

Exposing Market Mechanism Design Trade-offs via Multi-objective Evolutionary Search

Arjun Chandra*, Richard Allmendinger[†], Peter R. Lewis[‡], Xin Yao[‡] and Jim Torresen*

*Department of Informatics, University of Oslo, Norway
{chandra|jimtoer}@ifi.uio.no

[†]Advanced Centre for Biochemical Engineering,
Department of Biochemical Engineering, University College London, UK
r.allmendinger@ucl.ac.uk

[‡]Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA),
School of Computer Science, University of Birmingham, UK
{p.r.lewis|x.yao}@cs.bham.ac.uk

Abstract—Market mechanisms are a means by which resources in contention can be allocated between contending parties, both in human economies and those populated by software agents. Designing such mechanisms has traditionally been carried out by hand, and more recently by automation. Assessing these mechanisms typically involves them being evaluated with respect to multiple conflicting objectives, which can often be non-linear, noisy, and expensive to compute. For typical performance objectives, it is known that designed mechanisms often fall short on being optimal across all objectives simultaneously. However, in all previous automated approaches, either only a single objective is considered, or else the multiple performance objectives are combined into a single objective. In this paper we do not aggregate objectives, instead considering a direct, novel application of multi-objective evolutionary algorithms (MOEAs) to the problem of automated mechanism design. This allows the automatic discovery of trade-offs that such objectives impose on mechanisms. We pose the problem of mechanism design, specifically for the class of linear redistribution mechanisms, as a naturally existing multi-objective optimisation problem. We apply a modified version of NSGA-II in order to design mechanisms within this class, given economically relevant objectives such as welfare and fairness. This application of NSGA-II exposes trade-offs between objectives, revealing relationships between them that were otherwise unknown for this mechanism class. The understanding of the trade-off gained from the application of MOEAs can thus help practitioners with an insightful application of discovered mechanisms in their respective real/artificial markets.

Keywords—*automated mechanism design, welfare, fairness, redistribution, market based interaction, resource allocation.*

I. INTRODUCTION

Market mechanisms lay out rules governing the allocation of resources amongst contending parties. Such parties could be humans interacting in real economies, whereby allocations of commodities via market mechanisms like auctions are commonplace. The parties can also be software agents in a multi-agent system, autonomously interacting with each other through a market mechanism to exchange resources [1], [2] that allow them to carry out tasks. Furthermore, software agents can also act on behalf of humans, for example in the case of grid and cloud computing, where agents are responsible for computational tasks using processors on the grid, or services in the cloud. Scheduling tasks or admitting

access to services via market based interaction mechanisms thus becomes an option [3], [4].

With such a prevalence of market mechanisms both in human and agent economies, market mechanism design is likely to remain a topic of great interest in economics and multi-agent systems. This is primarily due to the idea that according to most common mechanism design objectives, there frequently exists no single dominant mechanism [5]. Given this situation, the key research question becomes: *what mechanism should be chosen, given a current set of design objectives?* This paper aims to contribute towards an answer to this question.

In particular, we study a class of mechanisms called *linear redistribution mechanisms*, with the design objectives of allocative efficiency or *welfare maximisation* and *fairness*¹. Fairness can further be defined in multiple ways, and we look at two such definitions², namely *envy* and *disproportionality*. Welfare and fairness can be conflicting goals. For example, when the welfare of a system of agents is high, this can be due to a minority of the agents having high utility as compared to the rest, while fairness may entail (depending on the definition of fairness considered), a more even spread of utilities. Thus we have a three-objective automated mechanism design problem.

This problem naturally lends itself to being tackled using multi-objective evolutionary algorithms (MOEAs), yet previous work has mainly considered designing mechanisms by hand, or using a single (or a single aggregated) objective optimisation approach. We thus pose the mechanism design problem as a multi-objective optimisation problem (MOOP), and apply NSGA-II to expose the trade-off that the considered objectives impose on linear redistribution mechanisms. This reveals insights into the nature of the linear redistribution mechanism space hitherto unknown, additionally revealing unknown relationships between the considered objectives. The

¹Allocative efficiency is also known as welfare, and is used interchangeably in the paper. Welfare is usually defined as some variant of the sum of individual agent utilities, where agents derive their utilities from their possessions, namely their possessed allocations and money. The greater the utility, the more an agent may be able to carry out the tasks it is employed for.

²We consider two and not a single definition in order to stay consistent with previous work for this particular class of mechanisms, for fair comparisons with that work, and to build on it.

potential of using MOEAs for the problem of mechanism design is thus revealed.

Section II gives some background on how market mechanisms are usually designed. Section III elaborates on the class of mechanisms and the multiple performance objectives considered in this paper. The application of NSGA-II and the experimental setup are described in Section IV. Experimental results are presented in Section V, and the paper is concluded in Section VI.

II. DESIGNING MARKET MECHANISMS

Examples of market mechanisms often employed in human and agent economies include the English, Dutch and Vickrey auctions [6], in which an auctioneer facilitates the bidding and determines the allocation of resources. Where scarcity exists on both the seller and buyer sides, double auctions such as the Continuous Double Auction and Clearing House provide an alternative approach [7].

The choice of mechanism to be used is often determined by assumptions concerning the participants and market conditions. For example, double auctions may be adapted by modifying participation fees, transaction overhead fees or bid and ask matching strategies [8], to appeal to different types of participants. Wang [9] compares auction-based mechanisms and other market mechanisms such as posted offer markets in human economies. He claims that auctions are more commonly used where there is a greater dispersal of valuations of the good amongst the buyers. Interestingly, previous work in automated mechanism design has suggested that many mechanisms which do not appear at all intuitive to human auction designers may lead to more efficient outcomes in particular circumstances [5], [10], [11]. Despite this, the dominant approach to mechanism design is still by hand.

