POLICY EVOLUTION FOR LARGE SCALE ELECTRIC VEHICLE CHARGING CONTROL

Stephan Hutterer[a], Michael Affenzeller[b], Franz Auinger[a]

(a) University of Applied Sciences Upper Austria, School of Engineering and Environmental Sciences, Stelzhamerstraße 23, 4600 Wels, AUSTRIA;
(b) University of Applied Sciences Upper Austria, School of Informatics / Communication / Media, Softwarepark 11, 4232 Hagenberg, AUSTRIA;

[a] {stephan.hutterer, franz.auinger}@fh-wels.at
[b] michael.affenzeller@heuristiclab.com

ABSTRACT

In future smart electric grids, the control of electric vehicles' charging processes will be a central aim of demand side management. While this control enables the avoidance of possible critical peak-load values, the optimal coordination with supply from fluctuating renewables offers promising possibilities for power grid operation. Within this work, an optimization approach will be proposed that uses evolutionary optimization for computing performant control-policies for all EVs within a complex system. These policies are able to satisfy the EV users' energy demand on the one hand, while guaranteeing secure operation of the power grid on the other hand. Considering a high amount of EVs allows further the optimal integration of e-mobility into large-scale distribution networks.

Keywords: Electric Vehicle Charging Control, Probabilistic Power Flow, Evolutionary Optimization, Simulation Optimization

1. INTRODUCTION

Optimal integration of electric vehicles (EVs) into modern power grids plays an essential role in future power system operation and control. Numerous investigations have been performed in order to identify optimal charging policies for meeting objectives like peak-shaving, optimization of power quality metrics or maximal usage of power from renewable sources. Especially this interaction of zero-emission supply plants and electrified vehicles is seen as central concern, since the usage of energy from renewables directly influences the reachable environmental benefit of electric vehicles. Here, both the supply as well as the demand side show nondeterministic behavior which has to be tackled in some way. Therefore, a simulation-based optimization approach will be demonstrated, that uses metaheuristic algorithms for finding optimal charging schedules of an electric car fleet within a given system. This approach is capable of considering both, the physical power grid as well as the individual electrified traffic through probabilistic simulation models, where all nondeterministic influences can be incorporated dynamically into the heuristic search process. Each solution candidate will be evaluated a sufficient number of times through simulation in order to increase the accuracy of the performance estimation within an uncertain environment.

2. OPTIMAL CHARGING CONTROL

Various researchers examine the problem of integrating electric vehicles optimally into power grids, where direct control of charging power is seen as advantageous for reaching optimal load characteristics (Clement, 2008; Clement, 2009; Sortomme, 2011). Central challenge beside the formulation and computation of the optimization problem itself is the consideration of the individual behavior that mainly characterizes electric vehicle charging load. Different approaches try to tackle this task using static load profiles (Clement, 2008), representations of behavior using Queuing Theory (Vlachogiannis, 2009) or simulation via Monte Carlo methods (Sortomme, 2011).

All these approaches generally have in common that they try to compute static load profiles that are later used within certain optimization methods. Thus, there is no interrelated process that incorporates probabilistic behavior during the search for optimal solutions. Especially when talking about optimization in uncertain systems, simulation-based optimization with heuristic algorithms has been applied to various fields of applications and will be the central approach within this work. Here, with probabilistic simulation models, the uncertain system can be modeled holistically consisting of traffic simulation, probabilistic models of renewable supply as well as the power grid simulation model for the computation of resulting load flows.

3. SIMULATION-BASED POLICY EVOLUTION

The complete system architecture is shown in Figure 1:
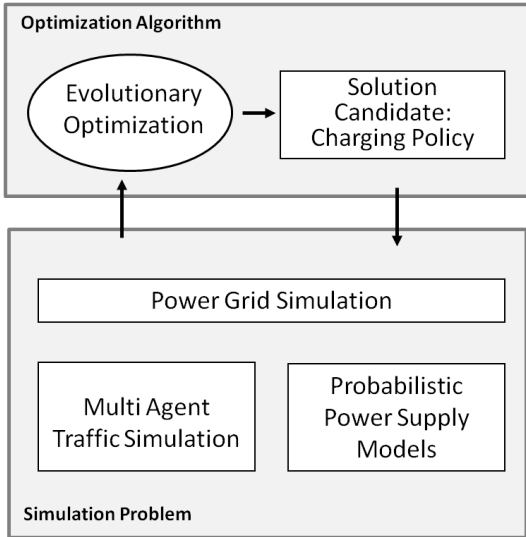
Figure 1: Simulation Optimization

Here, the problem represented by simulation consists of the distribution grid load flow simulation itself, probabilistic models for fluctuating supply from renewable plants as well as the simulation of the EV fleet. The solution candidate represented by a charging policy for all EVs is passed to the simulation for evaluation, returning its resulting fitness value to the optimization algorithm.

