The Effects of Variation on Solving a Combinatorial Optimization Problem in Collaborative Multi-Agent Systems

Christian Hinrichs and Michael Sonnenschein

Department of Computing Science, University of Oldenburg, Germany,
{christian.hinrichs,michael.sonnenschein}@uni-oldenburg.de

Abstract. In collaborative multi-agent systems, the participating agents have to join forces in order to solve a common goal. The necessary coordination is often realized by message exchange. While this might work perfectly in simulated environments, the implementation of such systems in a field application usually reveals some challenging properties: arbitrary communication networks, message delays due to specific communication technologies, or differing processing speeds of the agents. In this contribution we interpret these properties as sources of variation, and analyze four different multi-agent heuristics with respect to these aspects. In this regard, we distinguish synchronous from asynchronous approaches, and draw conclusions for either type. Our work is motivated by the use case of scheduling distributed energy resources within self-organized virtual power plants.

Keywords: Heuristics · Distributed Problem Solving · Self-Organization · Synchronous vs. Asynchronous Approaches · Virtual Power Plants

1 Introduction

According to Ferber [1], interaction between agents is a main component of multi-agent systems (MAS). A special case is *collaboration* situations, in which the agents in a MAS pursue a common goal, and have to join forces in a coordinated manner in order to reach that goal. This paradigm does not exclude self-interested or untrustworthy agents per se, as the effects induced by those properties can be tackled by proper coalition formation and incentivization mechanisms, cf. [2–4].

For example, such collaboration situations may arise in energy systems with a significant share of distributed energy resources (DER) like small scale combined heat and power (CHP) plants. As of today, in many European countries, especially Germany, DER operate under financial security of guaranteed electrical feed-in tariffs. However, in order to follow the goals as defined by the European Commission, this subsidy dependence should be reduced. The formation of virtual power plants (VPPs) as an aggregation level for small scale DER,

Published in: J.P. Müller, M. Weyrich, A.L.C. Bazzan, Multiagent System Technologies, Lecture Notes in Computer Science, Volume 8732, © Springer 2014

to overcome market barriers and to increase mutual reliability within a VPP by redundant dimensioning and adaptive compensation techniques [5], forms a possible integration path for these DER. In [6], a concept of self-organized VPPs is proposed, which involves both coalition formation and scheduling tasks as distributed optimization problems. The coalition formation process allows self-interested agents to join forces towards the common goal of offering reliable energy products at the market. This includes admissibility checks regarding the underlying power grid as well as monetary incentivization for the agents, as described in e. g. [7]. Each of those VPPs then represents a collaborative MAS as described above, as the participating DER have to coordinate their actions (e. g. their generation of electrical power) cooperatively in order to reliably deliver energy products as an aggregate.

In this paper, we focus on the effects of variation in collaborative MAS. Based on the findings of Ashby [8], who identified variation and error as important properties for the control of complex technical systems, Campbell et al. [9] studied variation in the context of MAS solving a distributed task allocation problem. Subsequently, Anders et al. [10] extended this work by introducing uncertainty from the environment of the MAS, and demonstrated the effects of variation on multi-agent algorithms solving the frequency stabilization problem in the power grid (cf. [11]). There, variation was modeled as a randomized threshold parameter for the activation of agents, thus affecting the participation of agents in the optimization process stochastically. The objective of the contribution at hand is to continue this research by addressing further types of variation:

- environmental effects, e. g. communication delays or arbitrary communication topologies and
- technical aspects, e. g. differing processing speeds of the participating agents.

Our study is motivated by the scheduling of DER in self-organized VPPs as described above. Because this task targets a critical infrastructure, robustness against variation is crucial here. Therefore, we examine different approaches for solving this scheduling problem with regard to the variation sources above, which are likely to be faced in deployed field applications.

The contribution proceeds as follows: In Sect. 2, we give a formal description of the optimization problem and subsequently present different solution strategies for this task in Sect. 3. Following, Sect. 4 describes the considered types of variation in detail. Sect. 5 then describes our evaluation setup and discusses the results. There we show that some approaches suffer significantly from variation, whereas others are basically unaffected or even benefit from certain types of variation. Finally, Sect. 6 concludes the paper.

2 The Multiple-Choice Combinatorial Optimization Problem

The motivation for this paper is the task of scheduling DER within self-organized VPPs. More specifically, we are referring to the use case of active power products traded on a day-ahead power market like the European Power Exchange (EPEX SPOT), cf. [12, 6]: Given an active power profile over a future planning horizon (e. g. the next 24 hours, discretized into 15 minute intervals) as scheduling target (in the following denoted as *active power product*), the task is to select a schedule for each participating DER for the planning horizon, such that the aggregation of all selected schedules within the VPP yields the active power product as close as possible. From a centralized point of view, this optimization problem can be expressed as *Multiple-Choice Combinatorial Optimization Problem* (MC-COP), an integer programming model that was already introduced in a similar form in [13]:

$$\min \quad \left\| \zeta - \sum_{i=1}^{|\mathsf{A}|} \sum_{j=1}^{|\mathsf{S}_i|} (\theta_{ij} \cdot x_{ij}) \right\|_1 \tag{1a}$$

$$\text{subject to} \quad \sum_{j=1}^{|\mathsf{S}_i|} x_{ij} = 1\,, \quad x_{ij} \in \{0,1\}\,, \quad i = 1 \ldots |\mathsf{A}|\,. \tag{1b}$$

Here, A denotes the set of agents in the considered MAS, i. e. the set of DER in a self-organized VPP. Each agent $a_i \in \mathsf{A}$ has an associated set of schedules $\mathsf{S}_i = \{\theta_{i1}, \theta_{i2}, \dots\}$ for the considered planning horizon. The task, as depicted in (1a), is to find a selection of schedules for the agents, such that the distance of the aggregation of the selected schedules to the active power product ζ, cumulated over all planning intervals, is minimized. The constraints in (1b) make sure that for each agent exactly one schedule is selected.