Cavallo [12] considers welfare as an objective and hand designs a mechanism within the linear redistribution mechanisms class (a class of auction mechanisms which is a variant of Vickrey auctions) that achieves high welfare. This class of mechanisms takes the idea of redistributing the auction revenue, that otherwise would remain with an auctioneer, back to the agents bidding for the auctioned resources. Scenarios where resources are jointly owned by agents, e.g. scheduling processor time on a jointly owned grid to computational tasks, such mechanisms can prove to be beneficial [3]. For the same class of mechanisms, Guo and Conitzer [13] take an automated approach and, via linear programming, optimise towards linear redistribution mechanisms to exhibit high welfare in the worst case. Both these approaches result in mechanisms that, upon evaluation w.r.t. welfare and fairness objectives, tend not to dominate each other. We describe both these works in more detail in Section III. In these prior works however, the (single) objective considered when evaluating a particular mechanism is allocative efficiency, or a variant of it.

This paper extends previous work in automated mechanism design, going beyond the approach of either considering the single objective of allocative efficiency (as discussed above) or combining multiple objectives into a single fitness measure [14]. In doing so, we present a general methodology for multi-objective automated mechanism design for arbitrary global objectives.

III. PRELIMINARIES

Let there be a set of agents $S = \{1, \dots, n\}$, and a set of outcomes O , where $o \in O$ is an n -tuple (o_1, o_2, \dots, o_n) , o_i being the allocation for agent $i \in S$. An agent i has a type $\theta_i \in \Theta$. The typespace Θ is the same for each agent, and represents the set of possible private information, e.g. preferences for various possible outcomes, that an agent might have. For n agents, the joint typespace can be denoted by Θ^n , where $\theta = (\theta_1, \theta_2, \dots, \theta_n) \in \Theta^n$. An agent of type θ_i has a valuation of $v_i(\theta_i, o_i)$ for an allocation o_i . The value functions are homogenous across the agents, such that $\forall i, j \in S, \forall \theta_j \in \Theta, \forall o \in O, v_i(\theta_j, o_j) = v_j(\theta_j, o_j)$.

We take the view of [15], and see a mechanism as a tuple (f, T) , where $f : \Theta^n \rightarrow O$ is a choice or allocation function, and $T = (T_1, \dots, T_n)$ represents an n -tuple of transfers (of currency amongst interacting agents), $T_i : \Theta^n \rightarrow \mathbb{R}$ associated with each agent i . Within a mechanism, agents report their types (which may or may not be their true types), and then the allocations and payment transfers are made according to f and T . Let o_i be denoted by $f_i(\theta)$, which is part of the outcome o that is allocated to agent $i \in S$ by f , for a type profile θ . As such, $f(\theta) = o = (o_1, o_2, \dots, o_n) = (f_1(\theta), f_2(\theta), \dots, f_n(\theta))$. We assume each agent $i \in S$ to be self-interested, and acting to maximise their *quasilinear utility*³ function. This utility for an agent is composed of the value v_i the agent obtains from an allocation o_i , as well as the transfers to and from the agent, encapsulated in T_i . More specifically, given the mechanism (f, T) , and true joint type θ , agent i obtains utility $v_i(\theta_i, f_i(\theta)) + T_i(\theta)$, provided all agents report their true types.

Let $v(\theta, o)$ denote $\sum_{i \in S} v_i(\theta_i, o_i)$, and $v^{-i}(\theta^{-i}, o)$ denote $\sum_{j \in S \setminus \{i\}} v_j(\theta_j, o_j)$. An efficient allocation function⁴ f^* is one which gives us an allocation o such that $\forall \theta \in \Theta^n, f^*(\theta) \in \arg \max_{o \in O} v(\theta, o)$. Let us call $f^*(\theta^{-i})$ an efficient allocation if we disregard the preferences of agent i . In other words, assuming that agent i did not exist ($i \notin S$), and regarding the allocation o_i , which would have been made to agent i , to be in contention amongst the remaining $n - 1$ agents, the efficient allocation for the remaining agents is $f^*(\theta^{-i}) \in \arg \max_{o \in O} v(\theta^{-i}, o)$, where $v(\theta^{-i}, o) = \sum_{j \in S, i \notin S} v_j(\theta_j, o_j)$. Vickrey-Clarke-Groves (VCG) mechanisms, variants of which are adopted in this work, comprise the efficient allocation function $f^*(\theta)$, and a transfer function $T_i(\theta) = v^{-i}(\theta^{-i}, f^*(\theta)) - v^{-i}(\theta^{-i}, f^*(\theta^{-i}))$. Note that $T_i(\theta)$ is the negative of the social cost $v^{-i}(\theta^{-i}, f^*(\theta^{-i})) - v^{-i}(\theta^{-i}, f^*(\theta))$ that an agent i imposes on the remainder of the agents, and thus a payment in this case.

We will solely be concerned with mechanisms that adhere to the following properties: *strategy-proof*, *non-deficit*, and *ex post individually rational*. A mechanism is *strategy-proof* if each agent is best off reporting its true type, regardless of its type and what other agents might report. A *non-deficit* mechanism is one where the total payment transfers aggregated across all agents never exceeds zero, i.e. $\sum_i T_i(\theta) \leq 0$. An *ex*

³Quasilinearity allows for agent valuations for an allocation to be interpreted as their *willingness to pay* [15] for it, and makes comparing between agent valuations possible.

⁴An efficient allocation function allocates resources to agents valuing them most.

post individually rational mechanism is one where no agent is worse off from participating. The VCG mechanism has these properties. Needless to say, such properties are often required in mechanisms. Strategy-proofness prevents strategic (including malicious) behaviour on the part of an agent, thus any cheating is prevented. A deficit would mean that an auctioneer would incur a loss or end up in debt from conducting the mechanism, and in the context of this paper, agents collectively receiving more money than what they spend for gaining access to resource units. Finally, an agent would not participate in a mechanism if they were to incur a loss from it, so mechanisms should ensure that the worst an agent can be from participating is to remain in the same position as they were before participating. We now describe the resource allocation scenario considered throughout this paper, the class of mechanisms (linear redistribution mechanisms) considered, and objectives (welfare and fairness) using which we expose trade-off mechanisms within the considered class.