In the end, the finally best found charging policy should satisfy end-users energy demand while considering probabilistic driving behavior of EVs (traffic simulation), guaranteeing secure distribution grid operation (power grid simulation) as well as maximizing usage of power from fluctuating renewables (probabilistic power supply models).

3.1. Policy Optimization

In existing research literature on integrating EV fleets into distribution grids, the common approach is to implement an optimization procedure that computes optimal charging schedules based on existing knowledge and forecasted system behavior in advance. But in an uncertain and volatile system such as the underlying one consisting of probabilistically behaving agents and intermittent power supply from renewable plants, it would be more appropriate to make charging decisions on the fly, reacting to dynamic situations quickly and in a flexible manner.

Therefore, a policy-based approach is the central aim of this work. Here, each agent (EV) receives a flexible policy rather than a static schedule that makes it react to its environment dynamically during operation, but in a globally optimal manner when deciding about the agent's charging. This policy is principally the same for all agents, but using individual data from agent's environment, it leads to agent-specific charging behavior.

The basic concept is indicated in Figure 2, where the policy evaluation is indicated for a given EV that

arrives at an arbitrary location which is equipped with charging infrastructure.
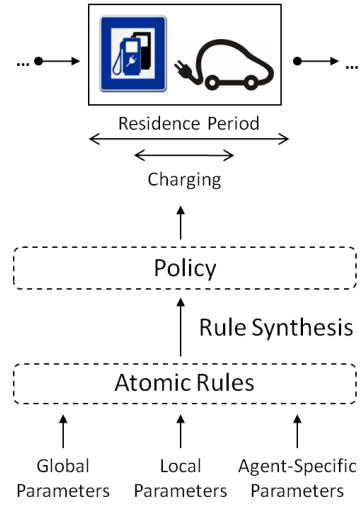
Figure 2: Policy Control

.

Principally, the optimized policy which finally decides the EV's charging power at a given time step is synthesized from atomic rules that consider agent-specific parameters from its environment. Out of these parameters, atomic rules are used to compute information out of them for evaluating EVs power demand as well as the state of its environment. Here, three different parameter classes can be distinguished from each other:

- Agent-specific parameters concern the EV's driving behavior, like its residence time at the actual charging station or its likelihood of getting parked at another charging spot later on.
- Local parameters consider other EVs immediately affecting the local situation in the power grid. For example, if the power grid is stressed locally because of a high amount of EVs charging at the same bus, their charging power has to be reduced in the next time step in order to avoid critical power flow conditions.
- Global parameters consider information describing the whole system's state, like the total load to the distribution grid, totally expected supply from renewables or financial aspects considering costs of electrical power supply.

Using these parameters as input for the atomic rules, each rule delivers a numeric result in the interval [0,1] that defines the agent's priority for charging. 0 would indicate that the corresponding EV should not charge at the actual time step, 1 advises it to charge with maximum power. Since a variety of criteria has to be taken into account for computing the optimal charging power of an EV, as can be estimated from the parameter

classes defined above, multiple rules have to be defined that finally have to be merged in some way. Table 1 gives an overview of all defined atomic rules.

Tabelle 1: Atomic Rules

Rule	Acronym
Total Residence Time so Far	RT
Estimated Time to Departure	ETTD
Passed Residence Time at Location	PRT
Actual Irradiance	AI
Past Irradiance During PRT	PI
Estimated Irradiance to ETTD	EI
Actual Wind Speed	AWS
Past Wind Speed During PRT	PWS
Estimated Wind Speed to ETTD	EWS
Actual Base Load	ABL
Past Base Load During PRT	PBL
Estimated Base Load to ETTD	EBL
Actual Price	AP
Past Price During PRT	PP
Estimated Price to ETTD	EP
Distance to Peak Load	DTB
Mean MVA Rating	MMVA
Number of EVs Same Location	NREVL
Mean number EVs Same Location During PRT	MNREVL
Number of EVs Charging Globally	NREVC
Number of EVs charging, Same Location	NREVCL
Mean Charging Rate per EV Globally	MCR
Mean Charging Rate, Same Location	MCRL
Agent's Already Charged Energy	ACE

More detailed information about the atomic rules can be obtained from (Hutterer, 2012). Summing up, all these rules combined are capable of considering not only the single EV's needs, but also describe the global system's state concerning power grid operation and behavior of the total EV fleet. These rules now have to be combined in an appropriate way in order to compute the final charging power of an affected EV.

