Policy Function Approximation for Optimal Power Flow Control Issues

Stephan Hutterer[(a)], Michael Affenzeller[(b)]

University of Applied Sciences Upper Austria, School of Informatics / Communication / Media, Softwarepark 11, 4232 Hagenberg, AUSTRIA

[(a)] stephan.hutterer@heuristiclab.com, [(b)] michael.affenzeller@heuristiclab.com

ABSTRACT

In nowadays operations research, dynamic optimization problems build a central and challenging research topic. Especially in real-world systems such as electric power grids, dynamic problems occur where robust solutions need to be found that enable (near-) optimal control over time in volatile as well as uncertain power grid operation. The authors of this work identified the application of policy function approximation for suchlike problems, where an analytic functions needs to be found that takes an arbitrary state of the dynamic system and outputs appropriate control actions aiming at system-wide goals. Such an approach is very fruitful for robust optimization over time.

Applying this approach to two different problem classes in power grid research, this work aims at summarizing this work and identifying potential future issues.

Keywords: simulation optimization, power flow control, dynamic stochastic optimization, policy-function approximation

1. INTRODUCTION

Taking a look at power grid optimization, tasks that necessitate fast and robust dynamic control lie on hand. Taking exemplarily the general optimal power flow (OPF) problem, the aim is to find the optimal configuration of all controllable units for satisfying a given load situation, using steady-state representation of the power grid. Thus, the solution of this problem addresses exactly one stationary state J_t, disregarding possible states in the near future or eventual uncertain conditions in the system. Considering the system one time step later (J_{t+1}) due to changing conditions of weather, customer behavior or any other influence, the power flow in the system would change, hence, requiring a new solution to the optimal power flow problem further necessitated by the non-linear behavior of an electric power distribution system. Such a new computation would require a robust and fast-converging solution method, that guarantees quick support with a new optimal solution, independent of system complexity and starting point, which cannot be guaranteed by traditional steady-state OPF methods (Wang, 2007). This concern is further complicated by the steady increase of the number of control variables in smart grid applications (Hutterer, 2013b).

Thus, electric power systems fundamentally represent applications that require dynamic optimization techniques, respectively methods that enable optimization over time. The general scheme of approximate dynamic control with policy function approximation builds a fruitful ground for suchlike issues (Powell, 2012).

The rest of the paper is organized as follows: Section 2 proceeds with discussing the principles of simulation-based policy function approximation using evolutionary algorithms. Therefore, genetic programming (GP) can be identified as suitable metaheuristic search technique for evolving powerful control policies. In order to demonstrate the application, Sections 3 and 4 continue with illustrating two empirical studies when applying GP-based policy evolution to two practical scenarios, namely dynamic optimal power flow control for generation unit scheduling on the one hand, and a demand-side management related scenario for controlling distributed electric vehicle charging processes on the other side. Section 5 finally concludes the work und gives outlook to future issues.

2. SIMULATION-BASED EVOLUTIONARY POLICY FUNCTION APPROXIMATION

Policy function approximation can be understood as the general principle of providing an anticipatory policy $P(J_t)$, that outputs (near-) optimal control actions at runtime. The great advantage of such policy-based control schemes is the avoidance of any reoptimization during runtime after change of some state variables. Thus, instead of computing a static optimal solution, the policy function is optimized such that it leads to minimal expected costs in each possible state. This principle is illustrated in Figure 1. The authors of this work already identified the usage of policy function approximation techniques for the sake of optimizing power flow related tasks in smart grid engineering. A special scheme was applied in (Hutterer, 2013a; Hutterer 2013c) that uses evolutionary simulation optimization to evolve optimal policies that are trained within a dynamic simulation model and minimize some expected cost function. The application of simulation therefore enables the integration of systems' uncertainties (like uncertain weather or customer conditions) into the optimization process. This simulation-based policy optimization shall now be applied to highly relevant scenarios in power grid engineering.

Proc. of the Int. Workshop on Simulation for Energy, Sustainable Development & Environment 2013,
ISBN 978-88-97999-27-0; Bossomaier, Bruzzone, Cunha, Janosy, Longo, Eds.