3 Solution Strategies for MC-COP

The MC-COP in (1) refers to a global view on the MAS. A central optimizer, with full knowledge about every S_i, could find an optimal solution using standard solving techniques for integer programs. However, in the considered use case, each S_i is represented by a self-interested agent acting on its own behalf. While a cooperative attitude for the agents is incentivized throughout the coalition formation process of the self-organized VPP (cf. [7]), the transfer of all S_i to a central instance should be avoided due to privacy aspects as well as technical difficulties, as discussed in e. g. [13–15]. Moreover, due to the nonseparability of the considered optimization problem with respect to the occuring constraints between schedule selections, approaches from the domain of Distributed Constraint Optimization Problems (DCOP) cannot effectively be applied here either, cf. [16]. Therefore we present a number of feasible approaches for MAS (i. e. each DER is represented by an agent) in the following, classified by their underlying coordination mechanism.

One possible solution strategy for the MC-COP is realizing a virtual market. In such a coordination mechanism, agents place bids on fulfilling (parts of) the power product ζ in a virtual marketplace. A central auctioneer then performs a market matching by selecting and combining appropriate bids. This process can be repeated iteratively, in order to approximate an optimal solution. Examples for this strategy can be found in [17–19]. Common to such market-based approaches is a tree topology with a central auctioneer as root node. The agents generally keep their search spaces private, and publish only selective parts of it in the form of bid proposals.

Another strategy that relies on a tree topology is the *Energy Plan Overlay Service* (EPOS), as proposed in [20]. Similar to the market setting, the root node acts as a global controller of the system and announces the aspired power product ζ to all agents. After that, at first the leaf agents send their whole search spaces to their respective superordinate agents in the topology. Subsequently, each intermediate agent executes the following actions: Upon receiving the search spaces of all its subordinate agents (and possibly also schedule selections from lower levels, see below), the intermediate agent calculates the best schedule combination from these search spaces with respect to ζ. Then, on the one hand, the selected schedules are sent back to the corresponding subordinate agents, thus informing them about their obligations. On the other hand, the intermediate agent sends the calculated schedule combination together with its own search space to its superordinate agent, which then executes the very same actions. In summary, in this approach agents select schedules for their subordinates within a tree topology, thus realizing a bottom-um planning with a certain degree of parallelism. Similar to market-based approaches, this planning can iteratively be repeated, in order to approximate an optimal solution.

An alternative strategy is the *Stigspace* approach [21]: All agents have access to a central information repository called the *Stigspace*, hence this is an instance of the black board coordination mechanism. In principle, the agents perform an iterative improvement process by adapting their schedule selection according to updated information in the *Stigspace*. After such an adaptation, an agent is obliged to publish its schedule selection by placing it into the *Stigspace* again, thus triggering adaptation in other agents regarding this choice. The process terminates either if no agent can improve the current situation any more with respect to fulfilling ζ, or if an external termination criterion holds (e. g. a predefined timespan, as used by the authors in [21]).

Finally, a completely distributed and asynchronous approach is given with the *Combinatorial Optimization Heuristic for Distributed Agents* (COHDA), see [13]. The key concept of COHDA is an asynchronous iterative approximate best-response behavior, where each agent reacts to updated information from other agents by adapting its own selected schedule with respect to the power product ζ. The agents are placed in an artificial communication topology (e. g. a *small world* topology), such that each agent is connected to a non-empty subset of other agents. To compensate for the resulting non-global view on the system, each agent a_i collects two distinct sets of information: on the one hand the believed current configuration $\gamma_i = \{\theta_1, \dots, \theta_{|\mathcal{A}|}\}$ of the system (that is, the believed set of currently selected schedules of all agents), and on the other hand the best known combination γ_i^* of schedules with respect to the power product ζ it has encountered so far. All agents $a_i \in \mathcal{A}$ initially only know their own respective set of schedules $\mathcal{S}_i$, and the difficulty of the problem is given by the distributed nature of the system in contrast to the task of finding a common allocation of schedules. Thus, the agents coordinate via message exchange. Beginning with an arbitrarily chosen representative of the self-organized VPP, each agent a_i executes the following three steps, cf. [13]:

1. (`update`) When an agent a_i receives information from one of its neighbors (say, a_j), it imports this information (γ_j and γ_j^*) into its own knowledge base by updating γ_i and, if better, replacing γ_i^* with γ_j^*.
2. (`choose`) The agent now adapts its own schedule according to the newly received information. If it is not able to improve the believed current system configuration γ_i, the agent reverts its current schedule to the one stored in γ_i^* (note that γ_i^* contains a schedule for each agent in the system and a_i takes its own of course).
3. (`publish`) If γ_i or γ_i^* has been modified in one of the previous steps, the agent finally publishes its knowledge base (γ_i, including its own selected schedule, and γ_i^*) to its neighbors.

The heuristic terminates when for all agents γ and γ^* are identical. At this point, γ^* is the final solution of the heuristic and contains exactly one schedule for each agent.

4 Sources of Variation

For the evaluation of the influence of variation on collaborative MAS solving the MC-COP, we considered different types of variation. We specifically focused on the effects that occur when finally implementing such a MAS in the targeted field application:

Communication topologies. We consider this as a source of variation, because it is not known in which topology such a system would operate in the field. For example, in our use case of self-organized VPPs, the participating DER might be connected to a given restricted communication network like power line carrier (PLC) [22] or a wireless mesh network [23].

Message delays. Another important aspect arising from the underlying communication technology is message delays. While it is possible to implement communication networks with real-time properties, a more likely scenario would be to reuse already available technologies like PLC or mesh networks as described above, or alternatively utilizing general purpose commmunication networks like broadband internet connections. In these cases it is important to know how an algorithm that heavily depends on communication behaves.

Reaction delays. Besides the communication technology, the agents themselves are a source of variation. In our use case of scheduling DER, agents would be implemented on different hardware platforms, depending on the manufacturer and model of the appliance under control. This would lead to different processing speeds of the agents. Combined with a dynamically changing environment for the executed agent with respect to available information (e. g. dynamically updated knowledge in the EPOS, *Stigspace* and COHDA approaches), this results in dynamically changing delays while reacting to incoming messages. While these differences may be quite small, their effects are still interesting with regard to the performance of collaborative MAS.

