triangleright$  Add all unpaired
          agents
15:        until  $\mathcal{A}_{\text{new}} = \mathcal{A}$ 

```

B. Bottom-Up Coalition Formation [5]

One less computationally expensive approach to approximately solve the coalition formation is proposed in [5]. The algorithm is summarised in Algorithm 1.

This algorithm cannot guarantee convergence to the optimal coalition structure, but is significantly less computationally intensive than solving (10) directly. The number of possible coalition structures of N agents grows according to the Bell numbers and is therefore combinatorial in N . Consequently, obtaining the optimal coalition structure by exhaustive enumeration quickly becomes intractable. Algorithm 1 avoids this combinatorial search by iteratively evaluating candidate pairwise merges among current coalitions. If m_k denotes the number of coalitions at iteration k , then the number of candidate merges at that iteration is $\binom{m_k}{2} = O(m_k^2)$.

C. Limited Information Coalition Formation

The approach in Algorithm 1 suffers from a crucial drawback, namely the need to solve the complete optimisation problem for every pair of agents, thereby sharing full local information with everyone. We wish to avoid sharing data on individual agents' consumption and generation behaviour. Not only does this raise privacy concerns, but also introduces a computational burden for every potential coalition. To remove this requirement, we recognise that *only an upper bound* on the solution to (9) is needed. We denote by $\bar{V}(\mathcal{C}_j)$, the upper bound on the joint cost received by coalition \mathcal{C}_j

$$\bar{V}(\mathcal{C}_j) \geq V(\mathcal{C}_j). \quad (11)$$

Then, we replace Line 8, with

$$\text{val}(\tilde{\mathcal{C}}) \leftarrow V(\mathcal{A}_j) + V(\mathcal{A}_k) - \bar{V}(\tilde{\mathcal{C}}), \quad (12)$$

and proceed with the rest of the algorithm unchanged, forming coalitions only if this new bound is positive.

The tighter this bound (by sharing more information), the better the final coalition structure. Hence, this trade-off gives agents a choice between information sharing and optimality. Once coalitions are formed, we assume the following (in Section III-D, we discuss how this can be achieved without sharing private information):

Assumption 1 (Value Knowledge of Formed Coalitions): Once a coalition \mathcal{C}_j is formed, its true value $V(\mathcal{C}_j)$ can be evaluated exactly.

We then have:

Theorem 1 (Limited Information Algorithm Properties): Consider Algorithm 1, with an upper bound $\bar{V}(\mathcal{C}_j)$ (satisfying (11)) on the joint cost in place of solving the optimisation problem in Line 8. Then the algorithm will satisfy the following properties

- 1) The resulting coalition structure satisfies Definition 1;
- 2) The algorithm terminates in a finite number of iterations;
- 3) The resulting coalition structure has a value at most equal to the decentralised optimum.

Proof: We consider each item in turn:

- 1) Since a coalition is accepted only if it does not overlap with a coalition already in the coalition structure (Line 12), and any coalitions not in a pair with another coalition are added at the end (Line 14), then the resulting coalition structure is valid.
- 2) We initialise the algorithm with $M = N$ coalitions (each agent in a singleton coalition). At each iteration, we either form new coalitions (strictly decreasing M) or terminate. M is bounded below by $\lceil \frac{N}{\mathcal{C}_{\max}} \rceil$ and decreases if coalitions form, so the algorithm terminates in at most $N - \lceil \frac{N}{\mathcal{C}_{\max}} \rceil$ iterations.
- 3) Our proposed algorithm forms coalitions between pairs of coalitions a, b only if the upper bound satisfies

$$\bar{V}(a \cup b) < V(a) + V(b). \quad (13)$$

By assumption, $\bar{V}(a \cup b)$ is a valid upper bound which satisfies (11). Combining these inequalities, we have

$$V(a \cup b) \leq \bar{V}(a \cup b) < V(a) + V(b), \quad (14)$$

so that any formed coalition offers a strict improvement in total cost compared to agents working individually. Since our proposed algorithm forms coalitions iteratively, starting from the decentralised solution, and each coalition formation strictly reduces the total cost, the final coalition structure will have a total cost at most equal to the decentralised solution. Equality holds only when no coalitions are formed. \blacksquare

1) A Practical Bound: A very computationally cheap option for this bound, that is also non-trivial, is to share only our current planned net grid trades (calculated in a decentralised manner). We then compare the grid consumption at each time step to check if some gain can be made. First, define

$$Q_j(t) = \sum_{i \in \mathcal{A}_j} G_i^{\text{tot}}(t), \quad (15)$$

the total grid trade by a coalition. Then, check if the sequence

$$\Xi_{jk}(t) = Q_j(t)Q_k(t), \quad (16)$$

contains any negative values (corresponding to one coalition buying (positive Q) and the other selling (negative Q) at the same time step).

