

Discrete-choice Multi-agent Optimization: Decentralized Hard Constraint Satisfaction for Smart Cities

Srijoni Majumdar
University of Leeds, Leeds, UK
s.majumdar@leeds.ac.uk

Chuhao Qin
University of Leeds, Leeds, UK
sccq@leeds.ac.uk

Evangelos Pournaras
University of Leeds, Leeds, UK
e.pournaras@leeds.ac.uk

ABSTRACT

Making Smart Cities more sustainable, resilient and democratic is emerging as an endeavor of satisfying hard constraints, for instance meeting net-zero targets. Decentralized multi-agent methods for socio-technical optimization of large-scale complex infrastructures such as energy and transport networks are scalable and more privacy-preserving by design. However, they mainly focus on satisfying soft constraints to remain cost-effective. This paper introduces a new model for decentralized hard constraint satisfaction in discrete-choice combinatorial optimization problems. The model solves the cold start problem of partial information for coordination during initialization that can violate hard constraints. It also preserves a low-cost satisfaction of hard constraints in subsequent coordinated choices during which soft constraints optimization is performed. Strikingly, experimental results in real-world Smart City application scenarios demonstrate the required behavioral shift to preserve optimality when hard constraints are satisfied. These findings are significant for policymakers, system operators, designers and architects to create the missing social capital of running cities in more viable trajectories.

KEYWORDS

distributed optimization, multi-agent systems, discrete choice, hard constraints, Smart Cities

ACM Reference Format:

Srijoni Majumdar, Chuhao Qin, and Evangelos Pournaras. 2023. Discrete-choice Multi-agent Optimization: Decentralized Hard Constraint Satisfaction for Smart Cities. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 8 pages.

1 INTRODUCTION

Setting hard constraints in how we consume, produce, distribute and manage urban resources becomes paramount for the sustainability of our cities [9]. Coordinated responses to climate change often aim to satisfy hard constraints for carbon footprint emissions and net-zero [29]. Smart Grid technologies are still under development because of challenges to satisfy hard operational constraints that can cause catastrophic power blackouts [24] (see Figure 1). The satisfaction of hard constraints is also the safeguards for safety and the social capital for trust in establishing autonomous self-driving cars at scale [7]. Currently, the complexity, scale and decentralization of socio-technical infrastructures in Smart Cities are a barrier for satisfying hard constraints by design. Instead, soft constraints prevail in the vast majority of optimization and learning

approaches applied to the broader spectrum of Smart City applications [13, 18, 22, 25]. This research gap is the focus and subject of this paper.




Figure 1: An illustrative optimization scenario. Baseline: Without any optimized demand response, the power peak causes a blackout. **Soft constraints:** Significant power peak reduction that does not though prevent the power blackout. **Hard constraints:** Guarantee the reduction of the power peak below the blackout threshold.

This paper studies the decentralized hard constraint satisfaction in discrete-choice multi-agent optimization, in particular distributed multi-objective combinatorial optimization problems. This is a large class of problems [22] in which agents autonomously determine a number of finite options to choose from (operational flexibility). The agents may have their own preferences over these alternatives, expressed with a *discomfort cost* for each option. However, such choices often turn out to be inter-dependent to minimize system-wide *inefficiency and unfairness costs* (non-linear cost functions) for which the agents may have (explicitly or implicitly) interest as well. These choices require coordination and computing the optimal combination of choices in an NP-hard problem [25]. This is especially the case when agents come with different levels of *selfish vs. altruistic* behavior, with which they prioritize the minimization of their individual discomfort cost over the collective inefficiency and unfairness cost. Smart Cities are full of emerging application scenarios that can be modelled as such optimization problems [22, 25, 28]: power peak reduction to avoid blackouts, shift of power demand to consume more available renewable energy resources, coordinated vehicle routing to decrease travel times, traffic jams and air pollution, coordinated swarms of Unmanned Aerial Vehicles (UAVs) for distributed sensing, load-balancing of bike sharing stations, and other.

So far, heuristics for solving these distributed multi-agent optimization problems mainly address soft constraints, which is the best effort to minimize all involved costs. This is because in the absence of complete information in the agents, it is simpler to design algorithms that search efficiently for good solutions even if these are not the optimal ones. Instead, satisfying hard constraints opens up a Pandora box: without full information, any autonomous agent choice can violate the hard constraints that can only be satisfied with certainty at an aggregate level. Letting agents make

conservative choices to avoid violating the hard constraints may significantly downgrade performance and optimality, while any rollback of choices violating these constraints is complex and costly.

To address this timely problem with impact on Smart Cities, a new decentralized hard constraint satisfaction model is introduced. The model constructs ranges of upper and lower bounds within which the aggregate choices and costs must remain, while optimizing for soft constraints. To solve the cold start problem in the initialization phase during which agents choices and costs are undergoing aggregation, a heuristic is introduced for the agents to make choices with the highest average likelihood to satisfy all hard constraints. As this heuristic is sensitive to the agents' self-organization (order) in decision-making and aggregation, agents keep reorganizing themselves after violations of hard constraints as long as a stopping criterion is not reached. After the cold start phase and after agents successfully satisfy the hard constraints, they can shift entirely to the optimization of the soft constraints while preserving the satisfaction of the hard constraints locally with a low cost.