A. Resource Allocation Scenario

The scenario considered is that of allocating a set $R = \{1, \dots, m\}$ of *homogenous* or *indistinguishable* resource units amongst n agents. Furthermore, an agent cannot be allocated more than one resource unit, also known as *unit demand*. This scenario has previously been considered in [13], [15], [16]. For the experimental results presented in this paper however, we focus on the case where $m = 1$ unit of resource, which may be in contention amongst n agents. Naturally, an agent cannot be allocated more than one resource unit in this case. An agent's type is specified by the value it associates with the unit of resource in contention. The scenario also implies that the agent's type can be given by a single value $\theta_i \in \mathbb{R}$. Furthermore, for comparisons with previous work [15], we let $\theta_i \in [0, 1]$.

1) *VCG mechanism in the scenario*:: As an example, let $n = 3$, and $m = 1$, with the agents having types $\theta_1 = 1.0$, $\theta_2 = 0.5$, and $\theta_3 = 0.3$ respectively. Under the VCG mechanism, this would lead to agent 1 getting the available unit. The payments, i.e. $v^{-i}(\theta^{-i}, f^*(\theta^{-i})) - v^{-i}(\theta^{-i}, f^*(\theta))$, by the agents, would be $-T_1(\theta) = 0.5 - 0.0 = 0.5$, while $-T_2(\theta)$ and $-T_3(\theta)$ remain zero. If agent i were the agent with the i th highest reported value θ_i , then we can write $\theta_1 \geq \theta_2 \geq \dots \geq \theta_n$. In the homogenous unit demand scenario with $n > m$, the payment by agent $i \forall i \in \{1, \dots, m\}$, $-T_i(\theta) = \theta_{m+1}$, and $\forall i \in \{m+1, \dots, n\}$, $-T_i(\theta) = 0$. The total VCG auction revenue (aggregate payment) in this case is $\sum_i -T_i(\theta) = m\theta_{m+1}$. In the case where $m = 1$, the VCG mechanism is also known as the Vickrey auction [6], wherein, the highest bidder wins the unit at a price determined by the second highest bid.

B. Redistribution Mechanisms

If the resources are jointly owned by the agents, instead of them being owned by a third party seller or auctioneer who can collect the auction revenue, redistribution mechanisms come to the fore. Allocating processor time on a jointly owned server for executing computational tasks, each task owned by a human who delegates an agent to have the task done [3], is one example of such joint ownership. In redistribution mechanisms, the auction revenue can be redistributed amongst

the agents in some fashion, instead of having it all thrown away. As an example, not redistributing the VCG revenue adversely affects the system by reducing its social welfare [15], as the payments are transferred outside the set of agents, thus lowering the utility of those making payments. Therefore, redistributing, as opposed to throwing away the revenue, has some merit. In this paper, redistribution mechanisms that allow returning portions of the VCG revenue back to the agents, are the ones we are interested in, provided such a redistribution keeps the mechanism *strategy-proof*, *non-deficit* and *individually rational*. Such redistribution mechanisms have previously been proposed, designed both by hand [12], [15], and by automated search [12], [16].

It is known that redistributing the VCG revenue naïvely (e.g. equal share of the revenue) can lead to the violation of the strategy-proof property [16]. This is because the second highest bidder gets incentivised to raise its own bid in order that the highest bidder pay more, thus receiving more from the redistribution. A general rule that can help prevent the violation of the strategy-proof property is to remove this incentive. It is also known that it is impossible to redistribute the whole VCG revenue amongst the agents, but only a large fraction of it [16]. Thus, a fraction of VCG revenue $m\theta_{m+1}$ gets redistributed amongst the agents, and the payment to agent i is not allowed to depend on θ_i , keeping its incentives unaffected by the redistribution, and the same as in the VCG auction.

1) Example redistribution mechanisms in the scenario

For $m = 1$, the redistribution mechanism proposed in [12], [17] first implements the VCG, and then redistributed the revenue in such a way that an agent i receives θ_2^{-i}/n , where θ_2^{-i} is the second highest reported value other than i 's (i.e. excluding θ_i). Thus, $\theta_2^{-i} = \theta_3$ for $i \in \{1, 2\}$ and $\theta_2^{-i} = \theta_2$ for $i \in \{3, \dots, n\}$. The total redistributed amount becomes $2\theta_3/n + (n-2)\theta_2/n$. This never exceeds the VCG revenue θ_2 [12], thus adhering to the non-deficit property. This mechanism was designed by hand with the objective of welfare under consideration. Let us call this mechanism RM_C . For $m = 1$, another redistribution mechanism, mentioned in [16], implements the VCG, and then redistributes the revenue in such a way that agent i receives $\theta_2^{-i}/(n-2) - 2\theta_3^{-i}/[(n-2)(n-3)]$, where θ_3^{-i} is the third highest reported value other than i 's.

2) *Linear redistribution mechanisms*:: An important point raised in [13], [16] suggests the potential for there to be better (with respect to an objective related to welfare) mechanisms that not only use the values θ_2^{-i} or both θ_2^{-i} and θ_3^{-i} , but also the values $\theta_j^{-i}, \forall j \in \{4, 5, \dots, n-1\}$, as part of the redistribution payment computations. In [13], this logic is applied to $m \geq 1$, and better mechanisms do result based on the single objective considered.

More precisely, a mechanism is sought where agent i would receive a payment $z_i = c_0 + c_1\theta_1^{-i} + c_1\theta_2^{-i} + \dots + c_{n-1}\theta_{n-1}^{-i}$. This is called a *linear redistribution mechanism* [13]. The transfer function for such redistribution mechanisms is $T_i(\theta) = v^{-i}(\theta^{-i}, f^*(\theta)) - v^{-i}(\theta^{-i}, f^*(\theta^{-i})) + z_i = -\theta_{m+1} + z_i$, while the allocation function remains the same as the VCG allocation function. This results in a mechanism space that can be characterised by a vector $(c_0, c_1, \dots, c_{n-1})$. We intend to explore this space of mechanisms with respect to multiple performance objectives, which includes a variant of the objective considered in [13], [16].