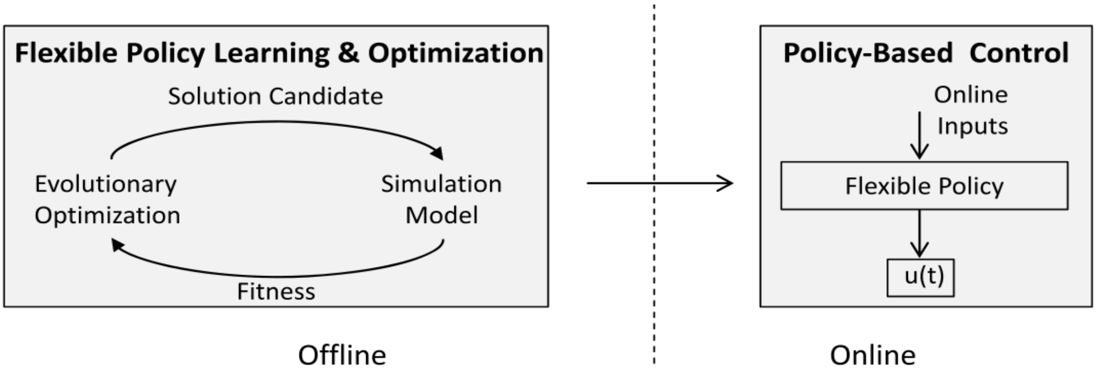

Figure 1: Principle of Policy Function Approximation

2.1. Formulating Simulation-Based Evolutionary Policy Function Approximation

As already discussed, policy function approximation is a method for dynamic optimization and seeks finding a generic function (policy) *P(x)* that returns a (near-optimal) action given a state. Most often, finding this function is the crucial step, where this technology usually gets applied when the structure of the policy is obvious (Powell, 2012). Equation 1 additionally aims at illustrating the principle of policy based control:

$$P = \arg\min_{P\in\Theta}[E(\sum_{t=1}^{K} F(P(J_t)))], \quad (1)$$

where J_t gives the system state at time *t, P* is the general denotation of a policy, $F(P(J_t))$ gives the fitness of a policy's action at time t and *E()* indicates that usually an estimate (over uncertain states) needs to be obtained. Since the aim is to find the best performing policy *P* (without knowing how this policy shall look like), it has to be optimized over the space of potential policies Θ.

Assuming that a sufficient description of a state is achievable (which serves as input to the aspired policy), the issue is now to find this policy. As already discussed and being illustrated by (Powell, 2012), policy-function approximation is especially attractive when the structure of the policy is obvious. However, what if we do not know about the structure a policy could take? What if we only know that a certain policy takes several input variables (that come from J_t) and combines them to a more complex analytic function (using manifold mathematical operators) in order to derive a control action? In this case, genetic programming (GP) provides a fruitful method for function approximation that does not need for a-priori knowledge on the aspired mathematical structure, but only knows about the input variables as well as a specific grammar for combining them. Applying a metaheuristic search process (genetic algorithms), GP finally searches for performant policies within Θ.

2.2. Function Approximation 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 trees, where leafs represent input variables describing the system state, that are combined by given mathematical operators which are incorporated by inner nodes. This kind of solution representation allows the evolution of arbitrary analytic functions without knowing their structure beforehand.

This approach has already been applied successfully to diverse applications, a general view on these works shall be provided herein.

3. APPLICATION TO THE GENERAL OPTIMAL POWER FLOW PROBLEM

The optimal power flow (OPF) problem has been stated some decades ago and is still the basic optimization problem in power system engineering. In its original formulation, the OPF considers steady-state situations, i.e. provides a static solution for exactly one considered discrete state. Since in volatile as well as uncertain power grids this consideration is some kind of "too optimistic", policy function approximation can be applied in order to build a general optimal power flow controller within simulated dynamic power grid environments. These simulations will be built based on real-world benchmark models, namely the IEEE distribution grid test cases[1]. For comparison reasons, the evolved approximate control policies will be evaluated with respect to exact interior point OPF solutions within steady-state situations created for testing reasons.

