

Multi-Objective Groundwater Management Using Evolutionary Algorithms

Tobias Siegfried, Stefan Bleuler, Marco Laumanns, Eckart Zitzler and Wolfgang Kinzelbach

Abstract

Sustainable management of groundwater resources is of crucial importance for regions where freshwater supply is naturally limited. In the light of a long-term planning of groundwater usage, computer-based decision support tools are necessary that on the hand are able to predict the complex system dynamics with sufficient accuracy and on the other hand allow to explore the space of usage scenarios with respect to different criteria such as sustainability, cost, etc. In this paper, we present a multi-objective evolutionary algorithm for groundwater management that optimizes the placement and the operation of pumping facilities over time, while considering multiple neighboring regions which are economically independent. The algorithm helps to investigate the cost trade-offs between the different regions by providing an approximation of the Pareto-optimal set which is demonstrated on a three-region problem. The application of the proposed methodology can also serve as a benchmark application as shown in the paper. The corresponding implementation is freely available as a precompiled module at <http://www.tik.ee.ethz.ch/pisa>.

I. INTRODUCTION

In many areas of limited freshwater resources, the productive potential of surface water such as rivers or lakes is not sufficient to cover the increasing sectoral demands for fresh water. Therefore, development of groundwater resources has greatly increased on a world-wide scale during the second half of the 20th century. Where available in appropriate quantity and quality, groundwater aquifers are a convenient freshwater storage. Due to over-abstraction, groundwater levels have regionally declined in different areas of the world over this time period. This phenomenon is an indication of non-sustainable resource utilization. It characterizes situations of resource mining where mean recharge to the resource is inferior to what is being abstracted over a prolonged period of time [1].

In most cases, falling groundwater levels have undesired consequences. First, lowered water tables induce rising provision costs due to increased energy requirements for water lift. Second, hydraulic change in the groundwater can cause pollution of pumped water where natural contamination or anthropogenic pollution sources are present. As an example, mining of aquifers in coastal regions can cause large-scale inland saltwater intrusion. These undesired and negative effects resulting from resource exploitation call for careful management of groundwater systems.

As any other economic activity, groundwater management should be carried out in an efficient, that is, optimal way so as to minimize negative consequences associated to exploitation. For this purpose, the managers have to choose appropriate pumping locations from where the water can be pumped and consequently conveyed to demand centers.

Tobias Siegfried and Wolfgang Kinzelbach are with the IFU, the Institute of Environmental Engineering, ETH Zurich, Switzerland ({siegfried, kinzelbach}@ifu.baug.ethz.ch). Eckart Zitzler and Stefan Bleuler are with the Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Switzerland ({zitzler, bleuler}@tik.ee.ethz.ch). Marco Laumanns is with the Institute for Operations Research (IFOR), ETH Zurich, Switzerland ({laumanns}@ifor.math.ethz.ch).

Any set of built boreholes with enough capacity, connected through an appropriate conveyance network, is sufficient for such purpose. According to the planners' preferences and resource as well as capital constraints, they might opt to concentrate a few boreholes locally. Similarly, environmental and capital constraints allowing, it might be preferable for them to spread the provision of the demanded quantity regionally over a large number of boreholes. Hence, each one of the decision-makers faces the following set of questions,

- How many pumping facilities should be installed?
- Where should such facilities be located?
- How should pumping be distributed over such facilities in time?
- What is the structure of the conveyance system that routs water from pumping facilities to the corresponding sink?

given demand and resource constraints so as to minimize the associated water provision costs, each one of the planners having his own set of preferences (see Figure 1 for a sketch of two such possible realizations of resource allocation). The groundwater management problem therefore belongs to the class of location-allocation problems as described by Daskin [2].

Indeed, this is a challenging task, namely for clumping or *centripetal* and dispersive or *centrifugal* forces that govern cost-minimal spatial patterns of groundwater utilization (see [3] for a discussion of such forces in an economic context). Transportation costs limit the range from which water can be conveyed to a demand location since net costs are proportional to the transportation distance. However, the nature of aquifer response with multiple drawdown cones developing and interfering at the pumping locations leads to increasing lift costs in time. Drawdown constraints imposed by a limited screen depth, local gradient constraints as well as an economic scarcity depth beyond which lift costs become forbiddingly high, all promote dispersion. That is, the relocation of pumping in the course of time to more distant places which, so far, are less affected by exploitation.