C. Mechanism Design as a MOOP

In [13], the linear redistribution mechanism search problem is transformed into a linear program with the introduction of various constraints. The objective used for this optimisation is the worst case welfare. By inspecting their optimised solution mechanisms, the authors were able to generalise the solutions into an analytical description spanning scenario settings, details of which can be found in [13]. Let us call this mechanism RM_{GC} . Considering VCG redistribution mechanisms, [15] proposes performance objective dealing with both welfare and fairness across the agents. A comparison of the RM_C mechanism with RM_{GC} reveals that neither of these dominates the other with respect to the objectives proposed in [15], which gives us further motivation to expose trade-off mechanisms in the linear redistribution mechanisms space. We also compare the exposed trade-off mechanisms with RM_C and RM_{GC} , which are the state-of-the-art in the linear redistribution mechanisms class.

The objectives from [15], which are computed with respect to the truthful outcome of the VCG mechanism, are as follows:

- **Welfare rate w :** the ratio of the social welfare of the agents including payments, to the social value of the efficient VCG allocation without payments. For a mechanism (f, T) and joint type $\theta \in \Theta^n$, the welfare rate $w((f, T), \theta)$ can be written as,

$$\frac{\sum_{i \in S} v_i(\theta_i, f_i(\theta)) + T_i(\theta)}{\sum_{i \in S} v_i(\theta_i, f_i^*(\theta))}$$

- **Envy rate e :** represents the average extent throughout the system of agents to which an agent prefers the outcome for another agent. Maximising over all agents j , if agent i were to experience a utility $u_i^{max} = \max_{j \in S} \{v_i(\theta_i, f_j(\theta)) + T_j(\theta)\}$ upon receiving j 's allocation and j 's payment, then, for a mechanism (f, T) and joint type $\theta \in \Theta^n$, the envy rate $e((f, T), \theta)$ can be written as,

$$\frac{1}{n} \sum_{i \in S} \frac{u_i^{max} - (v_i(\theta_i, f_i(\theta)) + T_i(\theta))}{u_i^{max}}$$

- **Disproportionality rate d :** As a dictator, an agent is seen as able to receive a utility composed of getting the *optimal allocation* for itself but with zero payment. $1/n$ of this utility is seen as a "fair share" [15]. The disproportionality rate represents the average extent throughout the system of agents to which an agent fails to obtain a "fair share" for itself. Thus, for agent i , if the ratio between its actual utility to its dictatorial utility is $1/n$, the agent gets a "fair share". For a mechanism (f, T) and joint type $\theta \in \Theta^n$, the disproportionality rate $d((f, T), \theta)$ can be written as,

$$\frac{1}{n} \sum_{i \in S} \max \left\{ 0, 1 - n \frac{v_i(\theta_i, f_i(\theta)) + T_i(\theta)}{\max_{o \in O} v_i(\theta_i, o_i)} \right\}$$

In the scenario we are considering, the optimal allocation $\max_{o \in O} v_i(\theta_i, o_i)$ amounts to the true value θ_i for agent i (i.e. dictatorial utility is θ_i).

It should be noted that the objective of maximising the worst case welfare in [13] is similar to the objective of maximising the welfare rate in [15] for the following reason. Both worst case welfare and the welfare rate are directly proportional to the aggregate redistribution payment. In both cases, the pressure for any search algorithm would be to make the aggregate redistribution payment move closer and beyond the VCG revenue. So, we can use the welfare rate objective as a proxy for the worst case welfare, although the latter is a much stronger search criterion. Nevertheless, we may omit the latter in our problem definition.

Constraints on the space of linear redistribution mechanisms of interest, which are characterised by $(c_0, c_1, \dots, c_{n-1})$, can be defined as follows. It is shown in [13], that for a mechanism to stay *ex post individually rational*, $z_n \geq 0$. This can also be written as $\sum_{i=m+1}^j c_i \geq 0$ for $j \in \{m+1, \dots, n-1\}$ [13]. Moreover, one can further limit the search space by the constraint suggesting that a mechanism is both *ex post individually rational* and *non-deficit* if $c_i = 0$ for $i \in \{0, 1, \dots, m\}$ [13] (please refer to Section IV for the specification of the value ranges of non-zero c_i). The *non-deficit* property can further be elaborated as a set of three constraints [13], namely: $(n - m - 1)c_{m+1} \leq m$, $(n \sum_{j=m+1}^{j=m+i-1} c_j) + (n - m - i)c_{m+i} \leq m$ for $i \in \{2, \dots, n - m - 1\}$, and $n \sum_{j=m+1}^{j=n-1} c_j \leq m$. The constraints make the welfare rate range between 0 and 1. The envy and disproportionality rates also range between 0 and 1. Moreover, in accordance with previous work [15], the envy rate is made to be 0 in the case where the numerator and denominator are 0.

Note that we are interested in designing mechanisms based on their *average case performance*. Since the joint type $\theta \in \Theta^n$ is not known beforehand, this uncertainty needs to be accounted for whilst searching for a Pareto optimal set of mechanisms. We assume θ_i s for the agents to be independently and identically distributed. We further assume the distributions to be uniform, and within the interval $[0, 1]$. This gives us a joint type that is distributed uniformly in the space $[0, 1]^n$. To account for the uncertainty in the joint type, we do the following. We take η_{mc} Monte Carlo (MC) samples from this joint type distribution, and evaluate the mechanism with respect to the specified multiple objectives η_{mc} times, each evaluation corresponding to a sample of the joint type. Following this, the average case performance of the mechanism is given by its welfare, envy and disproportionality rates, averaged across the η_{mc} evaluations. Let Θ_{mc} be the set of η_{mc} joint type samples. Thus, subject to the aforementioned constraints on the space of linear redistribution mechanisms, the search algorithm will need to optimise with respect to the following average case performance objectives:

- **(Maximise) Average case welfare rate g_w :**

$$g_w((f, T), \Theta_{mc}) = \frac{1}{\eta_{mc}} \sum_{\theta \in \Theta_{mc}} w((f, T), \theta)$$