3.1. Policy Formulation

When striving to obtain a general control policy *P(x)*, it needs to be assumed that *x* gives a sufficient representation of the system's state J_t. In order to derive a control action for optimal power flows, *x* would need to consider all dependent variables of a power flow model (see definitions of OPF formulations in order to get an overview of used dependent variables (Hutterer, 2013c)). In such a case, especially for real-world

[1] Christie, R. D.: Power systems test case archive, http://www.ee.washington.edu/research/pstca/

systems a policy *P(x)* would need to consider many hundred or even thousands of input variables.

Therefore, the authors of this work proposed the introduction of abstract information entities, so called “abstract rules”, which extract information from a system’s dependent variables and provide only the necessary data to a decision making unit that is crucial for this decision. In the case of handling a traditional optimal power flow (OPF) problem, a set of 7 rules (r) has been proposed in (Hutterer, 2013a) that is assumed to be both necessary as well as sufficient for making power flow decisions. Thus, a policy *P(r)* needs to be approximated, where for most systems $|r|<<|x|$ holds.

3.2. Experiments

In the context of this publication, empirical studies have been performed for different models provided by the IEEE test case archive. For illustration reasons, two of these networks shall be presented herein, namely the 14-bus (the smallest test case) and the 300-bus (the largest test case) problems. Out of these benchmark instances, dynamic problems are built according to (Hutterer, 2013c) for learning policies for dynamic optimization. In order to validate the approximated policies, a test-procedure has been created that is based on randomly created test states. Therefore, within the simulation of the dynamic power grid models, discrete-time states are expressed that represent one single state of the system each. For these states, the deterministic OPF solution is computed with interior point method in MATPOWER and compared to actions that obtained best found policies lead to within these states. This allows the definition of an error term (in means of fitness), hence the quantitative validation of the policies’ performance for approximate dynamic OPF.

3.3. Results

Table 1 lists the quality in means of error between the OPF fitness function value of the best found policy compared to the deterministically computed optimal solution within 10 arbitrary and independent test states. For both benchmark instances, it can be shown that the approximate policy-based control leads to near-optimal decisions, that are in mean only 0.6 respectively 0.51 % worse than the reference solutions.

Table 1: Performance Validation of OPF Policies

Time Step	**14-Bus**	**300-Bus**
1	0.0074	0.0160
2	0.0053	0.0201
3	0.0105	-0.0151
4	0.0045	0.0033
5	0.0041	0.0030
6	0.0071	0.0046
7	0.0071	0.0048
8	0.0064	0.0033
9	0.0034	0.0030
10	0.0042	0.0082
Mean Relative Error (MRE)	**0.0060**	**0.0051**

A more detailed discussion on OPF policy approximation can be obtained from (Hutterer, 2013c). This illustration shall demonstrate the power of policy function approximation for dynamic optimization in the context of the OPF-based generation unit scheduling and quantitatively shows its validity.

Further empirical studies applied policy-based control to demand side management; more detailed: to the control of distributed electric vehicle charging processes.

4. APPLICATION TO SMART ELECTRIC VEHICLE CHARGING CONTROL

Various researchers examine the problem of integrating plug-in electric vehicles (EV) optimally into power grids, where the control of charging power is seen as advantageous for reaching optimal power grid operation (Clement, 2008; Clement, 2009; Sortomme, 2011). Central challenge beside formulation and computation of the optimization problem itself is the consideration of the individual behavior that mainly characterizes electric vehicle charging load.

This PEV-charging control problem represents a dynamic optimization task that requires optimal control actions for high amounts of distributed EV-agents. Therefore, building a policy-based approach is the fundament of this show case. Here, each agent (EV) receives a flexible policy rather than static control decisions that makes it react to its environment dynamically, but in a globally optimal manner. This policy is principally the same for all agents, but using individual data from an agent’s environment, it leads to agent-specific charging control. The policies will be evolved using the presented policy function approximation approach. Here, a simulation model is built that represents a fleet of EVs within a given power grid area, which will be integrated into the simulation-based evolutionary optimization of PEV charging control policies.