Then, the upper bound on coalition cost $V(\mathcal{A}_j \cup \mathcal{A}_k)$ is

$$V(\mathcal{A}_j) + V(\mathcal{A}_k) - \sum_{t: \Xi_{jk}(t) < 0} (P^{\text{buy}}(t) - P^{\text{sell}}(t)) \min\{|Q_j(t)|, |Q_k(t)|\}, \quad (17)$$

i.e. the sum across time steps of possible trades, if no further optimisation is performed. From this we can see that

$$\text{val}(\tilde{\mathcal{C}}) = \sum_{t: \Xi_{jk}(t) < 0} (P^{\text{buy}}(t) - P^{\text{sell}}(t)) \min\{|Q_j(t)|, |Q_k(t)|\} \quad (18)$$

Corollary 1: The bound provided in (17) is a non-trivial upper bound to the coalitional cost.

Proof: We show that the quantity in (17) is achievable by explicitly constructing a feasible coalition strategy whose cost is exactly no greater than that bound. Since the optimal coalitional cost is no larger than the cost of any feasible strategy, this proves that (17) is an upper bound. We first prove the result for two individual buildings i and j . The extension to general coalitions follows by linearity of the aggregate net grid exchange, with variables replaced by their coalitional aggregate values (e.g. $Q_i(t)$ in place of $G_i(t)$).

We have that the strategies G_i, u_i and G_j, u_j must already be feasible strategies for agent i and j , since they are solutions to the decentralised optimisation (with $C_i = C_j = 0$). We fix u_i and u_j and do not alter these, and instead focus on shifting trades from the grid to the coalition.

For each time step t , if $\Xi_{ij}(t) < 0$, one building is buying from the grid while the other is selling at the same time step. In the decentralised solution, this results in simultaneous grid purchase at price $P^{\text{buy}}(t)$ and grid sale at price $P^{\text{sell}}(t)$. Thus, we consider moving $\Delta(t) = \min\{|G_i^{\text{tot}}(t)|, |G_j^{\text{tot}}(t)|\}$ trade from the grid to the coalition. By instead directly exchanging $\Delta(t)$ units of energy between the two buildings, we eliminate both the corresponding grid purchase and sale, yielding a cost reduction of $\Delta(t)(P^{\text{buy}}(t) - P^{\text{sell}}(t))$. We then have $C_i^{\text{tot}}(t) = \text{sign}(G_i^{\text{tot}}(t)) \cdot \Delta(t)$, and $C_j^{\text{tot}}(t) = -C_i^{\text{tot}}(t)$ so that $C_i^{\text{tot}}(t) + C_j^{\text{tot}}(t) = 0$, satisfying our coalition power balance. Since $(G_i(t), u_i(t))$ was feasible, and $\hat{G}_i^{\text{tot}}(t) + C_i^{\text{tot}}(t) = G_i^{\text{tot}}(t)$, then the modified strategy (\hat{G}_i, u_i, C_i) also satisfies the power balance (with demand and generation profiles unchanged).

