SPRINGER BRIEFS IN OPTIMIZATION

Slawomir Koziel
Stanislav Ogurtsov

Antenna Design by Simulation-Driven Optimization

Springer

SpringerBriefs in Optimization

SpringerBriefs in Optimization showcases algorithmic and theoretical techniques, case studies, and applications within the broad-based field of optimization. Manuscripts related to the ever-growing applications of optimization in applied mathematics, engineering, medicine, economics, and other applied sciences are encouraged.

For further volumes:
http://www.springer.com/series/8918

Slawomir Koziel • Stanislav Ogurtsov

Antenna Design by Simulation-Driven Optimization

Springer

Slawomir Koziel
School of Science and Engineering
Reykjavik University Engineering
Optimization & Modeling Center
Reykjavik, Iceland

Stanislav Ogurtsov
School of Science and Engineering
Reykjavik University Engineering
Optimization & Modeling Center
Reykjavik, Iceland

ISSN 2190-8354 ISSN 2191-575X (electronic)
ISBN 978-3-319-04366-1 ISBN 978-3-319-04367-8 (eBook)
DOI 10.1007/978-3-319-04367-8
Springer Cham Heidelberg New York Dordrecht London

Library of Congress Control Number: 2013958383

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

Design of contemporary antenna structures heavily relies on electromagnetic (EM) simulations. Accurate reflection and radiation responses of many antenna geometries can be obtained only with discrete full-wave EM simulation. On the other hand, the direct use of high-fidelity EM simulation in the design process, particularly for automated parameter optimization, results in high computational costs, often prohibitive. Other issues, such as the presence of numerical noise, may result in a failure of optimization using conventional (e.g., gradient-based) methods. In this book, we demonstrate that numerically efficient design of antennas can be realized using surrogate-based optimization (SBO) methodologies. The essence of SBO techniques resides in shifting the optimization burden to a fast surrogate of the antenna structure and the use of coarse-discretization EM models to configure the surrogate. A properly created and handled surrogate serves as a reliable prediction tool so that satisfactory designs can be found at the costs of a limited number of simulations of the high-fidelity EM antenna model. The specific SBO techniques covered here include space mapping combined with response surface approximation, shape-preserving response prediction (SPRP), adaptive response correction (ARC), adaptively adjusted design specification (AADS), variable-fidelity simulation-driven optimization (VFSDO), and surrogate-based optimization enhanced by the use of adjoint sensitivities. Multi-objective design of antennas is also covered to some extent. Moreover, we discuss practical issues such as the effect of the coarse-discretization model fidelity on the final design quality and the computational cost of the optimization process. Our considerations are illustrated using numerous application examples. Recommendations concerning application of specific SBO techniques to antenna design are also presented.

Contents

Chapter 1
Introduction

Design of modern antennas is undoubtedly a challenging task. An important part of the design process is the adjustment of geometry and material parameters to ensure that the antenna response satisfies prescribed performance specifications with respect to certain characteristics such as input impedance, radiation pattern, antenna efficiency, etc. (Volakis 2007; Schantz 2005; Petosa 2007; Balanis 2005). In this context, computationally inexpensive analytical models can only be used—in most cases—to obtain an initial estimate of the optimum design. This is particularly the case when certain interactions within the antenna itself and with the antenna environment (e.g., housing, installation fixture, feeding circuit, connectors) have to be taken into account. For these reasons, full-wave electromagnetic (EM) simulation plays an essential role in the design closure. Contemporary computational techniques—implemented in commercial simulation packages—are capable to adequately evaluate antenna reflection and radiation responses. On the other hand, full-wave simulations of realistic and finely discretized antenna models are computationally expensive: evaluation for a single combination of design parameters may take up to several hours. While this cost is acceptable from the design validation standpoint, it is usually prohibitive for design optimization that normally requires a large number of EM simulations of the antenna structure of interest.