Additional complexity arises due to political and legal systems that determine resource property rights. Political delineations usually do not follow natural groundwater catchment boundaries. Rather, many different types of decision makers at different societal and institutional scales compete with each other in finding inexpensive allocation solutions - from the small scale farmer to the rich land owner up to the districts and national water agency respectively. Due to the nature of the groundwater resource and its response to abstraction, each of the users best utilization strategy depends on the strategy of the other users. In economic parlance, users inflict stock externalities upon each other and face the classic problem of the commons (e.g. [4]). Thus, catchment-wide optimal groundwater management inherently involves multiple conflicting objectives.

In this paper, we present a multi-objective evolutionary algorithm for groundwater management that can be utilized as a decision support system for planners in complex situations as described above. The focus of the underlying model is on areas where groundwater is the sole resource that can be tapped in an economic manner. Groundwater flow is simulated using a finite-difference discretization of the aquifer domain. Since pumping rates can vary continuously over time at any particular location on the discrete domain, the optimization is of mixed discrete-continuous type.




Fig. 1. Management options in groundwater systems. The effects on the groundwater level h of different borehole configurations in the same aquifer are shown in profile. Left: A specified demand is pumped from one borehole, resulting in minimal infrastructural costs but leading to excessive drawdown. Right: The pumped quantity is spread over a large number of boreholes minimizing the drawdown while causing high installation costs.

Furthermore, moving boreholes are implemented to account for the fact that formerly chosen pumping locations may become inactive due to constraint violation and consequently have to be replaced. The application is, as the above discussion suggests, multi-objective in nature: The summed, discounted allocation costs arising for each groundwater user over the optimization horizon represent the different optimization criteria. The overall optimization task is to approximate the set of cost minimal allocation patterns on the basis of which a post-optimization trade-off analysis and the decision making process can be carried out.

The proposed multi-objective evolutionary algorithm not only represents a valuable tool for groundwater management per se, but also can be useful for the designers of multi-objective optimization algorithms. The application is implemented according to the platform-independent PISA standard (see [5]), which separates the application specific components, such as representation and variation of solutions, from the application-independent components, such as ranking and selection of solutions. Thereby, this application can serve as a real-world benchmark problem by means of which the performance of specific multi-objective algorithms can be assessed.

II. PROBLEM DESCRIPTION

A. Background

In groundwater, the piezometric head or level $h(t)$ at a certain time and location is a specific measure of total energy per unit weight above a datum and measured as a water surface elevation [6]. The physics of groundwater flow is such that ongoing pumping from a borehole in an aquifer induces a horizontal hydraulic gradient towards the borehole in the piezometric head. Because of this gradient, a localized cone of depression develops around the borehole. The dimensions of a cone of depression will depend not only on the pumping rate through time, but also on the hydrogeological