The proposed model is integrated into a collective learning algorithm, the *Iterative Economic Planning and Optimized Selections* (I-EPOS) [25]. This allows a comprehensive assessment of the decentralized hard constraint satisfaction model and its impact on the optimality of the soft constraints. For the first time, experimental evaluation with real-world data from Smart City scenarios disentangle the performance sacrifice as a result of satisfying hard constraints and the additional agents' altruism level required to mitigate such sacrifice. These findings are highly revealing for system operators, policymakers, system designers and architects. They can inform them about the additional social capital (incentives/rewards) that they require to build (and pay) to preserve the cost-effectiveness of socio-technical infrastructures operating with hard constraints.

The contributions of this paper are summarized as follows: (i) A model of decentralized hard constraint satisfaction on optimizing aggregate agents' choices and their aggregate costs. (ii) The instantiation of this model on a decentralized multi-objective combinatorial optimization algorithm of collective learning for multi-agent systems. (iii) The applicability of decentralized hard constraint satisfaction on three Smart Cities scenarios using real-world data: energy, bike sharing and UAVs swarm sensing. (iv) Insights about the optimality sacrifice as moving from soft to hard constraints and how this optimality loss is measured in terms of the required behavioral shift to preserve performance, i.e. restoring altruism deficit. (v) An open-source software artifact implementation of the model for the I-EPOS collective learning algorithm [16].

This paper is summarized as follows: Section 2 reviews related methods. Section 3 introduces the decentralized hard constraint satisfaction model. Section 4 illustrates the applicability of this model to the collective learning algorithm of I-EPOS and its implementation. Section 5 illustrates the experimental methodology and the evaluation. Section 6 concludes this paper and outlines future work.

2 COMPARISON TO RELATED WORK

We study a discrete-choice distributed multi-objective optimization problem for multi-agent systems with both soft and hard constraints.

In such systems, most related optimization approaches [12] operate as partially observable systems with agents communicating with their neighbors.

There are approaches that rely on an asynchronous hierarchical process using depth-first search to order agents that communicate with their parents to make choices and optimize objective functions [2]. Mailler et al. [15] cluster agents based on the constraints they attempt to satisfy, with a central controller that uses a branch and bound paradigm for searching solutions. These hierarchical approaches suffer from failure risks, performance bottlenecks, and potential privacy breaches in application scenarios involving sensitive personal data, e.g. location and health data.

Multi-agent reinforcement learning approaches with constraints on agent choices are earlier studied. For instance, Curran et al. [5] generate rewards for agents to optimize delay intervals that prevent air traffic congestion with greedy scheduling to implement *hard stop* (constraints) for agents when the delay surpasses a limit. Rollback to previous states (*warm restarts*) are earlier studied upon violation of global hard constraints in [20]. Simao et al. [30] learn non-violated execution by training using datasets containing constrained actions of the agents and corresponding global states of a centrally controlled environment. Even though the execution is decentralized, the learned model used to provide recommendations to the agents is an outcome of a centralized computation.

Decentralized and asynchronous versions of population search-based optimization methods, such as particle swarm optimization [1] or ant colony [8] algorithms show a slow convergence with high communication cost to rollback after violations of hard constraints, while improving global fitness and local search. This may slow down online real-time adaptations. Violation of hard constraints is prevented via message broadcasting that rolls back all choices made after a violation [2].

Other earlier approaches use tree overlay network structures for aggregation to aggregate messages from child nodes in form of hypercubes to reduce the frequency of message exchanges [4, 6, 14]. These methods use dynamic programming approach and thus storing all solutions increases the size of messages exponentially.

Other highly efficient discrete-choice multi-agent optimization methods, such as COHDA [11] and EPOS [25], address a large spectrum of NP-hard combinatorial problems in the domains of Smart Grids and Smart Cities [10, 18, 22, 26]. COHDA generalizes well in different communication structures among the agents that have full view of the systems, while EPOS focuses on hierarchical acyclic graphs such as trees to perform a cost-effective decision-making and aggregation of choices. Table 1 compares the design and efficiency of multi-agent optimization approaches, as well as how they address soft and hard constraints.

As COHDA shares full information between agents, it has higher communication cost. The computational cost is lower at global level for COHDA compared to EPOS because of the agents' brute force search to aggregate choices. Both COHDA and EPOS focus on satisfying soft constraints, like minimizing cost functions that satisfy balancing (minimum variance and standard deviation [25]) or matching (minimum root mean square error, residual sum of squares [26]) objectives. However, satisfying global hard constraints (Table 1) is challenging as agents need to additionally coordinate for choices, whose potential violations are only confirmed at an

Table 1: Comparison of self adaptive decentralized approaches

Attributes	I-EPOS [25]	EPOS [21]	COHDA [11, 17, 19]	H-DPOP [4, 14]
Plan Selection - Autonomy	Locally	Parent-Level	Locally	Parent-Level
Computational Cost	agent: $O(pt)$; system: $O(pt \log a)$	agent: $O(p^c)$; system: $O(p^c \log a)$	agent: $O(pt)$; system: $O(pt)$	agent: $O(p^c)$; system: $O(p^c)$
Communication Cost	agent: $O(t)$; system: $O(t \log a)$	agent: $O(p)$; system: $O(p \log a)$	agent: $O(at)$; system: $O(at)$	agent: $O(p^c)$; system: $O(p^c)$
Information Exchange	tree overlay; aggregate information;	tree overlay; aggregate messages;	k-connected graph; full information;	tree overlay; full information;
Soft Constraints	local (initialization), aggregated choices in global plan	local (initialization), aggregated choices in global plan	global	no soft constraints
Hard Constraints	local (initialization)	local (initialization)	local (initialization), global	local (initialization), global

p : number of plans (options), t : number of iterations, c : number of children, a : number of agents

aggregate level, which makes any rollback of choices to avoid violations particularly complex. An expensive rollback procedure earlier introduced in COHDA [17] performs complete rescheduling using a 0-1 multiple-choice combinatorial to find another solution that satisfy the constraints.