Automation of the antenna design process can be realized by formulating the antenna parameter adjustment task as an optimization problem with the objective function supplied by an EM solver (Special issue, IEEE APS 2007). Unfortunately, most of the conventional optimization techniques, including gradient-based (Nocedal and Wright 2000), e.g., conjugate-gradient, quasi-Newton, sequential quadratic programming, etc., and derivative-free (Kolda et al. 2003) methods, e.g., Nelder-Mead and pattern search techniques, require a large number of objective function evaluations to converge to a satisfactory design. For many realistic EM antenna models, where evaluation time per design reaches a few hours with contemporary computing facilities, the cost of such an optimization process may be unacceptably high. Another practical problem of conventional optimization techniques is numerical noise, which is partially a result of adaptive meshing techniques used

S. Koziel and S. Ogurtsov, *Antenna Design by Simulation-Driven Optimization*,
SpringerBriefs in Optimization, DOI 10.1007/978-3-319-04367-8_1,

by most contemporary EM solvers: even a small change of design variables may result in a change of the mesh topology and, consequently, discontinuity of the EM-simulated antenna responses as a function of designable parameters. The noise is particularly an issue for gradient-based methods that normally require smoothness of the objective function.

The aforementioned challenges result in a situation where the most common approach to simulation-driven antenna design is based on repetitive parameter sweeps (usually, one parameter at a time). This approach is usually more reliable than a brute-force optimization using built-in optimization capabilities of commercial simulation tools; however, it is also very laborious, time-consuming, and demanding significant designer supervision. Moreover, such a parameter-sweep-based optimization process does not guarantee optimum results because only a limited number of parameters can be handled that way. It is also difficult to utilize correlations between the parameters properly. Finally, optimal values of the designable variables can be quite counterintuitive.

In recent years, population-based search methods (also referred to as metaheuristics) (Yang 2010) have gained considerable popularity. This group of methods includes, among others, genetic algorithms (GA) (Back et al. 2000), particle swarm optimizers (PSO) (Kennedy 1997), differential evolution (DE) (Storn and Price 1997), and ant colony optimization (Dorigo and Gambardella 1997). Most of metaheuristics are biologically inspired systems designed to alleviate certain difficulties of the conventional optimization methods, in particular, handling problems with multiple local optima (Yang 2010). Probably the most successful application of the metaheuristic algorithms in antenna design resided so far in array optimization problems (e.g., Ares-Pena et al. 1999; Haupt 2007; Jin and Rahmat-Samii 2007, 2008; Petko and Werner 2007; Bevelacqua and Balanis 2007; Grimaccia et al. 2007; Pantoja et al. 2007; Selleri et al. 2008; Li et al. 2008; Rajo-Iglesias and Quevedo-Teruel 2007; Roy et al. 2011). In these problems, the cost of evaluating the single element response is not of the primary concern or the response of a single element is already available, e.g., with a preassigned array element. However, application of metaheuristics to EM-simulation-driven antenna design is not practical because corresponding computational costs would be tremendous: typical GA, PSO, or DE algorithm requires hundreds, thousands, or even tens of thousands of objective function evaluations to yield a solution (Ares-Pena et al. 1999; Haupt 2007; Jin and Rahmat-Samii 2007, 2008; Petko and Werner 2007; Bevelacqua and Balanis 2007; Grimaccia et al. 2007; Pantoja et al. 2007; Selleri et al. 2008; Li et al. 2008; Rajo-Iglesias and Quevedo-Teruel 2007; Roy et al. 2011).

The problem of high computational cost of conventional EM-based antenna optimization can be alleviated to some extent by the use of adjoint sensitivity (Director and Rohrer 1969), which is a computationally cheap way to obtain derivatives of the system response with respect to its design parameters. Adjoint sensitivities can substantially speed up microwave design optimization while using gradient-based algorithms (Bandler and Seviora 1972; Chung et al. 2001). This technology was also demonstrated for antenna optimization (Jacobson and Rylander 2010; Toivanen et al. 2009; Zhang et al. 2012). It should be mentioned, however, that adjoint sensitivities are not yet widespread in commercial EM solvers. Only CST (CST Microwave Studio 2011) and HFSS (HFSS 2010) have recently implemented this feature. Also, the use of adjoint sensitivities is limited by numerical noise of the EM-simulated response (Koziel et al. 2012c).

One of the most recent and yet most promising ways to realize computationally efficient simulation-driven antenna design is surrogate-based optimization (SBO) (Koziel and Ogurtsov 2011a, b; Forrester and Keane 2009; Queipo et al. 2005). The SBO main idea is to shift the computational burden of the optimization process to a surrogate model which is a cheap representation of the optimized antenna (Bandler et al. 2004a, b; Queipo et al. 2005; Koziel et al. 2006; Koziel and Ogurtsov 2011a). In a typical setting, the surrogate model is used as a prediction tool to find approximate location of the original (high-fidelity or fine) antenna model. After evaluating the high-fidelity model at this predicted optimum, the surrogate is updated in order to improve its local accuracy (Koziel et al. 2011c). The key prerequisite of the SBO paradigm is that the surrogate is much faster than the high-fidelity model. Also, in many SBO algorithms, the high-fidelity model is only evaluated once per iteration. Therefore, the computational cost of the design process with a well working SBO algorithm may be significantly lower than those with most of conventional optimization methods.