3.2. Rule Synthesis

Within this work, two approaches will be compared for constructing the final policy out of atomic rules. The aim of this so called "rule synthesis" is to compute a final value that describes the charging decision out of the set of atomic rules that are needed in order to consider all needed information from the agent's environment.

3.2.1. Synthesis with Linear Combination

The first approach uses a fixed mathematical structure given in Equation 1. Here, the agent's charging rate (CR) at time step i is computed using a linear combination of all rules r, each rule multiplied by a specific weight w, divided by the amount of rules j. This kind of rule synthesis is a common approach from production logistics as used in (Vonolfen, 2011) and

(Beham, 2009). Here, the control variables that are manipulated during the heuristic search process are the weights w_j that describe the impact of each rule. For this kind of real-valued optimization, evolution strategies according to (Beyer, 2002) are applied.

$$CR_i = \frac{\sum_{j=1}^{J} r_{j,i} * w_j}{j} \quad (1)$$

Even if the rule synthesis using a linear combination is quite intuitive and leads to competitive results, it seems to be inflexible, disregarding the possibility of identifying potential nonlinear relationships between atomic rules. Therefore, a second approach is introduced that allows a more flexible, nonlinear combination of atomic rules, namely genetic programming (GP) (Affenzeller, 2009).

3.2.2. Synthesis with Genetic Programming

Extending the principle concept of genetic algorithms, GP uses evolutionary-inspired concepts for the heuristic search process, but is able to evolve computer programs. Within the herein described work, these computer programs take the appearance of structured trees, where leafs represent rules as defined before, that are combined by arbitrary mathematical operators which are incorporated by inner nodes. This kind of solution representation allows arbitrary mathematical combinations of atomic rules.

To give some overview on GP, finding first research activities in the 1980s, the computationally expensive concept of GP was pushed majorly by the steady increase of computational power in the last two decades. One of the most important publications in this field was (Koza, 1992), stating GP as automated invention machine for numerous practical applications like the artificial ant problem or later applications of symbolic regression (Affenzeller, 2009), to name the most popular ones, while (Langdon 2002) finally provides profound analysis in the context of GA schema analysis. This ability of GP to automatically construct new solutions (programs) to a given problem is enabled by its special kind of solution representation, that is not restricted to a fixed structure (like fixed-length one-dimensional array as in standard GA), but forms a hierarchical computer program of variable length, consisting of functions and terminals. In the herein presented application, functions are inner nodes of the structured tree, while terminals can be constants or atomic rules. Figure 3 gives an exemplary tree that could represent a policy for the addressed problem.

Here, inner nodes that represent functions are indicated in dotted style, while terminal nodes are plotted in solid style. In this case, the policy would consider the estimated time to departure of the appropriate EV, the actual irradiance and thus the supply from photovoltaics, as well as the mean charging

rate of all other EVs in order to not stress the distribution grid with peak charging load. Out of this mathematical combination, finally a numerical value is derived that represents a charging decision.

Figure 3: Exemplary GP Solution

The great advantage of this kind of flexible solution representation compared to the application of a fixed-structure linear combination is that GP is able to find nonlinear coherences between atomic rules with variable length. Using any arbitrary combination of mathematical operators as inner nodes, the degree of freedom for finding performant policies gets increased drastically. Further, since GP is not constrained to use the hole set of atomic rules for solution creation, simpler policies can be found too. In the end, its disadvantage is that the possible solution space is increased drastically.

For overcoming this problem and pruning the solution space, the possible grammar is restricted to the following operators in Table 2. The grammar in this case defines the set of functions (inner nodes) that is applied for evolving solution candidates during the genetic search process. Further information on usable grammar can be obtained in the appropriate literature (Affenzeller, 2009) as well as in HeuristicLab (www.heuristiclab.com) which is used as heuristic optimization framework.

Table 2: GP Grammar

| Arithmetic Operators	{+, -, x, /}
Conditional Operators	Conditions {IfThenElse}
	Comparisons {<,>}

3.3. Policy Evaluation
The principal process of the policy evaluation can be obtained from Figure 4: in each time step of the simulation, if the agent remains at a charging station, the policy evaluation is initiated in order to compute the resulting charging power. After gathering all the information the agent needs (global, local as well as agent-specific parameters), the respective outcome of the atomic rules is computed. Combining these results according to the rule synthesis method, the final charging power can be derived. The atomic rules are generally constructed such that each rule as well as its respective weight results in a numeric value in the interval [0,1]. Thus, for the linear combination, the final results exists in the interval [0,1] as well. This value therefore is interpreted as charging rate and is

multiplied by the maximum possible charging power per EV. Thus, when using rule synthesis with linear combination, no invalid charging power can occur from the policy, as long as the decision variables are kept within [0,1].