4.1. Policy Formulation

A general policy shall be obtained. Its aim is to output a charging decision to an electric vehicle given the state of the system at a given time step. This optimized policy is derived from input variables that consider agent-specific parameters from its environment. Out of these parameters, the EV’s power demand as well as the state of its environment can be described sufficiently in order to derive a valid charging decision. Here, three different data classes can be distinguished from each other:

- Agent-specific data 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 data considers 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 data considers 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.

Out of these classes of input data, in the context of this work, the authors defined once more a set of abstract rules *(r)* that gathers all needed information for decision making and provides compressed information to control units (EVs in this case). Out of these rules, the general policy *P(r)* shall be learned that gives the decision on the charging power to a certain EV in the system.

4.2. Experiments

All experiments are based on the IEEE 33-bus distribution feeder. Within this feeder, 300 EVs are simulated to act individually. Additionally, renewable sources (wind-power and photovoltaics) are added to the system in order to create a dynamic and volatile scenario. A finally obtained charging policy has to derive robust charging decisions that provide system-wide near-optimal charging control over time.

Detailed discussions on this problem scenario as well as formal definitions are provided in (Hutterer, 2013b). Here, only the main issue of applying policy-function approximation to suchlike problems shall be illustrated.

4.3. Results

In this scenario, a general policy $P_{EV}(r)$ needs to be computed that derives accurate control actions for each single EV in the system, while considering both the agent's (EV) needs as well as system-wide goals of power grid operation. Equation 1 illustrates such a policy, which has been found in studies on the above mentioned system.

$$\begin{aligned} P_{EV} &= \\ c_1 &* ERT^{10} * AP * AWS * MCR * (PBL - c_2) \quad (1) \\ &+ c_3 * AI \end{aligned}$$

While the used input variables are discussed extensively in (Hutterer, 2012; Hutterer, 2013b), this policy considers an EVs remain time at a charging spot (ERT), the actual electricity price (AP), the actual wind speed (AWS), the mean charging rate of all other EVs during the previous time step (MCR), the past base load (PBL) at the previous time step as well as the actual solar irradiance (AI). Since all input rules are defined to give a value in the range [0,1], with the constants c1=17.86, c2=0.03 and c3=0.11 the policy finally provides a charging power value in the unit [kW].

Figure 2 depicts the mean charging power over all 300 EVs in the simulated system when applying the evolved policy to a test-scenario. From this illustration one can observe the principal functionality that this policy causes during operation, namely the principal shift of charging to time steps at night (where the grid-load is low). While this graph gives the mean charging power, the actual power of each EV differs and considers its individual behavior. However, over all EVs the system-wide constraints are satisfied that enable secure power grid operation, while the objective function considering total costs of energy supply is minimized.

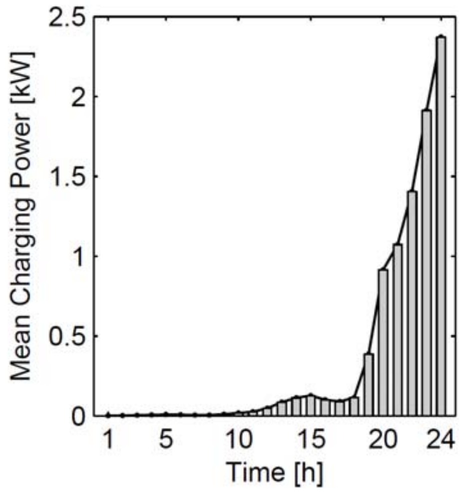

Figure 2: Charging Power in Simulated Test Scenario

However, this example shall demonstrate the evolution of control policies for dynamic optimization to an actual smart-grid relevant problem domain. More detailed discussions can be obtained from the referenced works, while this paper aims at providing a general view on the developed techniques. An outlook shall now depict several open issues.

5. OUTLOOK

Dynamic optimization with policies has the great advantage that it avoids the necessity of computing a specific solution to each state the dynamic system exhibits over time. Hence, dynamic adaptation of solutions seems to be not necessary, which is a major challenge in dynamic evolutionary optimization (Nguyen 2013a, Nguyen 2013b). However, this advantage only holds in a restrictive manner: A policy is able to make accurate decisions within situations that are sufficiently similar to those situations it has been trained to. For other situations, its extrapolation-ability is necessary to still make good decisions. As soon as specific situations are too different from the training simulation, obviously the policy-based control becomes useless. In such a case, the simulation model would need to be adapted in order to correspond to such situations, and the policy would need to be adapted / relearned.