There are two major types of surrogate models. The first one comprises function-approximation models constructed from sampled high-fidelity simulation data (Simpson et al. 2001). A number of approximation (and interpolation) techniques are available, including artificial neural networks (Haykin 1998), radial basis functions (Gutmann 2001; Wild et al. 2008), kriging (Forrester and Keane 2009), support vector machines (Smola and Schölkopf 2004), Gaussian process regression (Angiulli et al. 2007; Jacobs 2012), or multidimensional rational approximation (Shaker et al. 2009). If the design space is sampled with sufficient density, the resulting model becomes reliable so that the optimal antenna design can be found just by optimizing the surrogate. In fact, approximation methods are usually used to create multiple-use library models of specific components. The computational overhead related to such models may be very high. Depending on the number of designable parameters, the number of training samples necessary to ensure decent accuracy might be hundreds, thousands, or even tens of thousands. Moreover, the number of samples quickly grows with the dimensionality of the problem (so-called curse of dimensionality). As a consequence, globally accurate approximation modeling is not suitable for ad hoc (one-time) antenna optimization. Iteratively improved approximation surrogates are becoming popular for global optimization (Couckuyt 2013). Various ways of incorporating new training points into the model (so-called infill criteria) exist, including exploitative models (i.e., models oriented toward improving the design in the vicinity of the current one), explorative models (i.e., models aiming at improving global accuracy), as well as model with balanced exploration and exploitation (Jones et al. 1998; Forrester and Keane 2009).

Another type of surrogates, so-called physics-based surrogates, is constructed from underlying low-fidelity (or coarse) models or the respective structures. Because the low-fidelity models inherit some knowledge of the system under consideration, usually a small number of high-fidelity simulations are sufficient to configure a reliable surrogate. The most popular SBO approaches using physics-based surrogates that proved to be successful in microwave engineering are space mapping (SM) (Bandler et al. 2004a, b), tuning, tuning SM (Koziel et al. 2009a; Cheng et al. 2010), as well as various response correction methods (Echeverria and Hemker 2005; Koziel et al. 2009b; Koziel 2010a). To ensure computational efficiency, it is important that the low-fidelity model is considerably faster than the high-fidelity model. For that reason, circuit equivalents or models based on analytical formulas are preferred (Bandler et al. 2004a, b). The aforementioned methods (particularly space mapping) were mostly used to design filters or transmission-line-based components such as impedance transformers (Amari et al. 2006; Wu et al. 2004; Bandler et al. 2004a, b). Unfortunately, in case of antennas, reliable circuit equivalents are rarely available. For antennas, a universal way of obtaining low-fidelity models is through coarse-discretization EM simulations. Such models are relatively expensive, which poses additional challenges in terms of optimization.

The topic of this book is surrogate-based optimization methods for simulation-driven antenna design with the focus on surrogate-based techniques exploiting variable-fidelity EM simulations and physics-based surrogates. We begin, in Chap. 2, by formulating the antenna design task as an optimization problem. We also briefly discuss conventional numerical optimization techniques, including both gradient-based and derivative-free methods but also metaheuristics. In Chap. 3, surrogate-based optimization is introduced. In the same chapter, the SBO design workflow as well as various aspects of surrogate-based optimization is presented on a generic level. Chapter 4 is an exposition of the specific state-of-the-art physics-based SBO techniques that are suitable for antenna design optimization. The emphasis is put on methods that aim at minimizing the number of both high- and low-fidelity EM simulations of the antenna under design and thus reducing the overall design cost. Chapters 6–9 present applications of the methods discussed in Chap. 4 for the design of specific antenna structures. Variable-fidelity design exploiting adjoint sensitivity is presented in Chap. 10. Chapter 11 discusses multi-objective antenna design using surrogate models. Chapter 12 provides a discussion of open issues related to SBO antenna design with special focus on selecting simulation model fidelity and its impact on the performance and computational efficiency of the optimization process. The book is concluded in Chap. 13. Here, we formulate recommendations for the readers interested in applying presented algorithm and techniques in their antenna design and discuss possible future developments concerning mostly automation of simulation-driven antenna design.