We now only modify our trades as $\hat{G}_i^{\text{tot}}(t) = G_i^{\text{tot}}(t) - C_i^{\text{tot}}(t)$, and $\hat{G}_j^{\text{tot}}(t) = G_j^{\text{tot}}(t) + C_i^{\text{tot}}(t) = G_j^{\text{tot}}(t) - C_j^{\text{tot}}(t)$. Thus, at time step t , the combined cost reduction is

$$(P^{\text{buy}}(t) - P^{\text{sell}}(t)) [G_i^{\text{tot}}(t) + G_j^{\text{tot}}(t) - \hat{G}_i^{\text{tot}}(t) - \hat{G}_j^{\text{tot}}(t)] = \Delta(t)(P^{\text{buy}}(t) - P^{\text{sell}}(t)), \quad (19)$$

where the quantity on the left is exactly equal to the cost reduction for buildings i and j . After the modification, building i reduces its grid purchase by $\Delta(t)$, and building j reduces its grid sale by $\Delta(t)$. By summing across time

steps where $\Xi_{ij}(t) < 0$, we obtain feasible coalitional cost

$$V(\{i\}) + V(\{j\}) - \sum_{t: \Xi_{ij}(t) < 0} \left(P^{\text{buy}}(t) - P^{\text{sell}}(t) \right) \Delta(t).$$

Since the optimal coalitional cost $V(i, j)$ is no larger than the cost of any feasible coalition strategy, we obtain the bound in (17). Thus, this bound is achievable and is a valid upper bound on the coalitional cost. ■

D. Within Coalition Optimisation

Once a coalition structure is formed, we must also decide how to optimise the joint strategy of the coalition. Solving the optimisation problem centrally raises similar information sharing concerns to Algorithm 1, as it requires full information from all agents. To address this, we adopt a distributed solution based on ADMM [11], which enables agents to compute their local decisions while coordinating only through limited information exchange. We reformulate the coalition optimisation problem in a consensus form, allowing each agent to solve a local subproblem while enforcing agreement on shared variables through a coordinator.

We introduce auxiliary variables $\mathbf{z}_i \in \mathbb{R}^T$ for each agent i , representing the consensus coalition trade schedule across the prediction horizon, and dual variables $\lambda_i \in \mathbb{R}^T$. These variables enforce the coalition power balance constraint $\sum_i z_i(t) = 0$ at each time step while allowing agents to solve their optimisation problems locally. For each agent i in coalition \mathcal{C}_l , we make use of the following update steps

$$\mathbf{G}_{i,k+1}, \mathbf{u}_{i,k+1}, \mathbf{C}_{i,k+1} \quad (20a)$$

$$= \arg \min_{\{\mathbf{G}_i, \mathbf{u}_i, \mathbf{C}_i\} \in \mathfrak{X}_i} L_i(\mathbf{G}_i) + \lambda_{i,k}^\top \mathbf{C}_i^{\text{tot}} + \frac{c}{2} \|\mathbf{C}_i^{\text{tot}} - \mathbf{z}_{i,k}\|^2$$

$$\mathbf{z}_{i,k+1} = \mathbf{C}_{i,k+1}^{\text{tot}} - \frac{1}{|\mathcal{C}_l|} \sum_{j \in \mathcal{C}_l} \mathbf{C}_{j,k+1}^{\text{tot}} \quad (20b)$$

$$\lambda_{i,k+1} = \lambda_{i,k} + c(\mathbf{C}_{i,k+1}^{\text{tot}} - \mathbf{z}_{i,k+1}), \quad (20c)$$

where (20a) is performed in parallel by all agents, and (20b, 20c) are performed centrally (the coordinator initialises λ_i and shares the updated value at each iteration). We use $\mathbf{C}_i^{\text{tot}}$ to denote the stacked vector of coalitional trades for agent i , and c is a user-selected regularisation parameter (guidance on selecting this parameter is available in [11]). Note that this requires sharing only tentative coalition trades at each iteration, thereby ensuring that individual consumption and generation data remain private. Since the problem is convex and can be written in a consensus form, standard ADMM convergence results apply [11]

Since ADMM enforces the coalition balance constraint only asymptotically, early termination may result in violations of the constraint $\sum_{i \in \mathcal{C}_j} \mathbf{C}_i^{\text{tot}} = \mathbf{0}$. To ensure feasibility, we take the coalition trades from \mathbf{z}_i , which satisfy this constraint by construction. The local grid trades G_i are then recomputed to restore the power balance constraint for each agent, while satisfying their non negativity constraints and keeping the control inputs u_i fixed. This yields a feasible solution, at the expense of suboptimality. The local power balance is restored by updating the grid trades while keeping

the control inputs u_i unchanged, ensuring that the power-balance constraints are satisfied at each time step.

IV. MODEL PREDICTIVE CONTROL

Our methods up to this point have dealt with open-loop optimisation of the energy purchasing problem. We now consider applying our methods in a model predictive control (MPC) framework.