When using rule synthesis with GP, the policy directly outputs the desired charging power. Here, possibly invalid values may result from the solution candidate (negative charging power, too high charging power) because of the high degree of freedom when building the structured tree, which has to be managed in some way. A reason could be for example the addition of a constant or a multiplication of rules with some value. In this work, this is considered the following way: if the resulting value is less than 0 or greater than the maximum charging power, the value is set to 0 or the maximum value respectively and a penalty is added to the fitness term according to the degree of the violation. Since power-flow simulation may not converge in exceptional conditions (for example if the resulting charging load takes unmanageable values), this penalty is turned into a so called "death penalty" for the respective solution. Thus, if a solution candidate leads to non-convergence of the load-flow simulation, it is assumed to be useless.

3.4. Solution Evaluation
The evaluation procedure of a solution candidate is indicated in Figure 4.
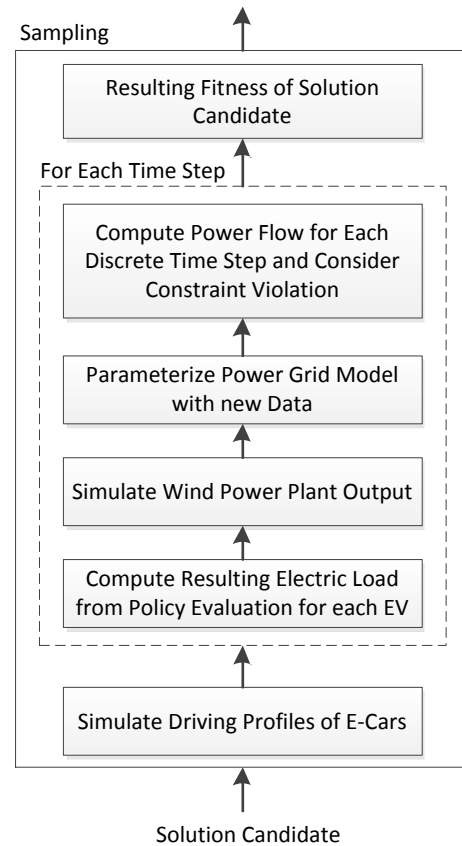
Figure 4: Solution Evaluation

When evaluating a solution candidate in simulation, first of all the traffic simulation is performed over the whole time interval, in order to describe the expected EV behavior. Having this behavior in form of computed driving profiles for each agent, its resulting charging power is computed for each time step that follows from the evaluation of the respective policy. With the charging load caused by all agents, further the power flow simulation is executed for each time step considering probabilistic injection from renewable sources, in order to compute the actually occurring load flow in the physical distribution grid. With the final power flow solution, all constraints as well as the objective function can be evaluated in order to derive the fitness of the solution.

3.5. Evaluation Under Uncertainty

When evaluating a solution candidate in an uncertain environment, estimating its real performance is a ubiquitous as well as challenging task which might be computationally expensive. This is due to the stochastic nature of the evaluation as well as the slow convergence of a performance measure estimator relative to the number of runs performed.

The obtained fitness from a simulation run can be formulized as

$$\tilde{f} = f + \varepsilon,$$

where the obtained fitness value $\tilde{f}$ deviates from the real fitness caused by some probabilistic noise ε. This noise is mainly defined in literature as being normally distributed (Stagge, 1998; Fu 2002) by $N(\mu, \sigma)$ with zero mean. While the evolutionary search proceeds rapidly it may happen that some other individual than the best is chosen as parent for the next generation, caused by an inaccurate estimate of f. This may lead to a decrease in progress velocity and may also lead the evolutionary search into unpromising regions of the search space. Thus, doing multiple evaluation runs and averaging over the obtained values of $\tilde{f}$ is important for estimating the candidates' real performance. Since The accuracy of this estimate cannot improve faster than $1/\sqrt{N}$, where N is the number of computed samples as indicated in Figure 4, choosing an appropriate value of N is essential when addressing computational costs of evaluation. Different approaches have been investigated both in literature as well as by the authors of this paper (Hutterer, 2012) for inferring N in an adaptive way during the search process. Within this paper, it is seen as sufficient to experimentally derive a performant value for N and fix it.