Summarizing, satisfaction of hard constraints during initialization phase, when agents accumulate information about other agents' choices (cold start problem) remains an open challenge. It is also unclear how the satisfaction of hard constraints degrades the performance of these efficient algorithms based on their soft constraints. Addressing these open questions is the focus of this paper.

3 HARD CONSTRAINT SATISFACTION MODEL

This section introduces the general optimization problem and the decentralized hard satisfaction model.

3.1 Optimization problem

Table 2 summarizes the mathematical symbols of this article. Assume a socio-technical systems of n users, each assisted by a software agent, i.e. a personal digital assistant. Each agent i has a number of k options to choose from. Each option j is referred to as a *possible plan*, which is a sequence of real values $p_{i,j} = (p_{i,j,u})_{u=1}^m \in P_i = (p_{i,j})_{j=1}^k, \forall i \in \{1, \dots, n\}$. Each agent selects one and only one plan $p_{i,s}$, which is referred to as the selected plan (i.e. the agent's choice). All agents' choices aggregate element-wise to the collective choice, the *global plan* $g = (g_u)_{u=1}^m = \sum_{i=1}^n p_{i,s}$ of the multi-agent system. A possible plan of an agent may represent a resource schedule or allocation, e.g. the energy consumed over time or the energy consumed from a certain supplier. Multiple possible plans for each agent represent alternatives and its operational flexibility. In the example of energy, the global plan represents the total energy consumption in the system over time or suppliers (see Figure 1).

Agents' choices are made based on different, often opposing criteria. Each agent has its individual preferences over the possible plans, measured by the *discomfort cost* $f_D(p_{i,j}) = D_{i,j}$ of each plan j , which also makes the mean discomfort cost in the system $f_D(p_{1,s}, \dots, p_{n,s}) = \frac{1}{n} \cdot \sum_{i=1}^n f_D(p_{i,s})$. Each agent can make independent choices to minimize their own discomfort cost. However, agents may also have interest to satisfy the following two general-purpose collective criteria: inefficiency cost $f_I(\sum_{i=1}^n p_{i,s}) = I_i$ and unfairness cost $f_U(D_{1,s}, \dots, D_{n,s}) = U_i$. If these cost functions are non-linear, meaning the choices of the agents depend on each other,

Table 2: Mathematical notations used in this paper.

Notation	Meaning
n	Number of agents
P_i	Set of possible plans for agent i
m	Plan size
k	Number of plans
$p_{i,j} \subset \mathbb{R}^m$	The j^{th} plan as sequence of m elements of agent i
$p_{i,s}$	Selected plan of agent i
$g = \sum_{i=1}^n p_{i,s}$	Global plan from selected plans of all n agents
β_i	Discomfort factor for agent i
α_i	Unfairness factor for agent i
r	Constraints satisfaction rate
I	Inefficiency cost
D	Discomfort cost
U	Unfairness cost
$f_D: \mathbb{R}^m \rightarrow \mathbb{R}$	Discomfort cost function
$f_I: \mathbb{R}^m \rightarrow \mathbb{R}$	Inefficiency cost function
$f_U: \mathbb{R}^m \rightarrow \mathbb{R}$	Unfairness cost function
$\mathcal{U} \subset \mathbb{R}^m$	Sequence of upper bound hard constraints
$\mathcal{L} \subset \mathbb{R}^m$	Sequence of lower bound hard constraints
$\mathbb{E}(p_{i,j}, \mathcal{U})$	Expected satisfaction for upper bound constraints
$\mathbb{E}(p_{i,j}, \mathcal{L})$	Expected satisfaction for lower bound constraints

the satisfaction of soft constraints, i.e. minimizing the inefficiency and unfairness cost, is a combinatorial NP-hard optimization problem [25]. Balancing (e.g. min variance) and matching objectives (e.g. min residual of sum squares) are examples for measuring inefficiency cost with a broad applicability in load-balancing application scenarios of Smart Cities: minimizing power peaks, shifting demand to times with high availability of renewable energy resources, rerouting vehicles to avoid traffic jams, etc. Table 3 show such a case, in which three agents have two options (plans). These plans may represent energy consumption choices while forming an optimal global plans that meets the available energy supply. The elements may signify the power consumption for the day and night. The agents choose plans with minimum dispersion between elements (soft constraints), but that leads to a suboptimal global plan of [7,13]. The global plan should also come with lower dispersion. The variance of the discomfort costs over the population of agents can measure the unfairness cost. Satisfying all of these (opposing) objectives depends on the selfish vs. altruistic behavior of the agents, e.g. whether they accept a plan with a bit higher discomfort cost to decrease inefficiency or unfairness cost. We can observe this in Table 3. If Agent C selects a plan ([6,2]) with higher energy requirement during the day, it achieves to minimize the inefficiency cost and an optimum global plan of [10,10] is achieved. Such multi-objective

trade-offs are modelled with the parameters α_i and β_i such that:

$$p_{i,s} = \underset{j=1}{\operatorname{argmin}}^k (1 - \alpha - \beta) \cdot f_1(p_{1,s} + \dots + p_{i,j} + \dots + p_{n,s}) \\ + \alpha \cdot f_U(D_{1,s}, \dots, D_{i,j}, \dots, D_{n,s}) \\ + \beta \cdot f_D(D_{1,s}, \dots, D_{i,j}, \dots, D_{n,s}). \quad (1)$$

From the above equation, it is apparent that the choice of a plan cannot be easily optimized without (i) information of the other agents' choices and (ii) coordination of the agents' choices for non-linear cost functions that depend on each other. The optimization heuristics (Section 2) address the satisfaction of such soft constraints via various sequential information exchange, information aggregation and communication schemes that coordinate agents' choices. See the baseline scenario of soft constraints in Table 3.