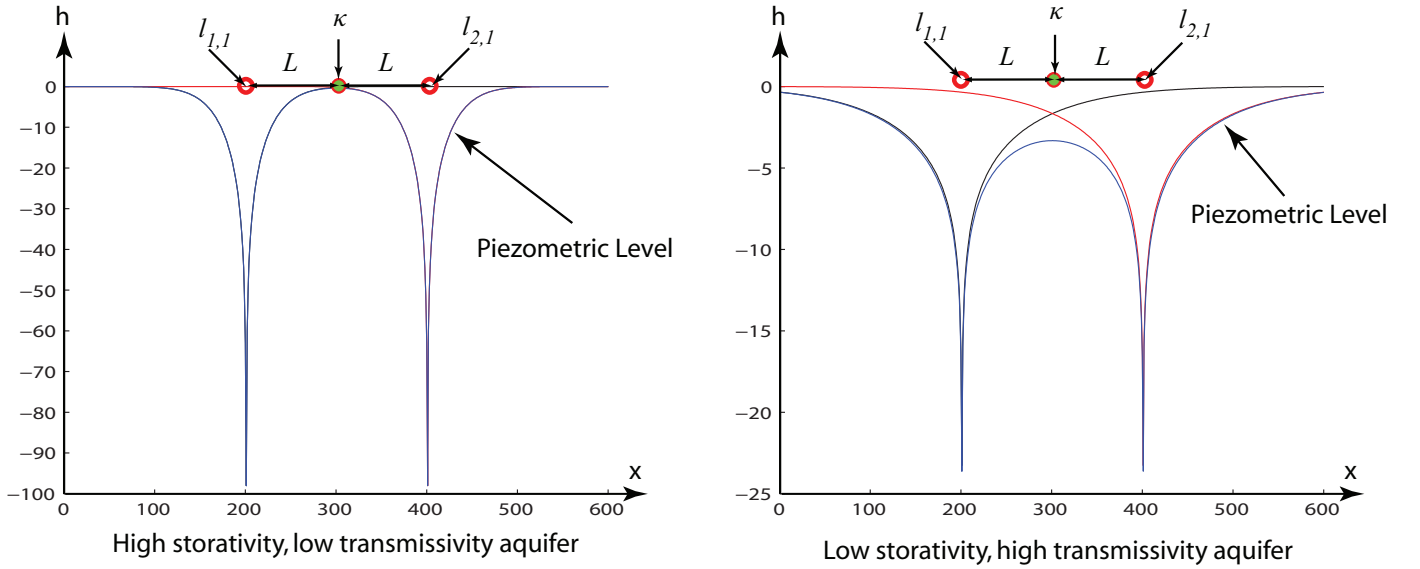


Fig. 2. Two sample aquifers with homogeneous transmissivity and storativity. Left: Small diffusivity aquifer; Right: High diffusivity aquifer. It is assumed that similar quantities of water are supplied from two boreholes located at $l_{1,1}$ and $l_{1,2}$ each one being at distance L from the demand center κ . The black and red lines show the individual drawdown components at time t as well as the magnitude of the borehole interference. The blue line, a superposition of the individual drawdown cones, shows the observable piezometric level $h(t)$ at time t . In the case of the small diffusivity aquifer, external stock effects are practically nil at the time of observation (notice the differing scaling of the vertical axis).

variables transmissivity T and the storage coefficient S by which the aquifer is characterized. With ongoing pumping, the drawdown cone expands over the domain and borehole interference is observed in case of multiple boreholes (see Figure 1).

For the purpose of illustration, Figure 2 shows an identical pumping configuration in 2 aquifers with similar geometry but differing characteristics. The aquifers diffusivity α determines the shape of the drawdown cones¹. In the case of small diffusivity, these drawdown cones are of limited spatial extent but very deep. Contrary to that and assuming similar pumping rates, the piezometric lowering in a high diffusivity aquifer is less deep but extending over a considerable area at the same time of observation. Whereas, the former aquifer exhibits private property characteristics, the latter shows far reaching stock effects and is therefore of common good characteristics in economic terms. Since drawdown cones grow in space and time, any aquifer will exhibit common property characteristics sooner or later.

In other words, the placement of boreholes in an aquifer system has a strategic component – namely the drawdown inflicted by one user upon the others. Imagine a catchment with only two groundwater resource users. A particularly interesting configuration for planner A might be – circumstances permitting – to place his boreholes on the border to country B to effectively divide the stock effects felt by himself by a factor of 2. Clearly, B would not be supportive of such a strategy. In fact, each spatial configuration is a trade-off between the individual objectives of the planners involved. In such a context, a rational, utility-maximizing agent should also consider what the choice of his own allocation strategy

¹ α is the ratio of transmissivity T and storativity S (see (Section II-C) for an explanation of these parameters that describe aquifer properties).

implies for the choices of the strategies of the other planners and what these in turn imply for his own choices and so on ad infinitum.