Chapter 2
Antenna Design Using Electromagnetic Simulations

In this chapter, we formulate the antenna design task as a nonlinear minimization problem. We introduce necessary notation, discuss typical objectives and constraints, and give a brief overview of conventional optimization techniques, including gradient-based and derivative-free methods, as well as metaheuristics. We also introduce the concept of the surrogate-based optimization (SBO) and discuss it on a generic level. More detailed exposition of SBO and SBO-related design techniques will be given in Chaps. 3 and 4.

2.1 Antenna Design Task as an Optimization Problem

Let $\boldsymbol{R}_f(\boldsymbol{x})$ denote a response of a high-fidelity (or fine) model of the antenna under design. For the rest of this book, we will assume that $\boldsymbol{R}_f$ is obtained using accurate full-wave electromagnetic (EM) simulation. Typically, $\boldsymbol{R}_f$ will represent evaluation of performance characteristics, e.g., reflection $|S_{11}|$ or gain over certain frequency band of interest. Vector $\boldsymbol{x} = [x_1\ x_2 \ldots x_n]^T$ represents designable parameters to be adjusted (e.g., geometry and/or material ones).

In some situations, individual components of the vector $\boldsymbol{R}_f(\boldsymbol{x})$ will be considered, and we will use the notation $\boldsymbol{R}_f(\boldsymbol{x}) = [R_f(\boldsymbol{x}, f_1)\ R_f(\boldsymbol{x}, f_2)\ \ldots\ R_f(\boldsymbol{x}, f_m)]^T$, where $R_f(\boldsymbol{x}, f_k)$ is the evaluation of the high-fidelity model at a frequency f_k, whereas f_1 through f_m represent the entire discrete set of frequencies at which the model is evaluated.

The antenna design task can be formulated as the following nonlinear minimization problem (Koziel and Ogurtsov 2011a):

$$\boldsymbol{x}^* = \arg\min_{\boldsymbol{x}} U\left(\boldsymbol{R}_f\left(\boldsymbol{x}\right)\right) \tag{2.1}$$

where U is the scalar merit function encoding the design specifications, whereas $\boldsymbol{x}^*$ is the optimum design to be found. The composition $U(\boldsymbol{R}_f(\boldsymbol{x}))$ will be referred to as the objective function. The function U is implemented so that a better design $\boldsymbol{x}$ corresponds to a smaller value of $U(\boldsymbol{R}_f(\boldsymbol{x}))$. In antenna design, U is most often

S. Koziel and S. Ogurtsov, *Antenna Design by Simulation-Driven Optimization*, SpringerBriefs in Optimization, DOI 10.1007/978-3-319-04367-8_2,

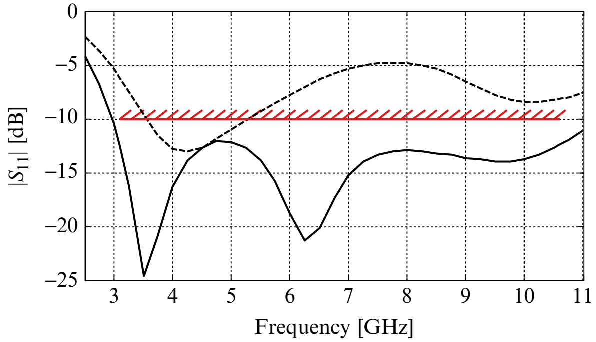

Fig. 2.1 Illustration of minimax design specifications, here, $|S_{11}| \leq -10$ dB for 3.1–10.6 GHz, marked with *thick horizontal line*. An example UWB antenna reflection response that does not satisfy our specifications (*dashed line*) (specification error, i.e., maximum violation of the requirements is about +5 dB) and another response that does satisfy the specifications (*solid line*)

implemented as a minimax function with upper (and/or lower) specifications. Figure 2.1 shows the example of minimax specifications corresponding to typical UWB requirements for the reflection response, i.e., $|S_{11}| \leq -10$ dB for 3.1–10.6 GHz. The value of $U(\boldsymbol{R}_f(\boldsymbol{x}))$ (also referred to as minimax specification error) corresponds to the maximum violation of the design specifications within the frequency band of interest.