However, introducing hard constraints on the aggregated choices g and their costs $D_{i,j}$, I_i , U_i is challenging. This is because there is no guarantee to satisfy the hard constraints in the absence of full information, which is usually the case for decentralized heuristics that require time to converge to full information. This is a particular cold start problem of initialization/exploration, during which the first choices are made under high uncertainty. As choices with high likelihood of violating hard constraints add up incrementally, it becomes increasingly hard to discover choices that will prevent such violations. Hence, agents need different and more conservative selection criteria that prioritize hard over soft constraints. Designing and evaluating these criteria is a contribution of this paper.

3.2 A heuristic for satisfying hard constraints

The heuristic for satisfying the hard constraints on aggregate choices and their costs is illustrated in this section. Table 3 also illustrates an example of applying the heuristic in practice.

3.2.1 Constraints on aggregate choices. For each element g_u of a global plan g , a hard constraint is defined by a range (envelope) of an upper $\mathcal{U} = (\mathcal{U}_u)_{u=1}^m$ and lower $\mathcal{L} = (\mathcal{L}_u)_{u=1}^m$ bound, where \mathcal{U} , \mathcal{L} are also sequences of real values of equal size $|\mathcal{U}| = |\mathcal{L}| = |g| = m$. Each value u of the upper bound denotes that $\mathcal{U}_u \geq g_u$, whereas for the lower bound it holds that $\mathcal{L}_u \leq g_u$.

The selected plan expected to satisfy all hard constraints at the initialization phase, during which the aggregate choices (global plan g) are not known, is estimated as follows:

$$p_{i,s} = \underset{p_{i,j} \in P_i}{\operatorname{argmax}} \mathbb{E}(p_{i,j}, \mathcal{U}, \mathcal{L}), \quad (2)$$

where the expectation of satisfaction is given by:

$$\mathbb{E}(p_{i,j}, \mathcal{U}, \mathcal{L}) = \sum_{u=1}^m (\mathcal{U}_u - p_{i,j,u}) + \sum_{u=1}^m (p_{i,j,u} - \mathcal{L}_u). \quad (3)$$

3.2.2 Constraints on aggregate costs. The modeling for the hard constraints on the aggregate costs is exactly the same as the one of the aggregate choices, where the expected satisfaction for each of the costs of $f_D(p_{1,s}, \dots, p_{n,s})$, $f_1(\sum_{i=1}^n p_{i,s})$ and $f_U(D_{1,s}, \dots, D_{n,s})$ is calculated for upper and lower bounds with $|\mathcal{U}| = |\mathcal{L}| = 1$.

3.2.3 Constraint satisfaction rate. The effectiveness of the hard constraint satisfaction heuristic is measured by the satisfaction rate (r). This is the total number of satisfactions achieved out of a total number of trials made. These trials are often existing parameters of the optimization algorithms, for instance, random repetitions, or

the order with which agents aggregate choices made to coordinate and optimize their own choices.

4 HARD CONSTRAINTS IMPLEMENTATION

The model of decentralized hard constraint satisfaction is implemented and integrated into the I-EPOS collective learning algorithm¹ [25]. I-EPOS solves a large class of optimization problems, as formalized in Section 3.1. It is chosen due to its large spectrum of applicability in Smart City scenarios [22] as well as its superior performance in satisfying soft constraints [25]. Efficient coordinated choices are made using a self-organized [21] tree topology within which agents organize their interactions, information exchange and decision-making. I-EPOS benefits from the fact that trees are acyclic graphs: communication cost is very low and all exchanges in I-EPOS are at an aggregate level without double-counting. The coordination is a result of a more informed decision-making: each agent makes a choice taking into account the choices of a group of other agents (the tree branch underneath during initialization) or the choices of all agents at a previous time point (after initialization). Coordination evolves in multiple learning iterations, each consisting of a *bottom-up* and *top-down* phase. During the bottom-up phase, each agent chooses based on the new choices of the agents below and the choices of all agents in the previous iteration. However, each agent has an information gap: It has no information about the subsequent choices of the agents above in the tree. This problem is solved during the top-down phase, in which agents roll back (back propagation) to choices of the previous iteration as long as no costs reduction is achieved. Further information about the design of the I-EPOS collective learning algorithm is out of the scope of this paper and can be found in earlier work [25].

The decentralized hard constraint satisfaction model is implemented by filtering out the possible plans in Equation 1 that violate the given upper and lower bounds. However, in the first learning iteration, it is not possible to determine these plans that violate the hard constraints with certainty because each agent only knows about the aggregate choices of the agents underneath (and not the ones above). As a result, the root agent may end up having no possible plan that does not violate the hard constraints. To prevent the likelihood of these violations, the agents make more conservative choices according to Equation 2 during the first iteration, aiming at maximizing the expected satisfaction of the hard constraints. Once the hard satisfactions are satisfied, the agents switch back to plan selection according to Equation 1, while keep filtering plans that violate the hard constraints. The agents cannot violate the hard constraints in these subsequent learning iterations because they always have the option to roll back to the choices made at the end of the first learning iteration during which hard constraints are satisfied (via the top-down phase).