Of course, there are more possible sources of variation one might want to consider. For example, regarding robustness, message losses and (temporary) node failures are such types. However, due to the *FLP impossibility proof* by Fischer et al. [24], we defer the task of handling these to the control layer of the communication protocol (e. g. as in [23]), and do not consider them in our study. Similarly, examining the adaptivity of an approach with respect to e. g. dynamically changing search spaces or optimization goals is an interesting aspect. For the former, we refer to the ongoing work in [6] and [25]. For the latter, a respective study is presented in [26, in press].

5 Evaluation

The objective of the paper at hand is to evaluate the effects of variation on collaborative MAS solving the MC-COP. For this, we will focus on the sources of variation introduced in Sect. 4 one by one, each with respect to the different solution strategies presented in Sect. 3. We do not study second order effects in this paper, hence the parameters are analyzed independently from each other. In general, we consider three different effect types: solution quality, run-time and communication expenses.

In cases where the arising effects are straightforward and easy to derive, we argue verbally about them. For the remaining cases, we present results from respective simulation experiments. These have been conducted using a system that is capable of simulating an asynchronous MAS with configurable parameters matching the considered sources of variation. But instead of restricting our study to the motivating use case, e. g. by simulating different types of DER, we use synthetic problem instances. This has the advantage that the inherent properties of the problem instances are known beforehand, so that simulation results can be interpreted independently from specific use cases. More specifically, we tailored the problem instance $\mathsf{P}(m, n, q)_{h/S}$ from [27] for our problem. Originally, the problem instances in [27] were designed for the multiple-choice knapsack problem (MC-KP). But as the MC-COP defined in (1) is closely related to the MC-KP (more specifically, the MC-COP is a generalized multiple-choice subset sum problem, MC-SSP, which in turn corresponds to the MC-KP without profits), we can reuse the construction method easily by neglecting the profit values that are present in the MC-KP. To preserve consistency with the referred work, we use the very same parameter values for m, n and q for constructing our problem instances as in [27, Sect. 5]. Following, the considered instances then comprise $m = 10$ agents with each $n = 5$ available elements to choose from. Referring to the symbols defined in Sect. 2, this means that $\forall a_i : \mathcal{S}_i = \{\theta_{i1}, \ldots, \theta_{i5}\}$. The parameter $q = 5$ defines the dimensionality of the elements, i. e. in our use case, the schedules of each agent would span a planning horizon of 5 intervals. Finally, the parameter h determines the position of the optimization target ζ in the solution space of the optimization problem, which is divided into S partitions. In our study, we set $S = 100$, and treat h as random variable that is uniformly distributed over $[1, S]$, such that the actual optimization goal varies with each executed simulation run. In summary, using fixed values for m, n and q, while applying varying values for h, yields problem instances that are directly comparable due to their identical basic structure, but nonetheless allow us to derive statistically sound conclusions. For more details on the construction of the problem instances, especially on how to define the concrete values of ζ and the elements θ_{ij}, please refer to [27, Sect. 4].

The considered sources of variation are each modeled as a stochastic process (which will be described in more detail in the respective sections below), thus we repeated each experiment for 100 times. An *experiment* is thereby defined as a specific parameter setting for the considered source of variation, i. e. the amount of variation that we impose on the system. For each variation type, the examined range of this amount is chosen such that the resulting effects could be clearly identified, respectively. Within each experiment, for each executed simulation, we recorded the number of simulation steps until the heuristic terminated and the total number of exchanged messages during the whole process. Additionally, the quality of the final solution is calculated according to (1a) for each simulation. This calculated value is then normalized regarding the theoretically best and worst solution possible (which in turn have been calculated using an exhaustive search method in advance). Thus, in the following, solution quality is expressed as remaining error in the range $[0, 1]$, i. e. as normalized distance between the final solution and the optimization target, so that lower values denote better solutions. Please note that a preliminary study using similar performance indicators has been published in [28]. However, that work focused on the COHDA heuristic only and is based on a simulation scenario specific for the use case of scheduling DER. Moreover, the interpretation of the results has been done there on a qualitative level only. In order to gain resilient knowledge about the effects of variation, we re-enacted the study using synthetic problem instances as described above. Further, we performed a regression analysis on the resulting data series to examine the effects quantitatively.

For each simulation experiment, we show a figure containing the results for the three performance indicators solution quality, run-time and communication expenses as boxplots, where the respective box spans from the upper to the lower quartile of the results. The median is shown as horizontal line within a box, whereas the whiskers span over $1.5 \times$ the interquartile range. Additionally, the average is denoted with a star marker and outliers are illustrated by plus markers.

Table 1. Transformations used in the linear regression tests

Method	Transformations	Regression equation	Predicted value
Logarithmic model	lin-log	$y = \beta_0 + \beta_1 \cdot \ln x$	$\hat{y} = \beta_0 + \beta_1 \cdot \ln x$
Power model	log-log	$\ln y = \beta_0 + \beta_1 \cdot \ln x$	$\hat{y} = \mathrm{e}^{\beta_0} \cdot x^{\beta_1}$
Linear	lin-lin	$y = \beta_0 + \beta_1 \cdot x$	$\hat{y} = \beta_0 + \beta_1 \cdot x$
Quadratic model	sqrt-lin	$\sqrt{y} = \beta_0 + \beta_1 \cdot x$	$\hat{y} = (\beta_0 + \beta_1 \cdot x)^2$
Exponential model	log-lin	$\ln y = \beta_0 + \beta_1 \cdot x$	$\hat{y} = \mathrm{e}^{\beta_0 + \beta_1 \cdot x}$

In the subsequent regression analysis, the goal was to identify the exact intercorrelation between the respective source of variation and each of the performance indicators. For that, we took the medians of each recorded data series, and applied different variable transformations to the resulting series of medians for each experiment, which are summarized in Tab. 1. For each transformed data series, we then performed a standard linear regression and calculated the according coefficients of determination $\mathrm{R}^2 \in [0, 1]$ based on Pearson's correlation coefficient R. For a given data series, the transformation method that yields the highest R^2 for this data then describes the estimated intercorrelation model.