To simplify notation, we will occasionally use the symbol $f(\boldsymbol{x})$ as an abbreviation for $U(\boldsymbol{R}_f(\boldsymbol{x}))$.

In reality, the problem (2.1) is always constrained. The following types of constraints can be considered:

- Lower and upper bounds for design variables, i.e., $l_k \leq x_k \leq u_k$, $k = 1, \ldots, n$
- Inequality constraints, i.e., $c_{\text{ineq}\cdot l}(\boldsymbol{x}) \leq 0$, $l = 1, \ldots, N_{\text{ineq}}$, where N_{ineq} is the number of constraints
- Equality constraints, i.e., $c_{\text{eq}\cdot l}(\boldsymbol{x}) = 0$, $l = 1, \ldots, N_{\text{eq}}$, where N_{eq} is the number of constraints

Design constraints are usually introduced to make sure that the antenna structure that is to be evaluated by the EM solver is physically valid (e.g., certain parts of the structure do not overlap). Also, constraints can be introduced in order to ensure that the physical dimensions (length, width, area) of the antenna do not exceed certain assumed values.

In this book, geometry constraints such as those described above are handled explicitly. Other types of constraints, particularly those that emerge due to converting initially multi-objective design problem into single-objective one, are handled through penalty functions. It should be mentioned though that the literature offers efficient ways of explicit handling expensive constraints; see, e.g., Kazemi et al. (2011), Basudhar et al. (2012).

Figure 2.2 shows the simulation-driven design optimization flowchart. Typically, it is an iterative process where the designs found by the optimizer are verified by

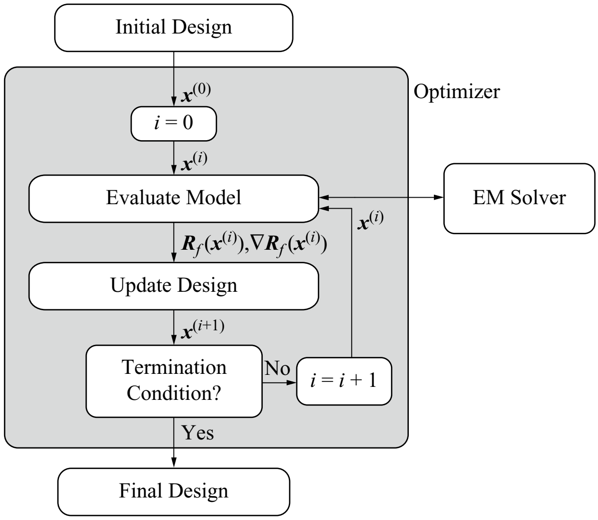

Fig. 2.2 Simulation-driven design through optimization. Generic optimization scheme is an iterative process where the new candidate designs are generated by the optimization algorithm and the high-fidelity model is evaluated through EM simulation for verification purposes and to provide the optimizer with information that allows searching for possible better designs. Depending on the type of the algorithm, the search process may be guided by the model response or (if available) by the response and its derivatives (gradient)

evaluating the high-fidelity model in the EM solver and—depending on a particular algorithm—the search process is guided either by the model response itself or the response of its gradients (if available). In Sects. 2.2–2.5, we briefly discuss conventional optimization approaches. In Chaps. 3 and 4, we discuss surrogate-based optimization methods, which are the main topic of this book.

2.2 Gradient-Based Optimization Methods

Gradient-based optimization techniques are the oldest and the most popular optimization methods (Nocedal and Wright 2000). In order to proceed toward the optimum design, they utilize derivative information of the objective function. Assuming that the objective $f(\boldsymbol{x})$ is sufficiently smooth (i.e., at least continuously differentiable), the gradient $\nabla f = [\partial f/\partial x_1\ \partial f/\partial x_2 \ldots \partial f/\partial x_n]^T$ gives the information about descent of f in the vicinity of the design at which the gradient is calculated. More specifically,

$$f(\boldsymbol{x}+\boldsymbol{h}) \cong f(\boldsymbol{x}) + \nabla f(\boldsymbol{x}) \cdot \boldsymbol{h} < f(\boldsymbol{x}) \tag{2.2}$$

for sufficiently small $\boldsymbol{h}$ as long as $\nabla f(\boldsymbol{x}) \cdot \boldsymbol{h} < 0$. In particular $\boldsymbol{h} = -\nabla f(\boldsymbol{x})$ determines the direction of the steepest descent. In practice, using steepest descent as a search direction results in a poor performance of the optimization algorithm (Nocedal and