4. SETTING UP EXPERIMENTS

As highlighted in the introductory chapter, many researchers are investigating this optimal integration of EVs into distribution grids nowadays. When considering the treated power grid levels, these researchers mainly focus on quite low-level integrations in mostly radial distribution feeders or even lower. The special advantage of the herein used policy-based control approach, as discussed extensively in (Hutterer, 2012), is that it enables the consideration of huge EV fleets and thus consider their integration from a higher level point of view. Hence, in this work, larger distribution networks are considered that will be discussed as follows.

4.1. Large-Scale Distribution Grid Testcases

In order to guarantee universality for the considerations within this work, the well known IEEE distribution grid testcases[1] will be used and modified for representing valid test instances. Throughout the grid, a huge EV fleet is modeled, where each single agent can produce a charging load of maximum 11kW, related to a three-phase charging process with 400V and 16A, as exemplarily possible when using a Mennekes VDE (Type 2) plug connector. This configuration, as existing for example when charging the well known Tesla Roadstar, is certainly one of the most important technical specifications in this field in actual developments and is seen to get a common standard throughout EV-manufacturers.

The power grid simulation model is being downscaled such that the cumulated charging power of all agents sums up to 20% of the daily peak load maximally in each considered case. For representing individual electrified traffic from a power grid point of view, the relevant behavior that has to be modeled describes time interval and location of each EV when being parked to a charging station and thus being ready for charging. Based on real-world traffic data from an Austrian survey[2], two most relevant driving patterns can be extracted for a week day, namely the pattern of full-time and half-time workers. Within each pattern, three different locations are modeled for parking at home, at work and at any location in free time (shopping, education, entertainment). For each location, different probabilities for the existence of a charging infrastructure are modeled, describing a possible future infrastructure scenario from an actual point of view: at home, each EV user has an own charging station. At work, there is a probability of 50% that an appropriate infrastructure is available. For locations where potential users remain in free time, this probability is assumed to be 25%.

The resulting charging load at a specific location is than being correlated to a corresponding bus within the distribution grid model. Within each simulation run, synthetic driving profiles are computed from prototype-

---

[1] Testcases provided by University of Washington, UW Electrical Engineering. (1999).
http://www.ee.washington.edu/research/pstca/

[2] Federal Ministry for Transport, Innovation and Technology, Verkehr in Zahlen 2007,
http://www.bmvit.gv.at/verkehr/gesamtverkehr/statistik/downloads/viz07gesamt.pdf, Retrieved 09.07.2012

profiles, being randomized in terms of driving time and residence time at specific locations. Thus, the probabilistic behavior of individual traffic can be modeled based on real-world data and incorporated into the evolutionary optimization process enabled by the simulation-based approach.

For modeling the power output of renewable sources, wind power plants as well as large-scale photovoltaic plants are added to the distribution network. For wind power modeling, the corresponding wind speed values at the plant sites are sampled from a Weibull-distribution as described in (Vlachogiannis, 2011), where their power curves are assumed such that each plant reaches its maximum output at cut-off windspeed. Using the sampled wind speed value, with the plant's power curve the resulting power output of the plant can be modeled.

Photovoltaic-plants follow a typical daily generation profile that is randomized in each time step with a standard deviation of 10%, considering a typical uncertainty in photovoltaic-generation forecasting. All renewable supply models are designed such that they cumulated produce in average 50% of the energy needed for all EVs in the system.

In order to consider realistic power grid conditions, the base load is modeled as described in the IEEE testcases, but randomized too for simulating a probabilistic demand side with a standard deviation of 4%.

A thorough discussion of the used modeling approach can be obtained from (Hutterer, 2012). The configurations for both test cases are shown in Tabelle 13, where the distribution of renewable plants throughout the grid model as well as the EV fleet are defined.

Table 3: Test Cases Configurations

14-Bus Testcase	
# EVs	960
# Photovoltaic Plants	3
# Wind Plants	2
Bus # with Photovoltaic Injection	6,8,10
Bus # with Wind Power Injection	3,12
Bus # with fixed Generation	2
Slack Bus #	1
118-Bus Testcase	
# EVs	4366
# Photovoltaic Plants	9
# Wind Plants	3
Bus # with Photovoltaic Injection	12,31,46,54,59,61,87,103,1 11
Bus # with Wind Power Injection	25,49,100
Bus # with fixed	10,26,65,66,69,80,89
Generation	
Slack Bus #	1