5.1 Communication Topologies

Depending on the considered solution strategy, varying communication topologies can have a more or less severe impact. For example, market based approaches as well as EPOS and the *Stigspace* approach all require a very specific communication topology to work properly (i. e. tree topologies in the former two cases and a star topology in the latter one). Hence, to be able to cope with arbitrary topologies in the field, techniques like overlay networks [29] have to be implemented in order to overcome the inherent restrictions of these approaches. Due to the thereby induced routing overhead, arbitrary communication topologies will result in possibly longer transmission times, but will have no *direct* effect on solution quality, run-time and communication expenses. Thus we refer to the examination of message delays in Sect. 5.2 for these approaches.

More interesting in this regard is the COHDA approach, as this heuristic is inherently able to cope with arbitrary topologies, as long as the topology forms a connected graph. A complete graph naturally yields the fastest spreading of information in the network. But as COHDA is an asynchronous heuristic where the agents become active upon receiving updated information from their neighborhood, and subsequently send messages back into their neighborhood, such a topology also leads to a maximal number of exchanged messages. On the other hand, if all agents are connected e. g. in a ring, the system will show the opposite behavior, as each agent is able to send messages to exactly two other agents. Hence, in such a topology, information spreads more slowly while exhibiting fewer messages.

In order to gain more detailed insight into the resulting effects, we studied this in terms of the *density* of the communication topology, i. e. different sizes of neighborhoods for the agents. For this, based on the ring topology and the fully connected graph as extreme cases, we simulated different increments of link densities ϕ. This parameter is based on the definition of small world networks in [30]: Starting with a ring topology, the density of the network is increased by adding up to $n \cdot \phi$ links to the topology. This is done by randomly choosing two agents in each of the $n \cdot \phi$ iterations of this construction process, and connecting these agents if they aren't already connected in the communication topology.

The results for a series of simulations with $\phi \in \{0, 0.1, 0.5, 1, 2, 4\}$ are summarized in Fig. 1. All results in this experiment generally show the same solution quality. Hence, the quality of the solutions produced by COHDA is almost independent from the density of the communication topology. Regarding the simulation length and the amount of communication, an opposing trend between those two is visible. Obviously, if more communication links are present, the simulation terminates faster while exhibiting a larger amount of messages, and vice versa. The regression analysis for these two effects is shown in Fig. 2. In either case, both the logarithmic and the power model yield very high coefficients of determination. A closer look at the estimated model parameters reveals that, in the power model, the elasticity coefficient in both cases is quite low ($\beta_1 = 0.14$ for the simulation steps data and $\beta_1 = 0.26$ for the messages data, not shown in the figure), yielding a very similar intercorrelation in comparison to the logarithmic model (cf. Tab. 1). Therefore, the sensitivity of the simulation steps and the messages to the communication topology, respectively, decreases with increasing link density.

In summary, the link density of the underlying communication topology acts as a trade-off parameter for COHDA, resulting in either less simulation steps or less messages sent in the process of the heuristic. However, this effect is less present in topologies with larger link densities. The effect on solution quality is minimal.

5.2 Message Delays

For the evaluation of the effects due to delayed messages in the communication layer, we have to distinguish synchronous from asynchronous approaches. In the context of our study, the former are characterized by the existence of synchronization points. These define algorithmic *phases*, such that all agent's actions within a specific phase have to be completed before the next phase can start. Moreover, the agent's actions do not depend on each other within a single phase. For example, in market-based approaches as described in Sect. 3, each *call for bids* forms such a phase. The central auctioneer waits and collects proposals (and refusals) until all agents have answered. Only then the answers are being evaluated. Hence, message delays have no influence on the structural process in such a situation, and there is no effect regarding solution quality or communication expenses. But it is easy to see that message delays indeed have a proportional influence on the run-time, as they directly affect phase durations. Besides market-based approaches, the EPOS approach is synchronous as well and thus shows the same effects.