Figure 2 illustrates the implementation of the hard constraints model on the open-source I-EPOS software artifact [16]. The implementation of the cost function interfaces is extended to filter out plans that violate the hard constraints, as well as the selection based on the expected satisfaction principle of Equation 2. The hard constraints are controlled via the main input parameter file of I-EPOS (Java Properties). Constraints on aggregate choices and costs can

¹Available at: <https://github.com/epournaras/EPOS>.

Table 3: An example of a discrete-choice combinatorial optimization problem with three agents ($n = 3$), each with two plans ($k = 2, m = 2$). All combinations of possible plan selections make $2^3 = 8$ possible global plans. Hard constraints with an upper bound on the aggregate choices (global plan g) are introduced with an expected satisfaction of $\sum_{u=1}^m (\mathcal{U}_u - p_{i,j,u})$. (1) The baseline scenario is the soft constraints that minimize the inefficiency cost $f_i(g) \approx |g_{i,j,1} - g_{i,j,2}|$. The global plan [10, 10] is the one with the minimum inefficiency cost. (2) Agents choose plans that maximize the expected satisfaction. This results in the global plan of [6, 15] that satisfies the hard constraint $\mathcal{U} = [9,]$. (3) Similarly, the hard constraint $\mathcal{U} = [, 9]$ is satisfied with the global plan of [14, 9]. (4) Both new hard constraints $\mathcal{U} = [10, 13]$ are satisfied with the global plan [7, 13]. (5) The second hard constraint $\mathcal{U} = [9, 9]$ is violated by the selected global plan [7, 13]. \checkmark : constraint satisfaction; \times : constraint violation in g

Constraints	Agent A		Agent B		Agent C		All Possible Global Responses							
Agents' Plans (p)	[3, 5]	[2, 7]	[1, 3]	[5, 2]	[6, 2]	[3, 5]	[10,10]	[14,9]	[7,13]	[11,12]	[9,12]	[13,11]	[6,15]	[10,14]
1. Soft Constraints (Baseline)														
Minimize Inefficiency Cost	[3-5]=2	[2-7]=5	[1-3]=2	[5-2]=3	[6-2]=4	[3-5]=2								
Selected Plans	\checkmark		\checkmark		\checkmark									
Selected Global Plan							\checkmark							
2. Hard Constraints - $\mathcal{U} = [9,]$														
Maximize Expected satisfaction	(9-3)+0=6	(9-2)+0=7	(9-2)+0=7	(9-5)+0=4	(9-6)+0=3	(9-3)+0=6								
Selected Plans		\checkmark	\checkmark			\checkmark								
Selected Global Plan											\checkmark			
3. Hard Constraints - $\mathcal{U} = [, 9]$														
Maximize expected satisfaction	0+(9-5)=4	0+(9-7)=2	0+(9-3)=6	0+(9-2)=7	0+(9-2)=7	0+(9-5)=4								
Selected Plans	\checkmark			\checkmark	\checkmark									
Selected Global Plan								\checkmark						
4. Hard Constraints - $\mathcal{U} = [10, 13]$														
Maximize expected satisfaction	(10-3)+ (13-5)=15	(10-2)+ (13-7)=14	(10-1)+ (13-3)=19	(10-5)+ (13-2)=17	(10-6)+ (13-2)=15	(10-3)+ (13-5)=15								
Selected Plans	\checkmark		\checkmark			\checkmark								
Selected Global Plan								\checkmark						
5. Hard Constraints - $\mathcal{U} = [9, 9]$														
Maximize expected satisfaction	(9-3)+ (9-5)=10	(9-2)+ (9-7)=9	(9-1)+ (9-3)=14	(9-5)+ (9-2)=11	(9-6)+ (9-2)=10	(9-3)+ (9-5)=10								
Selected Plans	\checkmark		\checkmark			\checkmark								
Selected Global Plan									\times					

be activated and deactivated. Two input .csv files are introduced, one for each type of hard constraints. Both contain the upper/lower bounds and the coding of the operators.




Figure 2: Implementation of decentralized hard constraints satisfaction in I-EPOS [16].

5 EXPERIMENTAL EVALUATION

Table 4 illustrates the application scenarios and settings of the experimental evaluation. A number of 1000 agents run the I-EPOS collective learning algorithm [25]. They are self-organized [21] in a height-balanced binary tree. The algorithm repeats 200 times, each time with a different random positioning of the agents in the tree. This introduces different decision-making order with which agents coordinate their optimized choices. At each repetition, the algorithm runs for 40 iterations, which is usually sufficient for convergence [25]. Evaluation is performed in three Smart City application

scenarios: (i) energy demand-response, (ii) bike sharing and (iii) UAV swarm sensing. The optimized inefficiency cost function and the generation of plans are also outlined in Table 4.

Table 4: I-EPOS [25] parameterization for the three datasets.

Parameter	Energy	Bike Sharing	UAV Swarm
Num. of agents (n)	1000	1000	1000
Num. of plans (k)	10	1 to 24	64
Plans size (m)	144	98	64
Num. of repetitions	200	200	200
Num. of iterations	40	40	40
Inefficiency cost f_i	Min VAR	Min VAR	Min RMSE

5.1 Smart City application scenarios

Energy demand-response. This dataset is based on energy disaggregation of the simulated zonal power transmission in the Pacific Northwest Smart Grid Demonstrations Project [23, 26]. It contains 5600 consumers with their energy demand recorded every 5 min in a 12h span of a day. The goal is to perform power peak shaving to prevent blackouts [25] by minimizing the variance.