- **(Minimise) Average case envy rate g_e :**

$$g_e((f, T), \Theta_{mc}) = \frac{1}{\eta_{mc}} \sum_{\theta \in \Theta_{mc}} e((f, T), \theta)$$

- **(Minimise)** Average case disproportionality rate g_d :

$$g_d((f, T), \Theta_{mc}) = \frac{1}{\eta_{mc}} \sum_{\theta \in \Theta_{mc}} d((f, T), \theta)$$

IV. APPLICATION OF MOEAS FOR DESIGNING MARKET MECHANISMS

The closest that previous work has come to applying MOEAs for designing market mechanism is by converting this multi-objective problem to that of a single objective [14]. We see that a truly multi-objective approach would seamlessly fit to this problem domain, especially since mechanisms often do not satisfy all objectives simultaneously. Exposing mechanisms which exist at different points within the trade-off thus helps enhance the understanding of the space of mechanisms considered, potentially resulting in insightful eventual application of discovered mechanisms in markets. We now describe how we apply an MOEA in the domain of market mechanism design.

A. Experimental Setup

We use NSGA-II [18] for searching the market mechanism design space, with the following modifications:

- 1) **Evaluation procedure:** Each candidate solution mechanism, given by the vector $(c_0, c_1, \dots, c_{n-1})$, whether in a parent or offspring population, is evaluated based on the average case welfare, envy and disproportionality rates defined in Section III-C. This requires MC sampling, which we described in Section III-C. We describe how we choose the MC sample size η_{mc} for such an evaluation in Section IV-A1.
- 2) **Selection and variation:** We use binary tournament selection with crowded-comparison to select one parent, which is then mutated as described in Section IV-A2. Such a selection is done as many times as the parent population size. We do not use crossover.

The evolved mechanisms are compared, given our three objectives, with state-of-the-art linear redistribution mechanisms, namely RM_{GC} and RM_C , as described in Section III-B.

1) Number of MC Samples for Average Case Performance Evaluation: To identify a suitable number of MC samples η_{mc} , we performed a preliminary experiment based on the idea of moving averages. More precisely, for a set of randomly generated solutions (i.e. mechanisms), we plotted the moving averages $g_k((f, T), \Theta_{mc})$ for $1 \leq \eta_{mc} \leq 6000$, for each of the objectives $k \in \{w, e, d\}$ as a function of the number of MC samples η_{mc} .

Fig. 1 shows the moving averages, together with the standard error, for a scenario setting with $n = 10$ agents using 10 randomly generated feasible solutions. The parameter η_{mc} can be set to the number of MC samples at which reliable convergence of the moving averages in all objectives becomes visually apparent. Such a convergence occurs around $\eta_{mc} = 2000$. Similar observations were made for the scenario settings given by $n = 5$ and $n = 8$, results not shown here.

As a comparison, previous work in this domain [15] considers $2000 \leq \eta_{mc} \leq 10000$ MC samples, depending on the scenario setting, in order to compute the average case




Fig. 1. Moving averages and standard error for $n = 10$ agents using 10 randomly generated linear redistribution mechanisms.

performance. Since more MC samples only make the performance estimate more accurate, we choose $\eta_{mc} = 6000$ for our experiments (during evolution), even though $\eta_{mc} = 2000$ would suffice. Moreover, in order to make fair comparisons (i) with previous work and (ii) among solutions exposed to different numbers of MC samples during evolution (which might occur due to the allowance of duplicate solutions), the performance of the mechanisms RM_{GC} and RM_C , as well as mechanisms represented by the final evolved population are evaluated on the *same* 10000 MC samples to form the new (averaged) objective values.

2) Managing Infeasible Solutions – Variation Operators and Numerical Limits on Solution Vectors: A solution mechanism is constrained by the inequalities which make it ex-post individually rational and non-deficit, as mentioned in Section III-C. These constraints define the feasible solution space. MOEAs use variation operators for the exploration of this space, and one can employ techniques commonly used for dealing with infeasible solutions, e.g. penalisation, clamping to the boundary, if a solution goes out of the feasible region. Penalisation, given constraint inequality violations, is possible, but we believe this can result in a loss in the search gradient if the variation operator generates too many infeasible solutions.

In the class of mechanisms considered, the constraints depend on n and m , thus different scenario settings impose different numerical limits on feasible solutions. As a first step, for the work reported in this paper, we take a simple approach to managing infeasible solutions. We explicitly set *approximate and pessimistic lower and upper bounds* on the numerical limit for each variable of the solution vector as follows. Considering the same 2000 MC samples, we compute the *average (across 2000 randomly generated feasible solutions) number of random draws from a joint uniform distribution $[-\lambda, \lambda]^{n-m-1}$ that are required to generate a feasible solution from within the bounds $[-\lambda, \lambda]$ on solution variables, given the scenario. The bounds $[-\lambda, \lambda]$ for a scenario setting n are chosen with one decimal precision to be those that result in this computed average \leq average number of random draws from*

the joint uniform distribution $[-1.0, 1.0]^3$ required to generate a feasible solution for $n = 5$ and $m = 1$ ⁵. The reasoning behind this is primarily to keep the time spent by the MOEA in coming up with a new population at each generation roughly the same across different scenario settings. The reasoning also results from the type of mutation operator considered. This *mutation operator* is as follows: with probability $1/(n-m-1)$ (the effective length of a solution vector), we replace a solution variable by a value drawn randomly from a uniform distribution within the interval $[-\lambda, \lambda]$.

Note that these choices indeed result in the MOEA not being allowed to search the entire feasible region, thus the resulting solutions are only approximately Pareto optimal within the restricted search space. We thus can expect to see the trade-off imposed by the considered multiple objectives on the found solutions, but within this smaller feasible space. The goals of the paper being the exposition of a new problem domain, and a simple application of an MOEA on the problem helping reveal insights for the domain in terms of trade-offs, we did not find the choices to be too restrictive.

3) Settings of Various Parameters for the Experiments:

The scenario settings in which we use NSGA-II for exposing trade-off mechanisms are specified in Table I, along with the used NSGA-II parameter values and the bounds on the search space that we impose, as discussed in Section IV-A2.