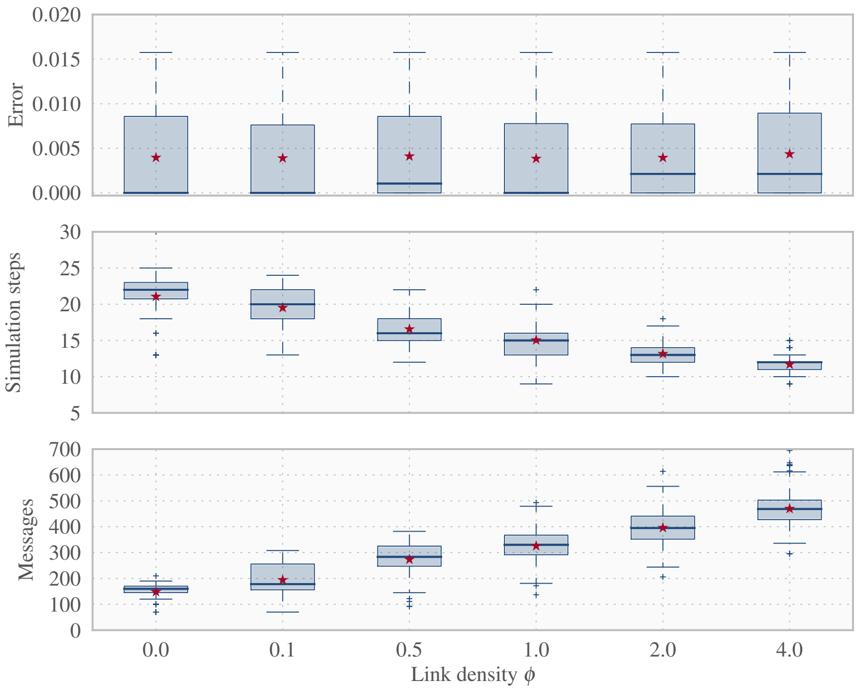

Fig. 1. Simulation results for different link densities in the COHDA approach

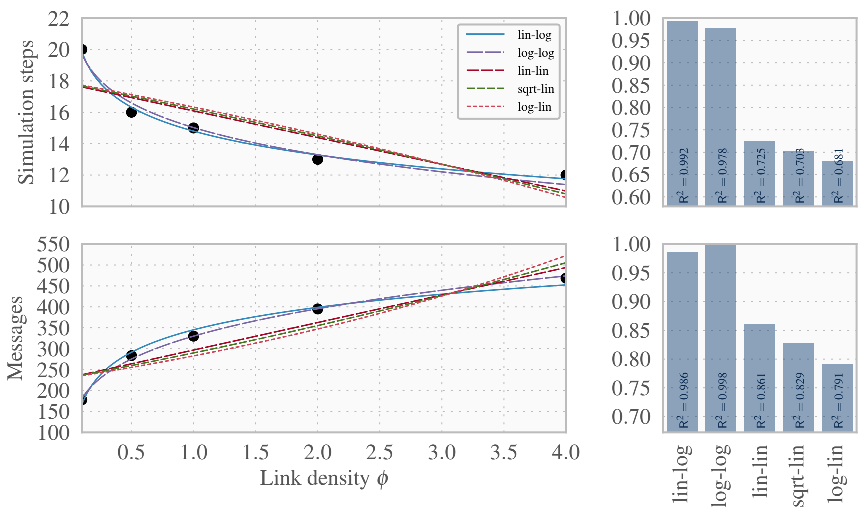

Fig. 2. Regression results for different link densities in the COHDA approach

On the other hand, asynchronous approaches are characterized by the absence of synchronization points. In these approaches, message delays can have a severe impact on the overall progress, because they may change the order of actions that exert influence on each other. In extreme cases, even the order of messages sent by a single agent can be disturbed, such that two subsequently transmitted messages arrive in the opposite order. From the presented approaches in Sect. 3, both the *Stigspace* and the COHDA approach are prone to such effects. For evaluating these, we set up our simulation environment as follows: The delivery of sent messages is delayed by the simulation core for a random number of simulation steps. The actual delay is determined for each message at runtime by calculating a uniformly distributed random number from the interval $[1, d_{\max}]$. Hence, the parameter $d_{\max}$ determines the maximal possible message delay per simulation run. We studied $d_{\max} \in \{1, 2, \dots, 10\}$.

For the *Stigspace* approach, the results are quite sobering. With no message delays, the approach exhibits synchronous behavior: All agents first read the *Stigspace* in parallel and subsequently write their adapted solutions (i. e. schedule selections in the considered use case) back into the *Stigspace*. As there is no further coordination mechanism, the system shows no convergence in this case. Instead, in almost half of the simulation runs, the system started oscillating between some solutions. In the other half of the simulation runs no oscillations occured, but no trends towards superior solutions were visible either, such that no convergence was possible and the simulations had to be stopped manually. By introducing message delays as defined above, the access to the *Stigspace* is partially being desynchronized, which effectively prevents oscillations, but still shows no convergence. Only in the other extreme, i. e. configurations with very large message delays in the order of $d_{\max} \approx m = 10$, a slight trend towards optimal solutions becomes visible. This indicates that the approach operates properly only with a *sequential* access paradigm for the *Stigspace*. Thus we conclude that the approach is not capable of handling variation in form of message delays at all.

For the COHDA approach, the results of the simulation study are summarized in Fig. 3. Similar to the experiment regarding varying link densities, all results in this experiment generally show the same solution quality with no noticeable trend. Hence, the quality of the solutions produced by COHDA is almost independent from possible message delays induced by the underlying communication technology. The run-time of the heuristic in terms of simulation steps rises constantly with increasing delays, while the amount of sent messages seems to converge to a fixed value. More information on this reveals the according regression analysis in Fig. 4. These results show a linear intercorrelation between message delays and simulation steps. Similar to the experiment regarding varying link densities, the most likely intercorrelation models regarding the effect of message delays on the amount of sent messages are both the logarithmic model and the power model. The power model here is estimated with a quite small elasticity coefficient $\beta_1 = 0.25$ and thus again rather describes a logarithmic intercorrelation.