Hence, if being learned accurately, policy-based control is valid for systems where their behavior, dynamics and uncertainty are adequately predictable. If such a system changes over long time, and the simulation no more matches the real-world, special techniques will need to be applied for learning a new policy or adapting the existing one, where numerous approaches already exist in literature, using memories of already evaluated solutions, sub-populations or immigration methodologies in order to adapt the evolutionary search to changing positions of the desired optima over time (Yang 2013). Hence, solution adaptation is avoided on a short-time scale, where the policy is able to derive decisions for uncertain and dynamic states. On a long-time scale, solution adaptation is still necessary in order to meet potential drifts of the system (i.e. a mismatch between simulation model and real-world).

Future work will need to concentrate on the examination of such drifts and needs to apply methodologies for policy adaptation.

5.1. Conclusion

This paper illustrated the application of policy function approximation for the sake of dynamic optimization under uncertainty in power grids. Summarizing related work from the authors and stating new results, two empirical studies have been outlined that show the application to central problem classes in power grid optimization. While policy-function approximation seems to be a fruitful technology for dynamic optimization, open issues have been identified that challenge new research questions in the context of dynamic systems.

ACKNOWLEDGMENTS

The work described in this paper was done within the Josef Ressel Centre for Heuristic Optimization Heureka! (http://heureka.heuristiclab.com/) sponsored by the Austrian Research Promotion Agency (FFG).

REFERENCES

Clement, K., Haesen, E., and Driesen, J. (2008): *The Impact of Uncontrolled and Controlled Charging of Plug-In Hybrid Electric Vehicles on the Distribution Grid*, Proceedings of the European Ele-Drive Conference, (EET-2008).

Clement, K., Haesen, E., and 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.

Hutterer, S., Affenzeller, M., and Auinger, F., (2012). *Evolutionary Optimization of Multi-Agent Control Strategies for Electric Vehicle Charging*. Companion Publication of the Genetic and Evolutionary Computation Conference, GECCO'12, Philadelphia, PA

Hutterer, S., Affenzeller, M., and Auinger, F., (2013). *Evolutionary Algorithm Based Control Policies for Flexible Optimal Power Flow Over Time*, Proceedings of the EvoStar – EvoEnergy Workshop Workshop, Vienna, Austria

Hutterer, S., Auinger, F., and Affenzeller, M., (2013). *Evolutionary Computation Enabled Controlled Charging for E-Mobility Aggregators*, IEEE CIS Symposium Series on Computational Intelligence, Workshop: CIASG, Singapore

Hutterer, S. and Affenzeller, M., (2013). *Genetic Programming Enabled Evolution of Control Policies for Dynamic Stochastic Optimal Power Flow*, Companion Publication of the Genetic and Evolutionary Computation Conference, GECCO'13, Amsterdam, Netherlands

Nguyen, T.T, Yang, S., and Branke, J. (2012). Evolutionary dynamic optimization: A survey of the state of the art. *Swarm & Evol. Comput.* 6, pp. 1–24.

Nguyen, T. T. and Yao X. (2012). Continuous Dynamic Constrained Optimization-The Challenges. *IEEE Transactions on Evolutionary Computation*, 16(6), pp. 769-786.

Powell, W. B., Simao, H. P., and Bouzaiene-Ayari, B. (2012). Approximate dynamic programming in transportation and logistics: A unified framework. *European Journal of Transportation and Logistics, 1(3),* pp. 237-284.

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, 1(1).*

Wang, H., Murillo-Sanchez, C. E., Zimmerman, R. D., and Thomas, R. J. (2007): On computational issues of market-based optimal power flow. *IEEE Transactions on Power* Systems, 22(3).

Yang, S. and Yao, X. (2013). *Evolutionary Computation for Dynamic Optimization Problems*. Studies in Computational Intelligence Vol. 490, Springer Berlin Heidelberg

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.