Any agreement on a basin-wide coordination of allocation ultimately depends on whether the outcome of such planning is an improvement over the *Status Quo* allocation, i.e. the payoff for each one of the planners if such agreement were to fail [7]. Instead of a single optimal allocation solution, a decision-supporting model for the allocation of a groundwater resource should therefore identify a set of compromise solutions for which none can be said to be better than the others in the absence of preference information from planners. Such a set is called a set of *Pareto-optimal* solutions upon which multilateral negotiations can be based.

Traditionally, such problems could only be solved by making oversimplifying assumptions about aquifer response and the shape of the cost functions. Furthermore, they relied on the aggregation of the objectives into a single, parameterized goal function. The optimizer would then systematically vary the parameters to achieve a set of solutions that approximate the Pareto-optimal set [8]. These approaches, however, are sensitive to the shape of the Pareto-optimal front and generally do not exploit synergies between solutions that could help to reduce computational time of the search [9]. This is a major drawback in cases where the search space is conceivably complex as with such location-allocation problems.

In the simulation-optimization groundwater literature, there exist several examples of the application of stochastic optimization techniques to nonlinear discrete-continuous problems. Dougherty and Marryott [10] present a simulated annealing algorithm to solve groundwater management models formulated as combinatorial optimization problems. McKinney [11] uses a genetic algorithm approach to solving groundwater management problems with discrete decision variables. The authors demonstrate the methodology for a variety of hypothetical examples and compare the genetic algorithm solution with a linear or nonlinear programming solution. Rogers and Dowla [12] combine artificial neural networks with a genetic algorithm to design remediation strategies for a hypothetical aquifer. In all of the above studies, the genetic algorithm solution compares favorably with the solution found using the gradient-based methods.

Wang and Ahlfeld [13] present a groundwater management model that identifies optimal borehole locations and pumping rates by defining the spatial coordinates of boreholes as decision variables to be determined along with the pumping rates. Their method uses a composite objective that minimizes pumping and also constrains the boreholes to be within a predefined sub-domain of the model. Using a hypothetical aquifer system, they study the problem of groundwater remediation design with the goal of identifying the minimum pumping remediation strategy. For their examples, the 'moving borehole' model outperformed the 'fixed borehole' model both in terms of total pumping and number of boreholes needed. Finally, Wang and Zheng [14] developed the *Modular Groundwater Optimizer* which is an evolutionary algorithm (EA) that can be readily coupled with the MODFLOW code [15] for the study of optimal site remediation. Problems of maximum contaminant removal for given costs or minimal costs for given contaminant removal can be tackled as shown by Zheng and Bennett [16].

None of the above studies however uses these recent optimization techniques in a context of multiple, conflicting

objectives. The paper by Erickson, Mayer and Horn [17] is a notable exception to that where the authors investigate pump-and-treat strategies to remove the most contaminant at the least cost within a small borehole catchment.

B. Model

The simulation-optimization model presented here has the following characteristics. We assume that a set of K spatially distributed sinks, i.e. demand centers exist to which m decision-makers have to provide water with $1 \leq m \leq K$. In other words, each decision-maker is in charge of at least 1 demand center $\kappa_m \in K$.² Furthermore, the total water quantity $Q_\kappa(t)$ demanded at a certain sink κ at time t is provided by supply from a finite number of I_κ boreholes each of which is positioned at a certain location $l_{\kappa,i}$ with $i \in I_\kappa$. Note that in our model each borehole supplies a unique demand center κ . Groundwater management refers to the m th decision-makers task of the definition of borehole specific pumping time-series $q(l_{\kappa,i}, t)$ for all i and κ so that the sum of the resulting present provision costs gets minimized given constraints.