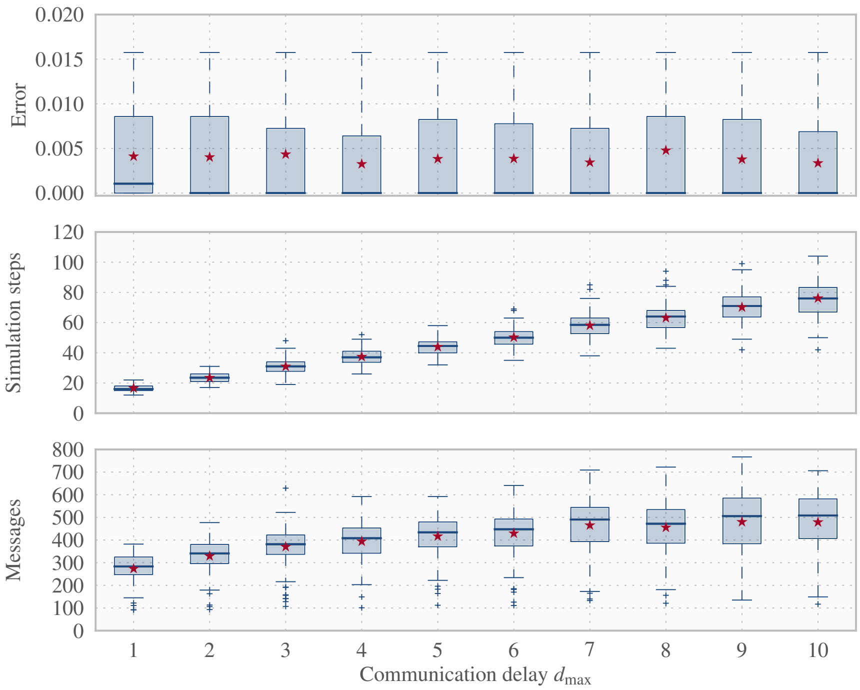

Fig. 3. Simulation results for different message delays in the COHDA approach

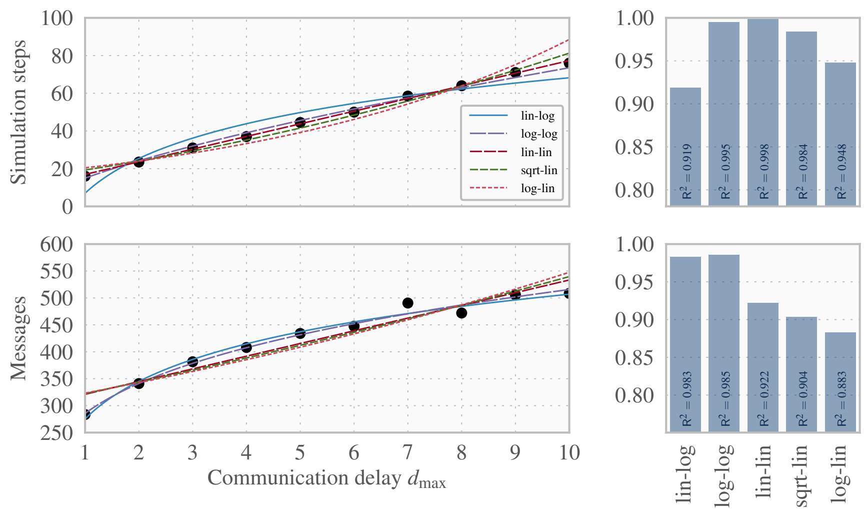

Fig. 4. Regression results for different message delays in the COHDA approach

In summary, the message delays induced by the underlying communication technology have a direct proportional influence of run-time of the COHDA approach in terms of simulation steps, while the amount of sent messages is less sensitive to this sort of variation. The effect on solution quality is minimal.

5.3 Reaction Delays

We define reaction delays as follows: In the progress of an approach, messages between agents are being delivered without delay, but instead the receiving agent will wait for a certain amount of time before processing the contents of the received message. As a side effect, an agent might receive multiple messages before it processes them all at once. Similar to message delays, we have to distinguish synchronous from asynchronous approaches in order to evaluate the effects of varying reaction delays. At first glance, reaction delays seem to have the same effects as message delays: For the synchronous approaches, there is no effect on the structural process, i. e. neither on solution quality nor on communication expenses, because the agent's actions within the same algorithmic phase are independent from each other (cf. Sect. 5.2). But again the run-time will increase proportionally with increasing reaction times due to the larger phase durations.

For the considered asynchronous approaches, we again have to look at the actual behavior of the approaches in simulation. Hence, in our simulation study, the actions of agents upon incoming messages are delayed by the simulation core for a random number of simulation steps. This is realized quite similar to the message delays above by calculating a uniformly distributed random number from the interval $[1, r_{\max}]$ as reaction delay. Hence, the parameter $r_{\max}$ determines the maximal possible delay per simulation run. We studied $r_{\max} \in \{1, 2, \ldots, 10\}$.

Simulations of the *Stigspace* approach show the very same behavior as with message delays. So indeed, for this approach, there is no difference between both sources of variation. Following the results from Sect. 5.2, we conclude that the *Stigspace* approach was designed with asynchronicity in mind, but lacks the necessary coordination mechanism to actually be able to handle the variation that will be present in a true asynchronous environment.

Finally, the simulation results for the COHDA approach regarding reaction delays are summarized in Fig. 5. Again, all results in this experiment generally show the same solution quality, and no specific correllation to the amount of reaction delays is visible. Hence, the quality of the solutions produced by COHDA is almost independent from differing reaction delays of the agents. Similar to the effects of varying message delays, the run-time in terms of simulation steps increases with larger reaction delays. Interestingly, the amount of messages *decreases* at the same time. Obviously, the COHDA heuristic benefits from variation by differing processing speeds in the deployed agents regarding the amount of coordination that is needed to converge to a joint solution in the collaborative MAS. The regression analysis for this experiment (see Fig. 6) again reveals a linear intercorrelation between reaction delays and simulation steps. Regarding the amount of sent messages, a logarithmic decrease is detected, indicating a less sensitivity here in comparison to the simulation steps.