5. PROBLEM FORMULATION

A formal description of the optimization problem shall now be stated in order to underline the application: given a fleet of EVs within a distribution grid, a vector $Pc = [Pc1,1,...,Pci,n]$ describes the active charging power of each EV n at time step i over a given time interval. At the end of this considered planning frame, each EV must have received a specific amount of energy for satisfying its daily demand:

$$E_{\min,n} \leq \sum_{i=1}^{I} Pc_{i,n} * \Delta t_i$$

This constraint is valid assuming that batteries are big enough and the one-way distance of an EV does not lead to a low state of charge. Since additional load caused by related charging of electric vehicles can endanger power grid security, constraints have to be satisfied that ensure secure distribution grid operation. Thus, within each time step i, power flow constraints have to be considered. Steady-state security constraints can be formulated (Wood, 1996) for ensuring lower and upper bounds for generator real and reactive power output Pg and Qg,

$$Pg_{j,\min} \leq Pg_j \leq Pg_{j,\max}$$
$$Qg_{j,\min} \leq Qg_j \leq Qg_{j,\max} ,$$

maximal power flows over transmission lines Pf,

$$Pf_k \leq Pf_{k,\max} ,$$

as well as admissible voltage deviations ΔV ,

$$\Delta V_{k,\min} \leq \Delta V_k \leq \Delta V_{k,\max} ,$$

for all buses $j=1...J$ and all transmission lines $k=1...K$.

While satisfying all formulated constraints, the objective function shall be defined of minimizing financial costs of power supply.

Since charging power is restricted to a maximum value, an additional constraint has to be formulated when using the GP-based policy synthesis as discussed in section 3.2, being formulized as:

$$Pc_{i,n} \leq P_{c,\max} .$$

What is important to mention at this point is that the vector Pc as introduced above containing the charging power of each EV at each time step is never present in a static manner. Each value Pci,n results from a single evaluation of the policy within the simulation of a certain time step. Therefore, as indicated in Figure 4, the load flow in the system is computed within each time step in order to check the constraints. The finally obtained fitness function is stated in Equation 2, where $\overline{CV(Pc)}_i$ is a vector containing the quadratic violations of each constraint, multiplied by $\bar{k}$ being a vector with fixed weights of the constraint violation value relative to the financial cost function value.

$$\sum_{i=1}^{24}[Cf(Pc)_i + \overline{k} * \overline{CV(Pc)_i}] \quad (2)$$

Since the objective concerns financial costs of energy supply for charging electric vehicles, a daily price profile is assumed that is taken from the European energy exchange as used in (Hutterer, 2012).

6. EXPERIMENTAL RESULTS

For the experiments performed within this work, evolutionary algorithms are used, depending on the applicated policy synthesis approach. For optimizing the weights for the fixed-structure linear combination, Evolution Strategies (ES) are applied according to (Beyer, 2002). ES are generally performant metaheuristics for real-valued optimization problems and proven to be suitable for simulation-based optimization (Hutterer, 2012). For the GP-based evolution of policies, Genetic Algorithms (GA) are applied. Both classes of algorithms are executed in HeuristicLab based on their standard implementations. The finally used configurations can be obtained from Table 4 and Table 5.

Table 4: Algorithm Configurations 14-Bus Testcase

Type	(5+15)-ES
Manipulator	SelfAdaptiveNormalAllPositions-Manipulator
Recombinator	Average Crossover
Parents per Child	2
Stopping Criterium	5000 Generations
Sampling	Sample Each Solution 3 Times
Type	GA
Manipulator	MultiSymbolicExpressionTree-Manipulator
Recombinator	SubtreeCrossover
Population Size	250
Mutation Probability	15%
Stopping Criterium	200 Generations
Sampling	Sample Each Solution 3 Times

Table 5: Algorithm Configuration 118-Bus Testcase

Type	(5+10)-ES
Manipulator	SelfAdaptiveNormalAllPositions-Manipulator
Recombinator	Average Crossover
Parents per Child	2
Stopping Criterium	5000 Generations
Sampling	Sample Each Solution 6 Times
Type	GA
Manipulator	MultiSymbolicExpressionTree-Manipulator
Recombinator	SubtreeCrossover
Population Size	150
Mutation	15%
Probability	
Stopping Criterium	200 Generations
Sampling	Sample Each Solution 6 Times

Further details on the used configurations can be obtained from HeuristicLab and from the appropriate literature respectively (Affenzeller, 2009).

6.1. Results for the 14 Bus Testcase

The obtained best solution is shown in Table 6, showing the obtained weights for the given atomic rules. It can easily be seen, that most rules are weighted near to 1, in order to construct the final policy out of them.