TABLE I. SCENARIO SETTINGS, PARAMETER VALUES FOR NSGA-II, AND BOUNDS ON THE NON-ZERO VARIABLES OF THE SOLUTION VECTOR FOR EACH SCENARIO SETTING.

Scenario Settings		
n	5, 8, and 10	
m	1	
NSGA-II Parameters		
generations	200	
population size $ P $	50	
population size $ Q $	50	
η_{mc}	6000	
mutation rate	$1/(n-m-1)$	
Bounds on the non-zero solution variables for each n		
5	$[-1.0, 1.0]$	
8	$[-0.2, 0.2]$	
10	$[-0.1, 0.1]$	

V. RESULTS AND DISCUSSION

We now present the results that the application of NSGA-II reveals.

A. Revealing New Mechanisms and Hidden Relationships between Objectives

Fig. 2 shows trade-off plots for NSGA-II outcomes for the three scenario settings $n \in \{5, 8, 10\}$. We show the trade-off imposed by the considered objectives by plotting graphs for the three pairs of combinations given these objectives, i.e. welfare vs. envy, welfare vs. disproportionality, and envy vs. disproportionality⁶.

⁵Note that the constraint inequalities in Section III-C impose restrictions on (c_0, \dots, c_m) to be zero, leaving us with a joint distribution over $n-m-1$ variables.

⁶This is for ease of visualising the trade-offs, which would be harder to ascertain in three dimensional plots.

Previous work that has looked at designing mechanisms by hand (RM_C) [12], [15] in the linear redistribution mechanisms class came up with a mechanism with both high welfare and relatively fair (in terms of both envy and disproportionality) mechanism. We can see this in Fig. 2, where the RM_C mechanisms can be seen to have low $-g_w$ (high welfare) and low g_e (envy) and g_d (disproportionality) across scenario settings. Mechanism RM_{GC} on the other hand, which was designed for high worst case welfare [13], [16], happens to do well on disproportionality across scenario settings, but results in high envy in the system. Both these works, which are state-of-the-art when it comes to designing mechanism in the linear redistribution class, did not however tell us much about the trade-off that the three objectives imposed on the designed mechanism. Thus, the relationship between welfare and fairness in the linear redistribution mechanisms class remained unknown. Applying NSGA-II on the problem indeed reveals this relationships.

More specifically, we can get the following main insights from the NSGA-II outcomes that can inform practitioners who want to design mechanisms and employ them in their respective markets:

- 1) We find mechanisms with lower envy w.r.t. the state-of-the-art in linear redistribution mechanisms, specifically when it comes to $n = 5$, and $n = 8$, as evident from Figs. 2(a) and 2(b). Such low envy based fairness comes at the expense of welfare of course. For $n = 10$, this does not seem to be the case (Fig. 2(c)), but remember we apply NSGA-II on a search space smaller than the feasible region (cf. Section IV-A2). This affects the outcomes we get. Searching in a larger feasible region, or indeed the entire feasible region remains future work. In any case, a practitioner could, in effect, decide to employ a mechanism with lower envy, at least for scenario settings with fewer number of agents, and yet keep the system with high enough welfare.
- 2) Welfare and disproportionality are negatively correlated with each other (Figs. 2(d), 2(e), and 2(f)). In other words, welfare and fairness (given by the disproportionality objective) do not conflict with each other. Yet, envy (which also measures fairness), does conflict with welfare. This difference in the relationship between welfare and disproportionality, and welfare and envy, has previously not been exposed for linear redistribution mechanisms. We also see from Figs. 2(g), 2(h), and 2(i), that envy conflicts with disproportionality. The relationship between these fairness objectives was hitherto unknown, and was possible to be revealed via the multi-objective optimisation approach. A practitioner would benefit from understanding such conflicting relationships between different fairness measures they consider mechanisms to satisfy, and indeed how one could be traded for the other. This can help them employ mechanisms with the desired kind and level of fairness in their respective markets.

Note that the multi-objective approach as applied in this paper is meant to only help reveal trade-offs. Searching for mechanisms in the entire feasible region could potential reveal