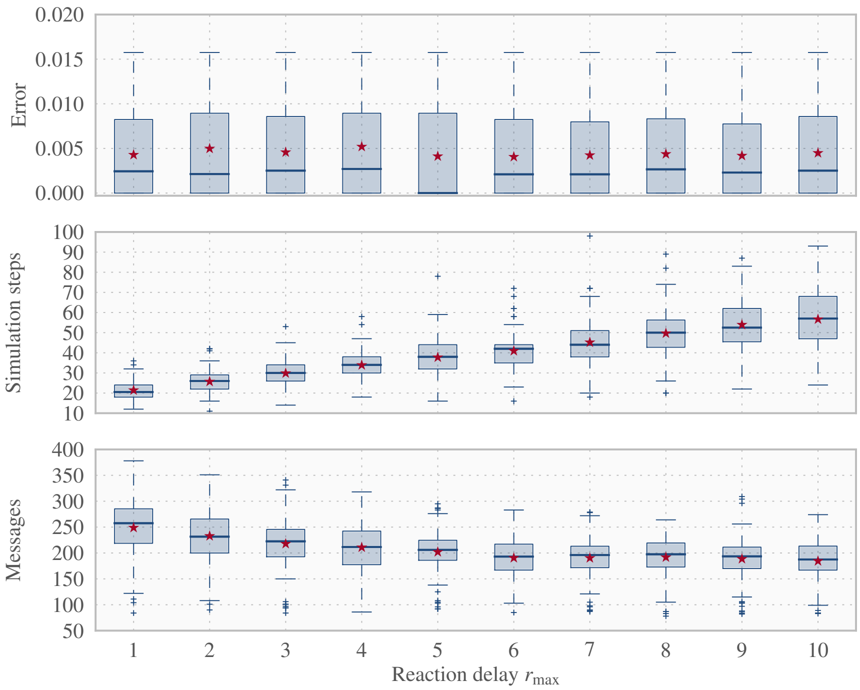

Fig. 5. Simulation results for different reaction delays in the COHDA approach

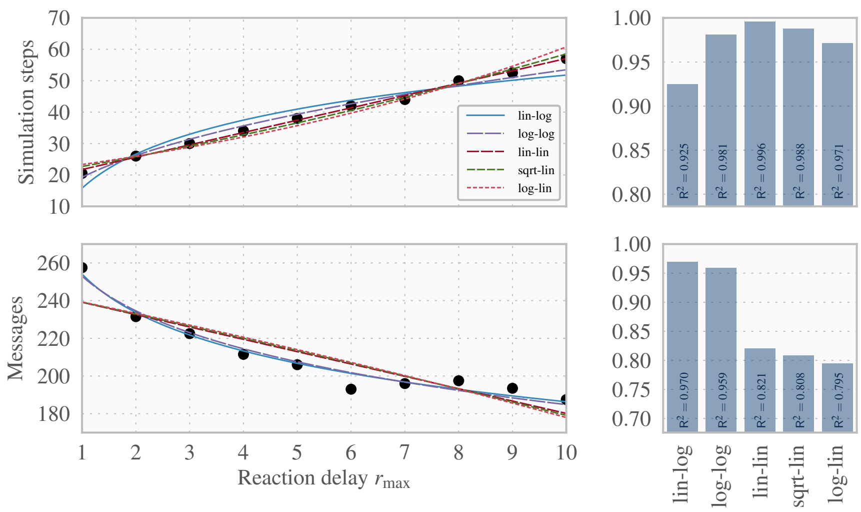

Fig. 6. Regression results for different reaction delays in the COHDA approach

In summary, the varying reaction delays induced by differing processing speeds in the deployed agents have a direct proportional influence on run-time in terms of simulation steps, while the amount of sent messages actually decreases with larger variation. The effect on solution quality again is minimal.

6 Conclusion

In this paper, we studied the effects of variation on collaborative multi-agent systems solving a combinatorial optimization problem. Our work is motivated by the use case of scheduling distributed energy resources within self-organized virtual power plants. For this, we presented the multiple-choice combinatorial optimization problem, MC-COP, and described four different solution strategies for such a task.

Our study is focused on the sources of variation that are likely to be faced in deployed field applications: arbitrary communication topologies, varying message delays, and differing reaction delays of the participating agents. In cases where the resulting effects were not obvious, we performed simulation experiments using a simulation system that is capable of simulating the considered variation types. We used synthetic problem instances that allowed us to interpret the simulation results independently from any specific use case. The simulation results were analyzed using descriptive statistics for determining qualitative properties regarding the general type of influence for each source of variation on the three performance indicators solution quality, run-time and amount of transferred messages, respectively. Further, regression analyses have been performed to detect the quantitative intercorrelations between those properties.

To interpret the results of our study, we categorized the presented solution strategies into synchronous and asynchronous approaches. Our findings indicate that synchronous and asynchronous approaches behave quite differently: Due to the existence of synchronization points, the structural process of the considered synchronous approaches is unaffected by both message delays and reaction delays. Therefore, those approaches only suffer in terms of run-time penalties from these types of variation. In contrast, the presented asynchronous approaches do not have inherent synchronization points, such that their structural process directly depends on the order of exchanged messages. The order is affected by both message delays and reaction delays. In the *Stigspace* approach, this effectively inhibits convergence and renders the approach infeasible in such situations. The COHDA approach, however, is able to handle these delays quite well: The solution quality exhibited by COHDA is almost independent from each considered source of variation, rendering the heuristic very robust in this regard. Both communication delays and reaction delays affect the run-time of the heuristic directly proportional, while the amount of messages is less sensitive to these delays. Interestingly, the amount of messages increases with larger communication delays, but decreases with larger reaction delays. This indicates that COHDA actually benefits from that source of variation regarding communication expenses. Moreover, from the considered approaches, COHDA is the only one that supports varying communication topologies. Here, the link density of the underlying communication network acts as a trade-off with respect to both simulation length and messages, where a higher density yields a faster convergence using more messages, and vice versa.

Our contribution intended to extend the research line of the influence of variation on multi-agent systems started by Campbell et al. [9], which was subsequently followed up by Anders et al. [10]. In summary, our study indicates that, depending on the type of the solution strategy, variation can have more or less severe impacts. Synchronous approaches tend to be robust against the considered sources of variation, but compensate this with increasing run-time. Asynchronous approaches, while having other advantages, are more prone to these factors. In this regard, the *Stigspace* approach suffers heavily from variation, whereas the COHDA approach was explicitly designed to overcome the difficulties induced by variation, and in one case actually benefits from it. But as we analyzed the effects from the considered sources of variation independently from each other, future work would be to study higher order effects as well.

We conclude that, in order to build reliable systems, the potential effects of variation should directly be accounted for when constructing collaborative MAS for field applications. Especially asynchronous approaches should be designed with care, as those can exert their possible benefits over synchronous approaches to the full extent only if the targeted environment is carefully considered in the design of the approach. Although we presented our work in the context of a specific use case, the general methodology of studying the considered sources of variation by conducting parameter variation experiments based on synthetic problem instances, followed by regression analyses, is applicable to a wide range of algorithms. We suggest adopting this methodology as a guideline for future developments that are targeted at field applications.

Acknowledgments. The simulations for this contribution were performed at the HPC Cluster HERO, located at the University of Oldenburg (Germany) and funded by the DFG through its Major Research Instrumentation Programme (INST 184/108-1 FUGG) and the Ministry of Science and Culture (MWK) of the Lower Saxony State.

We thank the anonymous reviewers and the program committee chairman for their detailed comments.