The finally best found solution for the GP-based rule synthesis cannot be visualized at this point, since it forms a structured tree of length 28 (number of used nodes) and depth 5. Some statements can even be done: the best found tree uses only 9 out of the 24 atomic rules for synthesizing the policy. This means that the rules are highly correlated to each other (which is obvious) and not all of them are needed for finding valid policies. Table 7 shows same numeric results.

Table 6: Best Solution 14 Bus Testcase

Rule	Weight	Rule	Weight
RT	1	AP	0.9166
ETTD	1	PP	0.6943
PRT	1	EP	0.9166
AI	0.9613	DTB	0.9572
PI	0.9959	MMVA	0.8560
EI	0.2370	NREVL	0.6080
AWS	0.9814	MNREVL	0.9393
PWS	1	NREVC	0.8132
EWS	0.5910	NREVCL	1
ABL	1	MCR	0.9939
PBL	0.9861	MCRL	0.9548
EBL	1	ACE	0.6544

Table 7: Numeric Results 14 Bus Testcase

Fitness	Standard Deviation of Fitness: 100 Replications
Synthesis with Linear Combination	
476.20	1.04%
GP-based Synthesis	
512.01	3.1%

The fitness addresses the resulting costs (in Euro) for supplying energy for charging the EV fleet. These costs only address the energy-generation costs, which vary in a range of around 0.03 to 0.08 per kWh over a typical day at the European energy exchange (EEX). The price that a consumer would have to pay additionally contains taxes as well as a fee to the power grid operator. Each solution is evaluated over 100 simulation runs for final results in order to obtain its robustness within the uncertain environment. Therefore, the standard

deviation of the obtained fitness over these 100 runs is used as robustness estimator.

As can be seen in Table 7, the synthesis with fixed structure linear combination outperforms GP-based synthesis in both metrics. Thus, the harder heuristic search caused by an increased solution space when trying to find a valid solution with GP dominates its advantage of finding nonlinear coherences between rules.

Nevertheless, both approaches are capable of finding feasible (all constraints are satisfied) solutions with optimized financial costs of energy supply.

6.2. Results for the 118 Bus Testcase

As can be seen from the algorithm configurations, for the second testcase, quite lower population sizes have been used. This is due to the fact that this testcase considers 4366 EVs and therefore the evaluation of the policy when simulating its performance has to be executed more than 4 times more often than in the smaller testcase. Thus, population size has been decreased in order to keep the optimization computationally tractable.

Table 8: Best Solution 118 Bus Testcase

Rule	Weight	Rule	Weight
RT	1.0000	AP	0
ETTD	0.2406	PP	0.0812
PRT	1.0000	EP	0.6627
AI	0.3997	DTB	0.2803
PI	1.0000	MMVA	1.0000
EI	1.0000	NREVL	0.3945
AWS	1.0000	MNREVL	1.0000
PWS	0.3103	NREVC	0.0654
EWS	0.6412	NREVCL	1.0000
ABL	0.2891	MCR	0.9603
PBL	0.4047	MCRL	0.4853
EBL	1.0000	ACE	0.8274

Table 9: Numeric Results 14 Bus Testcase

Fitness	Standard Deviation of Fitness: 100 Replications
Synthesis with Linear Combination	
2556.13	6.0%
GP-based Synthesis	
2612.20	6.8%

The best found solution as visualized in Table 8 differs with respect to the smaller testcase drastically, considering a much higher variation in the single weights. Taking a look at the numerical comparisons in Table 9, once more the given fixed structure synthesis outperforms the GP-based one. Comparing the reached fitness-values of both testcases, in the 14-bus case costs of 0.4958 result per single EV, while these costs are increased to 0.5854 in the 118-bus case for the best solution. Since in both cases same generation costs are

assumed for the power grid simulation, it can be stated that the solution of the smaller testcase has better overall quality, proving that the optimization task in the second case seems to be harder.

CONCLUSION

In this work, a simulation optimization framework has been proposed that is capable of computing optimal charging decisions for a huge fleet of electric vehicles for optimally integrating them into distribution grids.

These decisions are principally performed using flexible policies that enable the EVs to act individually within a dynamic environment. The respective policies are evolved using evolutionary computation techniques, where the search space is spanned through a simulation ensemble consisting of the electric power grid model, electric vehicle behavior models as well as probabilistic supply models. Two different representations have been formulated and compared to each other that enable the synthesis of the final policy using a set of atomic rules. Since these variants for rule synthesis majorly influence the evolutionary search, comparisons have finally been performed for evaluating reachable solution quality based on two practical test instances.