We introduce predicted consumption and generation profiles $\hat{d}(t+k|t), \hat{g}(t+k|t)$ for $k = 0, \dots, \min\{T_{\text{pred}}, T-t\}$, where $(\cdot|t)$ denotes a value predicted at step t . We assume that energy prices and battery dynamics are known precisely, so that this prediction is the only form of uncertainty.

Then, for every time step, we solve the open-loop optimisation problem using this predicted behaviour. We utilise the first planned battery charge signal $u(t)$ and (if applicable) the first coalition trade $C(t)$, and then purchase (or sell) energy as necessary to satisfy the power balance in (2). Note that, if our predictions are inaccurate, we may trade a different amount of energy with the grid than planned. Once we employ a strategy at time step t , we uncover the realised stage cost

$$R(\mathcal{C}_j)(t) = \sum_{i \in \mathcal{C}_j} P^{\text{buy}}(t) G_i^{\text{buy}}(t) - P^{\text{sell}}(t) G_i^{\text{sell}}(t). \quad (21)$$

When employing a coalition formation algorithm, we use predicted data to solve the coalition formation problem. Coalitions are reformed at every time step to adapt to changing conditions, and, after coalitions are formed, we check if any violate the following inequality

$$R(\mathcal{C}_j)(t) \leq \sum_{i \in \mathcal{C}_j} R(\{i\})(t). \quad (22)$$

If this condition is violated, then this coalition is dissolved and constituent agents work individually. In this way, we guarantee that our algorithm performs no worse than the decentralised optimum. As in the decentralised case, if predictions are inaccurate, we trade different amounts with the grid, maintaining coalition trades and charging signal.

Theorem 2 (MPC with Coalition Formation): Consider the predictions $\hat{d}(t|t) = d(t)$, $\hat{g}(t|t) = g(t)$. Let $\{\mathcal{C}_j(t)\}_{j=1}^{M(t)}$ be the coalition structure formed at time t , and define

$$J_{\text{coal}}^{\text{MPC}} = \sum_{t=1}^T \sum_{j=1}^{M(t)} R(\mathcal{C}_j(t))(t) \quad (23)$$

$$J_{\text{dec}}^{\text{MPC}} = \sum_{t=1}^T \sum_{i=1}^N R(\{i\})(t). \quad (24)$$

Where, in the coalitional case, any coalitions not satisfying (22) are dissolved and agents work individually on that time step. Then, despite coalitions being reformed at every time step,

$$J_{\text{coal}}^{\text{MPC}} \leq J_{\text{dec}}^{\text{MPC}}. \quad (25)$$

Proof: At each time step t , we enforce:

$$R(\mathcal{C}_j(t))(t) \leq \sum_{i \in \mathcal{C}_j(t)} R(\{i\})(t). \quad (26)$$

TABLE I: Average costs for open-loop (TO denotes timeout)

Scheme	Cost (euros)	
	70 Buildings	8 Buildings
Decentralised	239	454
Centralised	101	430
Limited Information	232	430
Bottom-Up	157	431
Optimal Coalition	TO	430

Summing over all time steps and coalitions we have

$$\sum_{t=1}^T \sum_{j=1}^{M(t)} R(\mathcal{C}_j(t))(t) \leq \sum_{t=1}^T \sum_{j=1}^{M(t)} \sum_{i \in \mathcal{C}_j(t)} R(\{i\})(t), \quad (27)$$

which gives $J_{\text{coal}}^{\text{MPC}} \leq J_{\text{dec}}^{\text{MPC}}$. ■

V. NUMERICAL RESULTS

We make use of a real-world dataset [15], [16] to perform experiments, with energy usage recorded for 15-minute time slots, predictions for both generation and demand, battery data and energy prices taken from this dataset¹. Whenever we seek to solve (9) for $V(\mathcal{C}_j)$, we perform ADMM until the primal residual is less than 10^{-5} , and use a penalty parameter $c = \frac{1}{2}$. The experiments, therefore, focus on assessing the impact of limited-information coalition formation, rather than differences in the underlying optimisation solver. We use a coalition size limit that allows computations to be completed within a reasonable time frame, while offering reasonable cost reductions (for our limited information approach, we compare different coalition sizes in Section V-D). All simulations were performed using 2 CPU cores at 2.6 GHz and 4 GB RAM.