Wright 2000; Yang 2010). Better results are obtained using so-called conjugate-gradient method where the search direction is determined as a combination of the previous direction $\boldsymbol{h}_{\text{prev}}$ and the current gradient, i.e.,

$$\boldsymbol{h} = -\nabla f\left(\boldsymbol{x}^i\right) + \gamma \boldsymbol{h}_{prev} \tag{2.3}$$

An example way of selecting the coefficient γ is a Fletcher-Reeves method with

$$\gamma = \frac{\nabla f\left(\boldsymbol{x}\right)^{\mathrm{T}} \nabla f\left(\boldsymbol{x}\right)}{\nabla f\left(\boldsymbol{x}_{prev}\right)^{\mathrm{T}} \nabla f\left(\boldsymbol{x}_{prev}\right)} \tag{2.4}$$

Having the search direction, the next design $\boldsymbol{x}^{i+1}$ is determined from the current one $\boldsymbol{x}^i$ as

$$\boldsymbol{x}^{i+1} = \boldsymbol{x}^i + \alpha \cdot \boldsymbol{h} \tag{2.5}$$

Here, the choice of the step size $\alpha > 0$ is of great importance (Nocedal and Wright 2000), and finding it is referred to as a line search.

It is also possible to utilize second-order derivative information, which is characteristic to so-called Newton methods. Assuming f is at least twice continuously differentiable, one can consider a second-order Taylor expansion of f:

$$f\left(\boldsymbol{x}+\boldsymbol{h}\right) \cong f\left(\boldsymbol{x}\right) + \nabla f\left(\boldsymbol{x}\right) \cdot \boldsymbol{h} + \frac{1}{2} \boldsymbol{h} \cdot \boldsymbol{H}\left(\boldsymbol{x}\right) \cdot \boldsymbol{h} \tag{2.6}$$

where $\boldsymbol{H}(\boldsymbol{x})$ is the Hessian of f at $\boldsymbol{x}$, i.e., $\boldsymbol{H}(\boldsymbol{x}) = [\partial^2 f/\partial x_j \partial x_k]_{j,k=1,\ldots,n}$. This means, given the current design $\boldsymbol{x}^i$ being sufficiently close to the minimum of f, that the next approximation of the optimum can be determined as

$$\boldsymbol{x}^{i+1} = \boldsymbol{x}^i - \left[\boldsymbol{H}\left(\boldsymbol{x}\right)\right]^{-1} \nabla f\left(\boldsymbol{x}\right) \tag{2.7}$$

If the starting point is sufficiently close to the optimum and the Hessian is positive definite (Yang 2010), the algorithm (2.7) converges very quickly to the (locally) optimal design. In practice, neither of these conditions is usually satisfied, so various types of damped Newton techniques are used, e.g., Levenberg-Marquardt method (Nocedal and Wright 2000). On the other hand, the Hessian of the objective function f is normally not available so quasi-Newton methods are used instead where the Hessian is approximated using various updating formulas (Nocedal and Wright 2000).

From the point of view of simulation-driven antenna design, the use of gradient-based methods is problematic mostly because of the high computational cost of accurate simulation and the fact that gradient-based methods normally require large number of objective function evaluations to converge, unless cheap way of obtaining sensitivity is utilized (e.g., through adjoints or automatic differentiation). Another problem is numerical noise that is always present in EM-based objective functions. Recently, adjoint sensitivity techniques have become available in some commercial EM solvers (CST 2013; HFSS 2010), which may revive the interest in this type of methods for antenna design because they allow calculation of sensitivity at little or no extra cost compared to a regular EM simulation of the antenna structure. On the other hand, automatic differentiation is usually not an option because source codes are not accessible whenever commercial solvers are utilized. It should also be mentioned that gradient-based methods exploiting a trust-region framework are usually more efficient than those based, e.g., on line search so that using trust region (Conn et al. 2000) is recommended whenever possible.

2.3 Derivative-Free Optimization Methods

In many situations, gradient-based search may not be a good option. This is particularly the case when derivative information is not available or expensive to compute (e.g., through finite differentiation of an expensive objective function). Also, if the objective function is noisy (which is typical for responses obtained from EM simulation) then the gradient-based search does not perform well.