Bike sharing. The Hubway Data Visualization Challenge dataset consists of the trip records of the Hubway bike sharing system in Paris [23, 25]. The data contain 2300 users, each with a varying number of possible plans for the bike stations from which bikes are picked up or returned (98 stations in total). The goal is to keep the bike sharing stations load balanced by minimizing the variance of the global plan.




Figure 3: Optimization under soft and three levels of hard constraints in the energy demand-response scenario. Light-grey shaded areas represent the upper bound and dark-grey shaded areas the lower bound. Arrows point to violations of hard constraints.

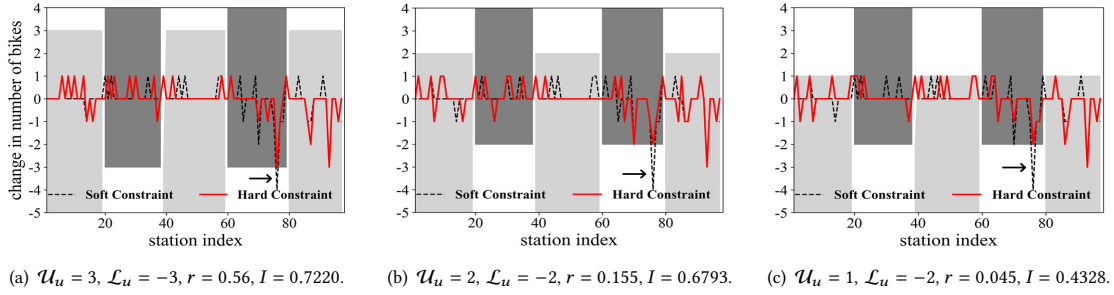


Figure 4: Optimization under soft and three levels of hard constraints in the bike sharing scenario. Light-grey shaded areas represent the upper bound and dark-grey shaded areas the lower bound. Arrows point to violations of hard constraints.




Figure 5: Optimization under soft and three levels of hard constraints in the UAV swarm sensing scenario. Light-grey shaded areas represent the upper bound and dark-grey shaded areas the lower bound. Arrows point to violations of hard constraints.

UAV swarm sensing. The dataset contains 1000 drones that can capture images or videos of vehicle traffic information on public roadways over 64 areas of interest (sensing cells) that are uniformly distributed in the city map [23, 27]. Drones aim to collect the required amount of sensing data (target plan) determined by a continuous kernel density estimation, for instance, monitoring cycling risk based on past bike accident data [3].

5.2 Hard constraint satisfaction works

For each application scenario, three incremental levels of hard constraints are set to the aggregate choices (global plan g). These levels are quantiles chosen empirically by observing the median global plan after several executions of I-EPOS based on soft constraints.

The agents are assumed here altruistic, such that: $\beta_i = 0, \alpha_i = 0, \forall i \in \{1, \dots, n\}$.

In the scenario of energy demand-response, hard constraints based on upper and lower bounds may represent an envelope of operation within which demand does not cause a blackout. In the scenario of bike sharing, hard constraints may represent limits on incoming or outgoing bikes that infrastructure operators may have, e.g. parking capacity. In the scenario of UAV swarm sensing, an upper bound of hard constraints may represent privacy-sensitive areas or regulated no-fly zones for drones. In contrast, a lower bound may represent minimal information required to monitor effectively a phenomenon, e.g. a forest fire or traffic jam.




Figure 6: Inefficiency and discomfort cost as a function of the altruism level for soft and hard constraints for the three Smart City application scenarios.




Figure 7: Required behavioral shift to mitigate the performance degrade of satisfying hard constraints for the three Smart City application scenarios. The satisfaction rate is also shown for each scenario. Performance comparison: β of soft constraints vs. β of hard constraints and satisfaction rate.

Figures 3, 4 and 5 show the global plans for the soft constraint (baseline) along with three incremental and alternating levels of hard constraints (upper/lower bounds). Under soft constraints, the upper and lower bounds are violated, whereas hard constraints prevent these violations. Stricter hard constraints prevent more violations, however, the satisfaction rate (r) drops significantly, while the inefficiency cost increases. This shows that strict hard constraints are likely to oppose the soft constraints. Last but not least, note that the scenario of UAV swarm sensing allows the satisfaction of a larger number of hard constraints, while preserving high satisfaction rates. This is because the agents have more operational flexibility by generating a larger number of plans ($64 > 24 > 10$).

5.3 Behavior shift can mitigate hard constraints

Satisfying hard constraints results in a degrade of the performance profile (lower inefficiency cost) achieved under soft constraints. The recovery from this degrade is measured here as the required social capital (behavioral shift) that agents need to offer such that soft and hard constraints have equivalent performance. The raise of social capital is measured by the reduction of the mean β_i value in the population of agents that makes them more altruistic, see Equation 1. The following method is introduced to measure the behavioral shift: (1) Perform parameter sweep on I-EPOS under soft constraints for $\beta_i = 0$, to $\beta_i = 1$, $\forall i \in \{1, \dots, n\}$ with a step of 0.025. (2) For each I-EPOS execution in Step 1 with a β_i value, a discomfort cost D and an inefficiency cost I , run I-EPOS under a hard constraint on the mean discomfort cost with an upper bound value equals

to D (the one of the I-EPOS execution under soft constraints). (3) Derive the increased inefficiency cost under the hard constraint on the discomfort cost. (4) Find the β_i value from Step 1 that has the closest inefficiency cost with the one derived in Step 3 under the hard constraint and (5) Compare the two β_i values in Step 2 and 4. The difference represents the required mean behavioral shift to mitigate the performance degrade of hard constraints.