A. Open Loop Optimisation

Results are presented in Table I for the open loop optimisation problem using first all 70 buildings, and then 8 buildings, with a maximum coalition size of 7 buildings, and planning for a single day ahead. For 70 agents, the optimal coalition problem was terminated after being unable to find a solution within 24 hours.

As can be seen, our algorithm offers an improvement over the decentralised scheme in both cases, and is even able to perform optimally when solving with 8 buildings. Despite the Bottom-Up algorithm outperforming our algorithm for the complete set of buildings, we will see in Section V-C, this comes at a cost of increased computation, and greater information sharing. The brute force approach to identify the true optimum coalition requires substantial computation (as discussed in Section III), and was unable to finish within 24 hours for 70 buildings.

¹Code is available at https://github.com/lukearcus/Coalition_grid.

TABLE II: Average costs for MPC (TO denotes timeout)

Scheme	Cost (euros)		Runtime (s)
	Cost (euros)	ADMM Iterations	
Decentralised	496	-	17
Centralised	475	14986	50700
Limited Information	489	1362	4067
Bottom-Up	478	18652	63432
Optimal Coalition	TO	-	-

B. Model Predictive Control

We now compare our techniques against existing methods when implementing a model predictive controller. In Table II we compare the cost achieved, number of ADMM iterations per time step (including both for coalition formation, and within-coalition solving) and runtime, for 10 buildings across the full horizon (1 day) with a 2-hour prediction horizon, and with coalitions restricted to a maximum of 6 buildings. We see that our limited information algorithm achieves a lower cost than the decentralised scheme, without incurring a significant computational burden. In Figure 3, we see how the formation of coalitions affects trades with the grid. Under the coalition formation scheme, trades follow a similar envelope but spike toward zero when trades are beneficial as agents trade with each other, not the grid. Once again, the brute force approach to find the optimal coalition was too computationally intensive to terminate in time.

C. Scalability

In Figure 2, we provide a comparison of the average number of ADMM iterations needed per time step for a varying number of buildings N , using the same setup for MPC as the previous section. We allowed a maximum of 72 hours total computation time for each algorithm. As can be seen, our limited information coalition formation technique outperforms both other algorithms as the number of buildings becomes larger, and is still able to maintain relatively few iterations until convergence, even with 24 buildings.

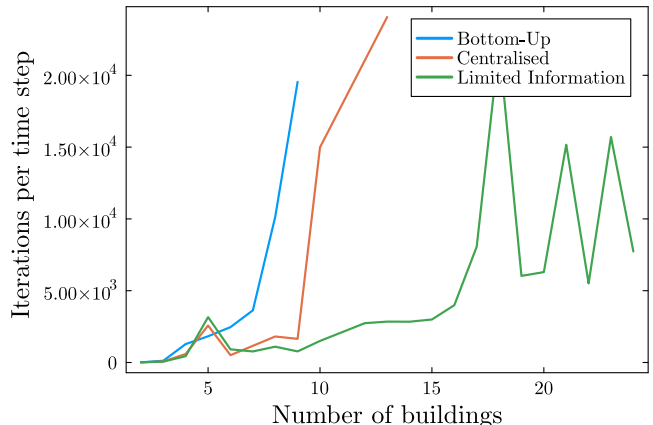




Fig. 2: Number of ADMM iterations vs Number of buildings for different schemes



(a) Decentralised



(b) Limited information coalitions

Fig. 3: Grid trades (in kWh, positive buy, negative sell) for a day under decentralised and limited information coalitions.

D. Coalition Size

For our proposed limited information scheme, we compare the effects of different choices for maximum coalition size in Figure 4. We used 20 buildings, a time horizon of 1 day, and a 2-hour lookahead. As we expect, the larger the coalition the more computationally intensive the problem becomes, but the lower the average cost. This suggests that a trade-off between coalition size and cost should be sought. For our previous experiments, we identified a maximum coalition size of six as offering a reasonable trade-off.