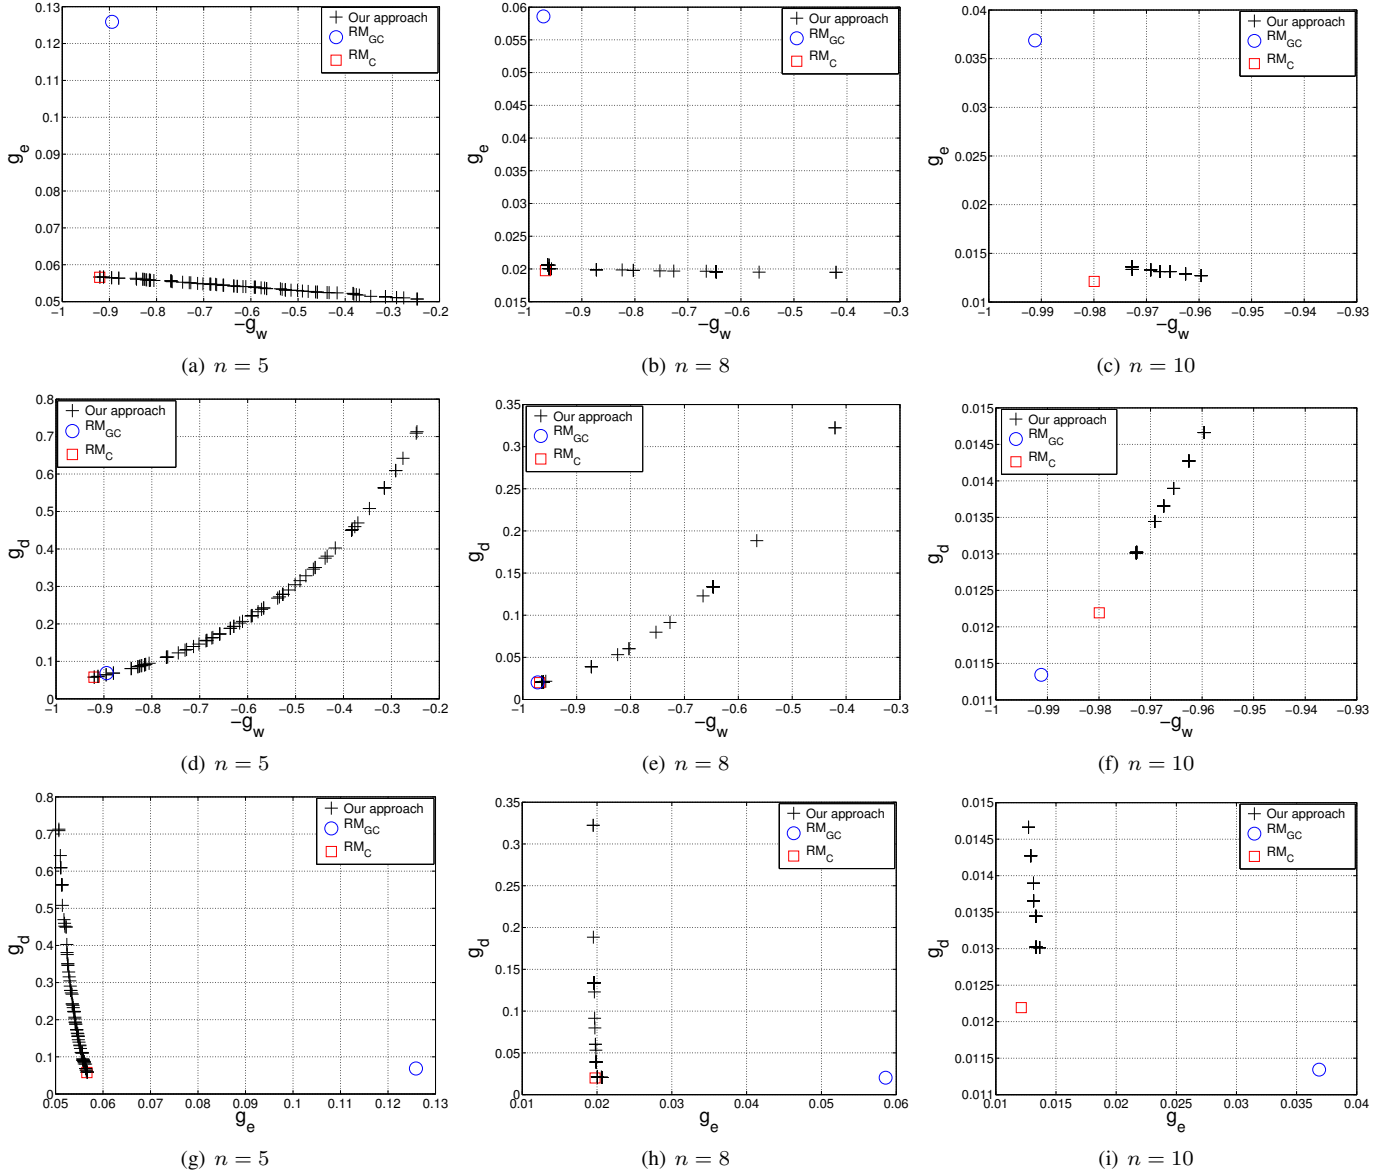


Fig. 2. Trends in the trade-off imposed by the average case performance objectives, specifically ((a), (b), (c)) average case welfare vs. envy, ((d), (e), (f)) average case welfare vs. disproportionality, and ((g), (h), (i)) average case envy vs. disproportionality, on evolved linear redistribution mechanisms across scenario settings given by $n \in \{5, 8, 10\}$ ($m = 1$ for all). Plotted are the average case performances of the evolved as well as state-of-the-art linear redistribution mechanisms (RM_{GC} and RM_C), based on the same 10000 MC samples. For the evolved mechanisms, each plot shows the non-dominated set of mechanisms generated from the collated final populations of 10 independent runs of NSGA-II.

Pareto-superior mechanisms as compared to state-of-the-art, but this remains future work.

B. Potential Impact of Exposing Trade-off Market Mechanisms

It is not clear why the potential of MOEAs in exposing trade-off mechanisms has remained untapped to date. For linear redistribution mechanisms, previous work to which we compare our work has designed mechanism based on a single objective, but analysed the performance of these mechanisms on other relevant objectives. From the trade-offs exhibited in Fig. 2, and the revelations in Section V-A, it seems reasonable to exploit the multi-objective philosophy and carry out such an analysis at design time via MOEAs.

The virtue of designing mechanisms by hand is that an

analytical solution can be more general in that it usually is a function of the scenario setting, e.g. RM_C [12], [15], thus one function spans scenario settings. On the other hand, optimisation using a single objective or as in our case using multiple objectives, is done specific to a scenario setting. Even so, as shown in [13], [16], it is possible to inspect and generalise the optimisation outcomes into an analytical description spanning scenario settings, e.g. RM_{GC} . Even if a single analytical description for a mechanism can be constructed by inspecting optimisation outcomes, one would miss out on possible trade-off analytical descriptions which could potentially be constructed upon inspecting the multi-objective optimisation outcomes. The potential of the multi-objective approach in this problem domain should thus not be undermined.

Moreover, in multi-agent systems research, specifically tailored towards resource allocation problems in decentralised systems, employing market mechanisms for such allocations is becoming commonplace [1], [19], [20]. Designing mechanisms for these systems with an a priori specification of preferences over the objectives may not always be feasible. This is because the requirements on the nature of allocation of resources in these systems may change in time, sometimes requiring efficiency/welfare and at other times requiring a fair distribution (proportionality or low envy). Such changes in requirements would necessitate different mechanisms to be employed. Knowing the trade-off imposed by the objectives on the discovered mechanisms can thus be useful for making such choices.