Figure 6 illustrates the inefficiency and discomfort cost as a function of β_i under soft and hard constraints and the three different application scenarios. It becomes apparent that hard constraints require a minimum and significant level of altruism, otherwise, inefficiency cost rapidly explodes, especially in the scenarios without significant operational flexibility. This is also the reason why the discomfort cost becomes easier to reduce in the scenario of UAV swarm sensing, which comes with higher operational flexibility.

Figure 7 shows the required behavior shift to restore the performance loss as a result of satisfying hard constraints. For energy demand-response, the agents need on average 44.29% higher altruism under hard constraints to meet the performance of the soft constraints. The bike sharing scenario suggests an almost complete shift from selfish to altruistic behavior. Strikingly, the scenario of UAV swarm sensing shows performance gain as a result of satisfying hard constraints. As the number of plans is significantly higher for the UAV dataset, the search space is larger, which affects the optimality of the collective iterative learning paradigm in I-EPOS.

The satisfaction rate for the energy demand-response (Figure 7(a)), bike sharing (Figure 7(b)) and UAV swarm sensing (Figure 7(c)) are

54.45%, 19.44% and 61.21% respectively. For $\beta \leq 0.475$ and $\beta \leq 0.25$, the satisfaction rate is 100% for the UAV swarm sensing and energy demand-response respectively. The operational flexibility by higher number of possible plans increases the constraints satisfaction rate.

6 CONCLUSION AND FUTURE WORK

To conclude, this paper shows that the decentralized satisfaction of global hard constraints is feasible. It is a significant enabler for sustainability and resilience in several Smart City application scenarios such as energy demand-response to avoid blackouts, load balancing of bike sharing stations to make more accessible low-carbon transport modalities as well as improved sensing quality and efficiency by swarms of energy-constrained drones. Results show that hard constraints can be easily violated when optimizing exclusively for soft constraints. Instead, the proposed model prevents to a very high extent such violations.

Results also reveal the performance cost when hard constraints are introduced and how this cost can be mitigated by a behavioral shift towards a more altruistic behavior that sacrifices individual comfort for collective efficiency. These findings are invaluable for informing policy makers and systems operators of the required social capital that they need to raise to satisfy ambitious policies such as net-zero.

The open-source software artifact implementation of the proposed model to the I-EPOS collective learning algorithm is a milestone to encourage further research and application scenarios based on decentralized hard constraint satisfaction. Future work includes the applicability of the model to other decentralized optimization algorithms. The proposed heuristic is designed to satisfy all hard constraints together, which may be a limitation for high numbers of such opposing constraints, i.e. sacrifice of optimality and low satisfaction rates. Instead, a more incremental (and possibly partial) satisfaction of the hard constraints is part of future work. The further understanding of how to recover missing social capital to preserve both efficiency and fairness is also subject of future

ACKNOWLEDGEMENTS

This work is supported by a UKRI Future Leaders Fellowship (MR/W009560/1): ‘Digitally Assisted Collective Governance of Smart City Commons–ARTIO’, the Alan Turing Fellowship project ‘New Edge-Cloud Infrastructure for Distributed Intelligent Computing’ and the SNF NRP77 ‘Digital Transformation’ project ‘Digital Democracy: Innovations in Decision-making Processes’, #407740_187249.

REFERENCES

- [1] S Burak Akat and Veysel Gazi. 2008. Decentralized asynchronous particle swarm optimization. In *2008 IEEE Swarm Intelligence Symposium*. IEEE, 1–8.
- [2] Graham Billiau, Chee Fon Chang, and Aditya Ghose. 2012. SBDO: A new robust approach to dynamic distributed constraint optimisation. In *Principles and Practice of Multi-Agent Systems: 13th International Conference, PRIMA 2010, Kolkata, India, November 12–15, 2010, Revised Selected Papers 13*. Springer, 11–26.
- [3] David Castells-Graells, Christopher Salahub, and Evangelos Pournaras. 2020. On cycling risk and discomfort: urban safety mapping and bike route recommendations. *Computing* 102 (2020), 1259–1274.
- [4] Ziyu Chen, Chen He, Zhen He, and Minyou Chen. 2018. BD-ADOPT: a hybrid DCOP algorithm with best-first and depth-first search strategies. *Artificial Intelligence Review* 50 (2018), 161–199.
- [5] William J Curran, Adrian Agogino, and Kagan Tumer. 2013. Addressing hard constraints in the air traffic problem through partitioning and difference rewards.

In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. 1281–1282.