Optimization techniques that do not use derivative data in the search process are referred to as derivative-free methods. Formally speaking metaheuristics (Sect. 2.4) as well as many surrogate-based approaches (Chaps. 3 and 4) also fall into this category. In this section, however, we only mention the basic idea of the local search methods. Figure 2.3 illustrates the concept of the pattern search (Kolda et al. 2003), where the search of the objective function minimum is restricted to the rectangular grid and explores a grid-restricted vicinity of the current design. Failure in making the step improve the current design leads to refining the grid size and allowing

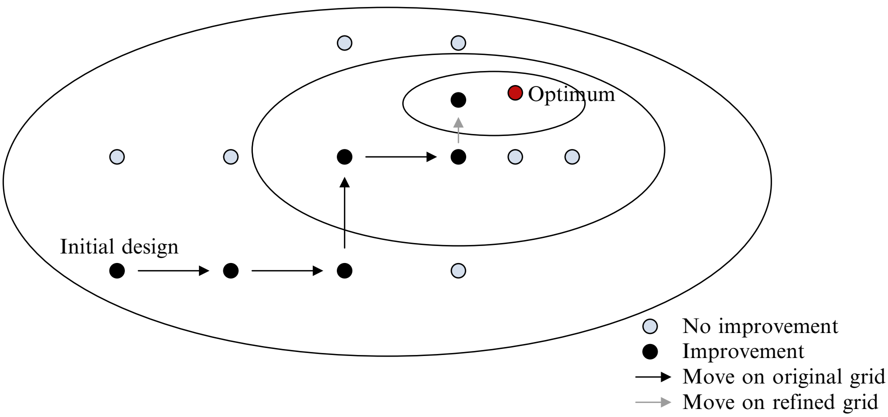

Fig. 2.3 The concept of pattern search. The search is based on exploratory movements restricted to the rectangular grid around the initial design. Upon failure of making the successful move, the grid is refined to allow smaller steps. The actual implementations of pattern search routines also use more sophisticated strategies (e.g., grid-restricted line search)

smaller steps. Various variants of the pattern search methods are available (see, e.g., Torczon 1997; Kolda et al. 2003). With sufficiently large size of the initial grid, these techniques can be used to perform a quasi-global search.

One of the most famous derivative-free methods is the Nelder-Mead algorithm (Yang 2010) also referred to as the simplex method. Its search process is based on moving the vertices of the simplex in the design space in such a way that the vertex corresponding to the worst (i.e., highest) value of the objective function is replaced by the new one at the location where the objective function value is expected to be improved.

Pattern search and similar methods are usually robust although their convergence is relatively slow compared to gradient-based routines. Their fundamental advantage is in the fact that they do not use derivative information and, even more importantly, they are quite immune to the numerical noise. An excellent and mathematically rigorous treatment of derivative-free optimization techniques, including model-based trust-region derivative-free methods, can be found in Conn et al. (2009). Many pattern search methods and their extensions possess mathematically rigorous convergence guarantees (Conn et al. 2009). An interesting extension of pattern search to constrained black-box optimization is Mesh Adaptive Direct Search (MADS) (Audet and Dennis 2006).

2.4 Metaheuristics and Global Optimization

Metaheuristics are global search methods that are based on observation of natural processes (e.g., biological or social systems). Most metaheuristics process sets (or populations) of potential solutions to the optimization problem at hand in a way that these solutions (also called individuals) interact with each other so that the optimization process is capable to avoid getting stuck in local optima and converge—with reasonable probability—to a globally optimal solution of the problem. At the same time, metaheuristics can handle noisy, non-differentiable, and discontinuous objective functions.

The most popular types of metaheuristic algorithms include genetic algorithms (GAs) (Goldberg 1989), evolutionary algorithms (EAs) (Back et al. 2000), evolution strategies (ES) (Back et al. 2000), particle swarm optimizers (PSO) (Kennedy et al. 2001), differential evolution (DE) (Storn and Price 1997), and, more recently, firefly algorithm (Yang 2010). A famous example of metaheuristic algorithm that processes a single solution rather than a population of individuals is simulated annealing (Kirkpatrick et al. 1983).

The typical flow of the metaheuristic algorithm is the following (here, P is the set of potential solutions to the problem at hand, also referred to as a population):

1. Initialize population P.
2. Evaluate population P.
3. Select parent individuals S from P.