References

1. Ferber, J.: Multi-Agent Systems: Introduction to Distributed Artificial Intelligence. Addison-Wesley – An Imprint of Pearson Education (1999)
2. Shehory, O., Kraus, S.: Methods for task allocation via agent coalition formation. Artificial Intelligence **101**(1-2) (1998) 165–200
3. Shehory, O., Kraus, S.: Feasible formation of coalitions among autonomous agents in nonsuperadditive environments. Computational Intelligence **15**(3) (1999) 218–251

4. Griffiths, N., Luck, M.: Coalition formation through motivation and trust. In: Proc. of the 2nd Int. Joint Conf. on Autonomous Agents and Multiagent Systems, ACM (2003) 17–24
5. Tröschel, M., Appelrath, H.J.: Towards reactive scheduling for large-scale virtual power plants. In Braubach, L., Hoek, W., Petta, P., Pokahr, A., eds.: Multiagent System Technologies. Volume 5774 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2009) 141–152
6. Nieße, A., Lehnhoff, S., Tröschel, M., Uslar, M., Wissing, C., Appelrath, H.J., Sonnenschein, M.: Market-Based Self-Organized Provision of Active power and Ancillary services: An Agent-Based Approach for Smart Distribution Grids. In: Complexity in Engineering (COMPENG), 2012. (2012)
7. Beer, S., Sonnenschein, M., Appelrath, H.J.: Towards a Self-Organization Mechanism for Agent Associations in Electricity Spot Markets. In Heiß, H.U., Pepper, P., Schlingloff, H., Schneider, J., eds.: Informatik 2011 – Proceedings der 41. GI-Jahrestagung. GI-Edition – Lecture Notes in Informatics (LNI) P-191, Bonner Köllen Verlag (2011)
8. Ashby, W.: Requisite variety and its implications for the control of complex systems. Cybernetica **1**(2) (1958) 83–99
9. Campbell, A., Riggs, C., Wu, A.S.: On the impact of variation on self-organizing systems. In: Fifth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2011), IEEE Computer Society (2011) 119–128
10. Anders, G., Hinrichs, C., Siefert, F., Behrmann, P., Reif, W., Sonnenschein, M.: On the Influence of Inter-Agent Variation on Multi-Agent Algorithms Solving a Dynamic Task Allocation Problem under Uncertainty. In: Sixth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2012), IEEE Computer Society (2012) 29–38
11. UCTE: UCTE Operation Handbook – Policy 1: Load-Frequency Control and Performance. Technical Report UCTE OH P1, Union for the Co-ordination of Transmission of Electricity (2009)
12. Abarrategui, O., Marti, J., Gonzalez, A.: Constructing the active european power grid. The Online Journal on Power and Energy Engineering (OJPEE) **1**(4) (2010)
13. Hinrichs, C., Lehnhoff, S., Sonnenschein, M.: A Decentralized Heuristic for Multiple-Choice Combinatorial Optimization Problems. In: Operations Research Proceedings 2012, Springer (2014) 297–302
14. Lisovich, M.A., Mulligan, D.K., Wicker, S.B.: Inferring Personal Information from Demand-Response Systems. IEEE Security & Privacy Magazine **8**(1) (2010) 11–20
15. AlAbdulkarim, L., Lukszo, Z.: Integrating information security requirements in critical infrastructures: smart metering case. International Journal of Critical Infrastructures **6**(2) (2010) 187–209
16. Chapman, A.C., Rogers, A., Jennings, N.R., Leslie, D.S.: A unifying framework for iterative approximate best-response algorithms for distributed constraint optimization problems. The Knowledge Engineering Review **26**(04) (2011) 411–444
17. Kok, J.K., Warmer, C.J., Kamphuis, I.G.: Powermatcher: Multiagent control in the electricity infrastructure. In: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems (AAMAS '05), ACM Press (2005)
18. Wedde, H.F., Lehnhoff, S., Rehtanz, C., Krause, O.: Bottom-Up Self-Organization of Unpredictable Demand and Supply under Decentralized Power Management. In: 2008 Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems, IEEE (2008) 74–83

19. Anders, G., Siefert, F., Steghöfer, J.P., Seebach, H., Nafz, F., Reif, W.: Structuring and Controlling Distributed Power Sources by Autonomous Virtual Power Plants. In: IEEE Power and Energy Student Summit (PESS 2010), IEEE Power & Energy Society (2010)
20. Pournaras, E., Warnier, M., Brazier, F.M.: Local agent-based self-stabilisation in global resource utilisation. International Journal of Autonomic Computing **1**(4) (2010)
21. Li, J., Poulton, G., James, G.: Coordination of Distributed Energy Resource Agents. Applied Artificial Intelligence **24**(5) (2010) 351–380
22. Horowitz, S.H., Phadke, A.G.: Power System Relaying. 3 edn. John Wiley & Sons Ltd. (2008)
23. Geelen, D., van Kempen, G., van Hoogstraten, F., Liotta, A.: A wireless mesh communication protocol for smart-metering. In: Computing, Networking and Communications (ICNC), 2012 International Conference on, IEEE (2012) 343–349
24. Fischer, M.J., Lynch, N.A., Paterson, M.S.: Impossibility of distributed consensus with one faulty process. Journal of the ACM **32**(2) (1985) 374–382
25. Nieße, A., Sonnenschein, M.: Using Grid Related Cluster Schedule Resemblance for Energy Rescheduling. In Donnellan, B., Martins, J.a., Helfert, M., Krempels, K.H., eds.: SMARTGREENS 2013 – 2nd International Conference on Smart Grids and Green IT Systems, INSTICC Press (2013) 22–31
26. Hinrichs, C., Lehnhoff, S., Sonnenschein, M.: Paving the Royal Road for Complex Systems: On the Influence of Memory on Adaptivity. In Pelster, A., Wunner, G., eds.: International Symposium Selforganization in Complex Systems: The Past, Present, and Future of Synergetics, Springer (in press).
27. Han, B., Leblet, J., Simon, G.: Hard Multidimensional Multiple Choice Knapsack Problems, an Empirical Study. Computers & Operations Research **37**(1) (2010) 172–181
28. Hinrichs, C., Sonnenschein, M., Lehnhoff, S.: Evaluation of a Self-Organizing Heuristic for Interdependent Distributed Search Spaces. In Filipe, J., Fred, A., eds.: International Conference on Agents and Artificial Intelligence (ICAART 2013). Volume 1., INSTICC Press (2013) 25–34
29. Doval, D., O'Mahony, D.: Overlay Networks: A Scalable Alternative for P2P. IEEE Internet Computing **7**(4) (2003) 79–82
30. Strogatz, S.: Exploring complex networks. Nature **410**(6825) (2001) 268–76