- [6] Yan Chen Deng, Ziyu Chen, Dingding Chen, Xingqiong Jiang, and Qiang Li. 2019. PT-ISABB: A hybrid tree-based complete algorithm to solve asymmetric distributed constraint optimization problems. In *International conference on Autonomous agents and multi-agent systems*. 1281–1282.
- [7] Huiying Du, Ge Zhu, and Jiali Zheng. 2021. Why travelers trust and accept self-driving cars: An empirical study. *Travel behaviour and society* 22 (2021), 1–9.
- [8] Arushi Gupta and Smriti Srivastava. 2020. Comparative analysis of ant colony and particle swarm optimization algorithms for distance optimization. *Procedia Computer Science* 173 (2020), 245–253.
- [9] Dirk Helbing, Farzam Fanitabasi, Fosca Giannotti, Regula Hänggli, Carina I Hausladen, Jeroen van den Hoven, Sachit Mahajan, Dino Pedreschi, and Evangelos Pournaras. 2021. Ethics of smart cities: Towards value-sensitive design and co-evolving city life. *Sustainability* 13, 20 (2021), 11162.
- [10] Christian Hinrichs et al. 2017. A distributed combinatorial optimisation heuristic for the scheduling of energy resources represented by self-interested agents. *International Journal of Bio-Inspired Computation* 10, 2 (2017), 69–78.
- [11] Christian Hinrichs, Sebastian Lehnhoff, and Michael Sonnenschein. 2014. CO-HDA: A combinatorial optimization heuristic for distributed agents. In *Agents and Artificial Intelligence: 5th International Conference, ICAART 2013, Barcelona, Spain, February 15–18, 2013. Revised Selected Papers 5*. Springer, 23–39.
- [12] ELSY Kaddoum. 2011. Optimization under Constraints of Distributed Complex Problems using Cooperative Self-Organization. *PhD thesis* (2011).
- [13] Shahzad Khan, Devashish Paul, Parham Momtahan, and Moayad Aloqaily. 2018. Artificial intelligence framework for smart city microgrids: State of the art, challenges, and opportunities. In *2018 third international conference on Fog and Mobile Edge Computing (FMEC)*. IEEE, 283–288.
- [14] Akshat Kumar, Adrian Petcu, and Boi Faltings. 2008. H-DPOP: Using Hard Constraints for Search Space Pruning in DCOP. In *AAAI*. 325–330.
- [15] Roger Mailler and Victor Lesser. 2004. Solving distributed constraint optimization problems using cooperative mediation. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004. AAMAS 2004*. IEEE, 438–445.
- [16] Srijoni Majumdar, Chuhaio Qin, and Evangelos Pournaras. 2023. *EPOS Hard Constraints Support*. <https://doi.org/10.5281/zenodo.7791326>
- [17] Astrid Nieße et al. 2014. Conjoint dynamic aggregation and scheduling methods for dynamic virtual power plants. In *2014 Federated conference on computer science and information systems*. IEEE, 1505–1514.
- [18] Astrid Nieße, Michael Sonnenschein, Christian Hinrichs, and Jörg Bremer. 2016. Local soft constraints in distributed energy scheduling. In *2016 Federated Conference on Computer Science and Information Systems (FedCSIS)*. IEEE, 1517–1525.
- [19] Astrid Nieße, Michael Sonnenschein, Christian Hinrichs, and Jörg Bremer. 2016. Local soft constraints in distributed energy scheduling. In *2016 Federated Conference on Computer Science and Information Systems (FedCSIS)*. IEEE, 1517–1525.
- [20] P Parnika, Raghuram Bharadwaj Diddigi, Sai Koti Reddy Danda, and Shalabh Bhatnagar. 2021. Attention actor-critic algorithm for multi-agent constrained co-operative reinforcement learning. In *International Conference on Autonomous Agents and MultiAgent Systems, ACM*.
- [21] Evangelos Pournaras. 2013. *Multi-level reconfigurable self-organization in overlay services*. PhD. Dissertation. Delft University of Technology. School of Technology Policy and Management.
- [22] Evangelos Pournaras. 2020. Collective learning: A 10-year odyssey to human-centered distributed intelligence. In *2020 IEEE International Conference on Autonomous Computing and Self-Organizing Systems (ACSOS)*. IEEE, 205–214.
- [23] Evangelos Pournaras. 2023. Agent-based Planning Portfolio. (1 2023). <https://doi.org/10.6084/m9.figshare.7806548.v5>
- [24] Evangelos Pournaras and Jose Espejo-Urbe. 2016. Self-repairable smart grids via online coordination of smart transformers. *IEEE Transactions on Industrial Informatics* 13, 4 (2016), 1783–1793.
- [25] Evangelos Pournaras, Peter Pilgerstorfer, and Thomas Asikis. 2018. Decentralized collective learning for self-managed sharing economies. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 13, 2 (2018), 1–33.
- [26] Evangelos Pournaras, Mark Yao, and Dirk Helbing. 2017. Self-regulating supply-demand systems. *Future Generation Computer Systems* 76 (2017), 73–91.
- [27] C Qin, F Candan, L Mihaylova, and E Pournaras. 2022. 3, 2, 1, Drones go! A testbed to take off UAV swarm intelligence for distributed sensing. In *Proceedings of the 2022 UK Workshop on Computational Intelligence*. Springer Nature.
- [28] Chuhaio Qin and Evangelos Pournaras. 2022. Coordination of Drones at Scale: Decentralized Energy-aware Swarm Intelligence for Spatio-temporal Sensing. *arXiv preprint arXiv:2212.14116* (2022).
- [29] Anu Ramaswami, Kangkang Tong, Josep G Canadell, Robert B Jackson, Eleanor Stokes, Shobhakar Dhakal, Mario Finch, Peraphan Jittrapirom, Neelam Singh, Yoshiki Yamagata, et al. 2021. Carbon analytics for net-zero emissions sustainable cities. *Nature Sustainability* 4, 6 (2021), 460–463.
- [30] Thiago D Simão et al. 2021. AlwaysSafe: Reinforcement learning without safety constraint violations during training. In *International Conference on Autonomous Agents and MultiAgent Systems, ACM*.