VI. CONCLUSIONS AND FUTURE WORK

Mechanism design has traditionally been carried out by hand, and recently via automation, albeit as a single objective optimisation problem or as an optimisation problem with the aggregation of multiple objectives into a single objective. Considering a class of mechanisms known as linear redistribution mechanisms, and the objectives of welfare and fairness (which is represented by the objectives envy and disproportionality), we pose the automated mechanism design problem as a naturally existing multi-objective (in our case three-objective) optimisation problem. We apply a modified NSGA-II to this problem, and show how the trade-off imposed by the objectives, which has previously never been exposed for linear redistribution mechanisms, provides insights into the problem domain. This revealed trade-off can be seen to be useful in at least the following ways: it can (i) provide practitioners with mechanisms exhibiting lower envy w.r.t. state-of-the-art linear redistribution mechanisms, and (ii) can reveal previously unknown relationships between the objectives considered. The application of the MOEA is rather simple, even though novel, and leaves many open directions for future research. As such, we will look at a more elaborate way of applying MOEAs which allows for searching the entire feasible region of the mechanism space considered. Further, scenario settings with $1 < m < n$ will also be looked at. The assumption of agent types distributed uniformly in the space $[0, 1]^n$ will also be relaxed. We will also consider different mechanism spaces (e.g. allowing non-linear redistribution of revenue) and objectives in order to gain insight into those spaces and potential trade-offs that the objectives impose on those spaces.

ACKNOWLEDGMENT

The research leading to these results was conducted in the EPiCS project (Engineering Proprioception in Computing Systems) and has received funding from the European Union Seventh Framework Programme under grant agreement n° 257906. <http://www.epics-project.eu/>

REFERENCES

- [1] L. Esterle, P. R. Lewis, X. Yao, and B. Rinner, "Socio-economic vision graph generation and handover in distributed smart camera networks," *ACM Transactions on Sensor Networks*, vol. 10, no. 2, 2014, to appear.
- [2] R. K. Dash, N. R. Jennings, and D. C. Parkes, "Computational-mechanism design: A call to arms," *IEEE Intelligent Systems*, vol. 18, no. 6, pp. 40–47, November 2003.
- [3] R. Cavallo, "Handling self-interest in groups, with minimal cost," in *Proc. 21st National Conference on Artificial Intelligence (AAAI'06)*. Boston, MA: AAAI Press, 2006, pp. 1585–1588.
- [4] F. Faniyi and R. Bahsoon, "Self-managing SLA compliance in cloud architectures: A market-based approach," in *Proc. International ACM Sigsoft Symposium on Architecting Critical Systems (ISARCS'12)*. Bertinoro, Italy: ACM, 2012, pp. 61–70.
- [5] S. Phelps, P. McBurney, and S. Parsons, "Evolutionary mechanism design: a review," *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 2, pp. 237–264, September 2010.
- [6] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.
- [7] D. P. Friedman and J. Rust, *The Double Auction Market: Institutions, Theories, and Evidence*. Boulder, Colorado, USA: Westview Press, 1993.
- [8] J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield, "JCAT: a platform for the TAC market design competition," in *Proc. of the 7th International Conference on Autonomous Agents and Multi-agent Systems: Demo papers (AAMAS'08)*. Estoril, Portugal: International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 1649–1650.
- [9] R. Wang, "Auctions versus posted-price selling," *The American Economic Review*, vol. 83, no. 4, pp. 838–851, September 1993. [Online]. Available: <http://www.jstor.org/stable/2117581>
- [10] D. Cliff, "Explorations in evolutionary design of online auction market mechanisms," *Journal of Electronic Commerce Research and Applications*, vol. 2, no. 2, pp. 162–175, 2003.
- [11] V. Walia, A. Bye, and D. Cliff, "Evolving market design in zero-intelligence trader markets," in *Proc. IEEE International Conference on E-Commerce*. IEEE Computer Society, 2003, pp. 157–163.
- [12] R. Cavallo, "Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments," in *Proc. 5th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS'06)*. Hakodate, Japan: ACM, 2006, pp. 882–889. [Online]. Available: <http://doi.acm.org/10.1145/1160633.1160790>
- [13] M. Guo and V. Conitzer, "Worst-case optimal redistribution of VCG payments in multi-unit auctions," *Games and Economic Behavior*, vol. 67, no. 1, pp. 69–98, September 2009. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0899825608001267>
- [14] S. Phelps, S. Parsons, E. Sklar, and P. McBurney, "Applying multi-objective evolutionary computing to auction mechanism design," Department of Computer Science, University of Liverpool, Liverpool, UK, Tech. Rep. ULCS-02-031, 2002.
- [15] R. Cavallo, "Fairness and welfare through redistribution when utility is transferable," in *Proc. 26th AAAI Conference on Artificial Intelligence (AAAI'12)*, Toronto, Canada, 2012, pp. 1306–1312. [Online]. Available: <https://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5115/5462>
- [16] M. Guo and V. Conitzer, "Computationally feasible automated mechanism design: General approach and case studies," in *Proc. 24th AAAI Conference on Artificial Intelligence (AAAI'10)*, Atlanta, GA, 2010, pp. 1676–1679. [Online]. Available: <http://www.aaai.org/ocs/index.php/AAAI/AAAI10/paper/view/1868>
- [17] M. J. Bailey, "The demand revealing process: To distribute the surplus," *Public Choice*, vol. 91, no. 2, pp. 107–126, 1997. [Online]. Available: <http://www.jstor.org/stable/30024208>
- [18] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.
- [19] A. Chandra, K. Nymoen, A. Voldsund, A. R. Jensenius, K. Glette, and J. Torresen, "Enabling participants to play rhythmic solos within a group via auctions," in *Proc. International Symposium on Computer Music Modeling and Retrieval (CMMR'12)*, 2012, pp. 674–789. [Online]. Available: http://cmmr2012.eecs.qmul.ac.uk/sites/cmmr2012.eecs.qmul.ac.uk/files/pdf/papers/cmmr2012_submission_64.pdf
- [20] S. H. Clearwater, Ed., *Market-based Control: A Paradigm for Distributed Resource Allocation*. Singapore: World Scientific, 1996.