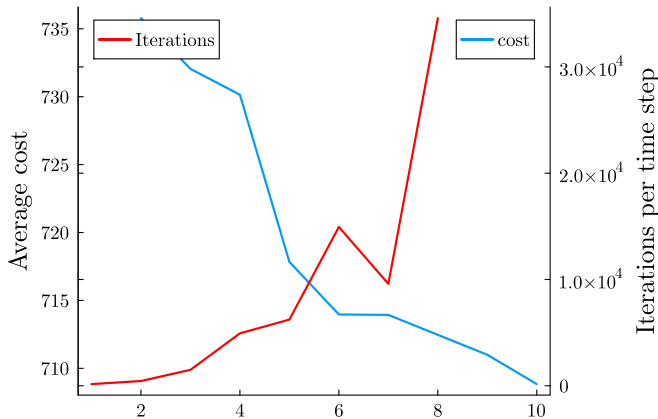


Fig. 4: Maximum coalition size affect on number of ADMM iterations, and average cost.

VI. CONCLUSIONS

Our results demonstrate that coalitional formation is achievable with limited information sharing while maintaining improved performance over the decentralised case. These preliminary results raise several avenues for future work, including analysing the trade-off between information sharing and optimality, exploring stronger guarantees than improvement over the decentralised optimum, extending to more general coalition formation problems and addressing uncertainty in future predictions for MPC.

REFERENCES

- [1] T. B. Nadeem, M. Siddiqui, M. Khalid, and M. Asif, "Distributed energy systems: A review of classification, technologies, applications, and policies," *Energy Strategy Reviews*, vol. 48, p. 101096, 2023.
- [2] W. L. Filho, L. V. Trevisan, A. L. Salvia, J. Mazutti, T. Dibbern, S. R. de Maya, E. F. Bernal, J. H. P. P. Eustachio, A. Sharifi, M. del-Carmen Alarcón-del Amo, and I. Kushnir, "Prosumers and sustainable development: An international assessment in the field of renewable energy," *Sustainable Futures*, vol. 7, p. 100158, 2024.
- [3] Y. Parag and B. K. Sovacool, "Electricity market design for the prosumer era," *Nature Energy*, 2016.
- [4] F. Fele, J. M. Maestre, and E. F. Camacho, "Coalitional control: Co-operative game theory and control," *IEEE Control Systems Magazine*, vol. 37, no. 1, pp. 53–69, 2017.
- [5] F. Fele, J. M. Maestre, and E. F. Camacho, "Coalitional control: A bottom-up approach," in *ACC*, pp. 4074–4079, IEEE, 2015.
- [6] G. Giaconi, D. Gunduz, and H. V. Poor, "Privacy-aware smart metering: Progress and challenges," *IEEE Signal Processing Magazine*, vol. 35, no. 6, pp. 59–78, 2018.
- [7] R. Razavi and A. Gharipour, "Rethinking the privacy of the smart grid: What your smart meter data can reveal about your household in ireland," *Energy research & social science*, vol. 44, pp. 312–323, 2018.
- [8] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, pp. 15–27, 2012.
- [9] P. R. Baldvieso-Monasterios and P. A. Trodden, "Coalitional predictive control: Consensus-based coalition forming with robust regulation," *Automatica*, vol. 125, p. 109380, 2021.
- [10] A. Sánchez-Amores, P. Chanfreut, J. Maestre, and E. F. Camacho, "Robust coalitional model predictive control with negotiation of mutual interactions," *Journal of Process Control*, vol. 123, pp. 64–75, 2023.
- [11] S. P. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [12] L. Lanza, T. Faulwasser, and K. Worthmann, *Distributed optimization for energy grids: a tutorial on ADMM and ALADIN*, ch. Chapter 5, pp. 121–145.
- [13] Y. Zhang, M. Hong, E. Dall'Anese, S. V. Dhople, and Z. Xu, "Distributed controllers seeking AC optimal power flow solutions using ADMM," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4525–4537, 2018.
- [14] A. Joshi, H. Kebriaei, V. Mariani, and L. Glielmo, "A sufficient condition to guarantee non-simultaneous charging and discharging of household battery energy storage," *CoRR*, vol. abs/2104.06267, 2021.

- [15] A. Le Franc and M. Ittel, "The EMSx dataset." <https://zenodo.org/records/5510400#.YUizGls69hE>, 2021.
- [16] A. Franc, P. Carpentier, J.-P. Chancelier, and M. De Lara, "EMSx: a numerical benchmark for energy management systems," *Energy Systems*, vol. 14, pp. 1–27, 02 2021.