REFERENCES

Affenzeller, M., Wagner, S., Winkler, S., & Beham, A., 2009. *Genetic algorithms and genetic programming: Modern concepts and practical applications*. Chapman & Hall, London, UK.

Beham, A., Kofler, M., Wagner, S., Affenzeller, M., Heiss, H., & Vorderwinkler, M., 2010: Enhanced Priority Rule Synthesis with Waiting Conditions, Proceedings of the *European Modeling and Simulation Symposium*, pp. 65-70, Morocco.

Clement K., Haesen, E., & Driesen, J.,2008. The impact of uncontrolled and controlled charging of plug-in hybrid electric vehicles on the distribution grid. *European Ele-Drive Conference*, Geneva, Switzerland.

Clement, K., Haesen, E., & Driesen, J., 2009. Stochastic analysis of the impact of plug-in- hybrid electric vehicles on the distribution grid. *20th International Conference on Electricity Distribution CIRED*, Prague, Czech Republic.

Langdon, W. B. & Poli, R, 2002. *Foundations of Genetic Programming*. Springer, Berlin Heidelberg New York.

Stagge, P., 1998. Averaging efficiently in the presence of noise. In A. Eiben (Ed.), *Parallel Problem*

Solving From Nature – PPSN V, Lecture Notes in Computer Science, Vol. 1498. pp. 188-197. Berlin/Heidelberg, Germany, Springer.

Sortomme, E., Hindi, M. M. , McPherson, J., & Venkata, M., 2011. Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses. *IEEE Transactions on Smart Grid*, 2(1), 198-205.

Vlachogiannis,John G., 2009. Probabilistic Constrained Load Flow Considering Integration of Wind Power Generation and Electric Vehicles, *IEEE Transactions On Power Systems*, 24 (4), pp. 1808-1817.

Vonolfen, S., Affenzeller, M., Beham, A., Lengauer E., & Wagner, S., 2011: Simulation-based evolution of resupply and routing policies in rich vendor-managed inventory scenarios, *Central European Journal of Operations Research (CEJOR)*, Accepted (available online: DOI: 10.1007/s10100-011-0232-5).

Fu, M. C., 2002. Feature article: Optimization for simulation: Theory vs. practice, *INFORMS Journal on Computing*, 14(3), pp. 192-215.

Beyer, H. G. and Schwefel, H. P., 2002. Evolution Strategies - A Comprehensive Introduction, *Natural Computing*, 1, pp. 3-52, Springer.

Hutterer, S., Auinger, F., & Affenzeller, M., 2012. Metaheuristic optimization of electric vehicle charging strategies in uncertain environment, *Proceedings of the International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)*, Istanbul Turkey.

Hutterer, S. and Affenzeller, M., 2012. Probabilistic Electric Vehicle Charging Optimized With Genetic Algorithms and a Two-Stage Sampling Scheme. *International Journal of Energy Optimization and Engineering* (accepted). IGI Global.

Koza, J. R., 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press, 1992.

Wood, A. J., & Wollenberg, B. F., 1996. *Power generation, operation, and control. Second edition*. Wiley-Interscience, New York, US:

AUTHORS BIOGRAPHY

Stephan Hutterer received his Diplomingenieur (FH) (graduate engineer) in Automation Engineering at the University of Applied Sciences Upper Austria. Since 2009, he is PhD candidate at the Johannes Kepler University Linz, Austria, working in the field of evolutionary optimization for intelligent power grids. His special research interests are simulation-based optimization for the integration of distributed renewable supply as well as electric vehicle fleets into power grids. He is member of the Heuristic and Evolutionary Algorithms Laboratory at the University of Applied Sciences Upper Austria.

Michael Affenzeller has published several papers, journal articles and books dealing with theoretical and practical aspects of evolutionary computation, genetic algorithms, and meta-heuristics in general. In 2001 he received his PhD in engineering sciences and in 2004 he received his habilitation in applied systems engineering, both from the Johannes Kepler University of Linz, Austria. Michael Affenzeller is professor at the University of Applied Sciences Upper Austria, Campus Hagenberg, and head of the Josef Ressel Center Heureka! at Hagenberg.

Dr. Franz Auinger received his Diplomingenieur (Master) in electrical engineering at Vienna University of Technology. PhD thesis in the field of automation objects and generic components for mechatronic engineering. Until february 2003 head of the group for intelligent automation and holonic at PROFACTOR. Since February 2003 professor for industrial informatics at the University of Applied Sciences Upper Austria at campus Wels in the field of embedded & real-time systems and since July 2006 head of the department for electrical engineering. Dr. Auinger holds publications in the area of holonic and distributed control systems, as well as smart electric grids.