The model is multi-objective with a problem dependent number m of deterministic cost objectives to be minimized. For each decision-maker, it associates present water provision costs with a groundwater allocation pattern by the coupling of the finite difference groundwater model with a cost module. Facility and pipeline installation costs as well as pumping and conveying costs are taken into account on the supply side. It involves both, discrete (i.e. borehole locations $l_{\kappa,i}$, finite-difference aquifer geometry and model time discretization t) and continuous decision variables (the discrete time-varying control policies $q(l_{\kappa,i}, t)$, i.e. yields, and the state variables $h(l_{\kappa,i}, t)$, i.e. the piezometric levels) subject to conditions which any feasible solution must satisfy. The model implements the option of moving boreholes in case of constraint violation. Territorial restrictions such as political boundaries within a catchment can be accounted for. Optimization is carried out without any prior preference information given by the decision-makers³. Finally, the resulting set consists of solutions which are equally important and could form the basis for negotiation.

The optimization goal is to *minimize* the vector valued cost function \mathbf{C} with

$$\mathbf{C} = f(\mathbf{q}(t), \mathbf{h}(t), \mathbf{L}) \tag{1}$$

where each entry $\mathbf{C} = (C_1, C_2, \dots, C_m)$ is the discounted present cost from the initial time $t = 0$ to the simulation time horizon $t = T$ for the corresponding planner subject to

²The subscript m in κ_m uniquely identifies each demand center's affiliation with a corresponding decision-maker. For notational simplicity, it is dropped where not explicitly needed.

³Note however that preference knowledge is used by the EA in the sense that it provides a complete ordering of sets of solutions. This ordering determines the optimization process (see [18]).

$$h(\mathbf{l}, t) \geq \underline{h}(\mathbf{l})$$

$$\nabla h(\mathbf{l}, t) \geq 0 \tag{2}$$

$$Q_\kappa(t) = \sum_i q(l_{\kappa,i}, t) \quad \forall \kappa$$

In Equation 1, $\mathbf{q}(t)$ are the m decision matrices of pumping whereas $\mathbf{h}(t)$ are the corresponding matrices of piezometric levels. Each entry describes pumping or piezometric head respectively at a certain time t and location $l_{\kappa,i}$. \mathbf{L} are the associated matrices of distances over which water is being conveyed. In the latter, each entry contains the conveying network routing length from each borehole at $l_{k,i}$ to the matching sink κ .

Equation 2 sums up head, gradient and pumping constraints and defines the set of feasible $(\mathbf{l} \times T)$ decision matrices $\mathbf{q}(t)$ and states $\mathbf{h}(t)$. \mathbf{l} is a matrix of all locations where pumping is taking place at a specific time t with $t \leq T$. Hence, if $h(\mathbf{l}, t)$ are the heads at all \mathbf{l} at time t then $\underline{h}(\mathbf{l})$ denote the drawdown constraints for the corresponding locations. $\nabla h(\mathbf{l}, t) \geq 0$ are gradient constraints where applicable. The gradient constraint criterion can be thought of as being head constraint criterion in the finite difference cells that are neighboring cells of potential contamination sources. Finally, $Q_\kappa(t)$ is the total water quantity demanded by the m th decision-maker at time t at a particular demand location κ . $\sum_i q(l_{\kappa,i}, t)$ indicates that this quantity is supplied by the i boreholes that deliver water to κ .

Let us denote boreholes where a constraint violation occurs at time t with $\hat{l}_{\kappa,i}(t)$ and $0 > t > T$. In such cases, the affected boreholes have to be temporarily switched off until a head recovery is observed so that the drawdown constraints as defined by Equations 2 are no longer violated. After such a termination of operation however, the demanded quantity $Q_\kappa(\tau)$, $\tau > t$, has to be redistributed over the remaining pumping stations $l_{\kappa,i}$. Unless offset by decreasing overall demand or a quick drawdown recovery at switched off locations, pumping rates at boreholes in operation consequently rise as more and more boreholes have to be switched off due to the lowering of the piezometric levels at increasingly fewer places. In the most extreme case, this would lead to a final concentration of pumping at one particular place while $t \leq T$. Clearly, such a model would not resemble much any real world allocation strategies of decision-makers that would rather replace temporarily switched off pumping locations by alternative ones. As such, this allocation dynamics becomes part of realistic future groundwater resource utilization strategies. We refer to this dynamics as *moving boreholes* (see also [19] for more details). A particular heuristics for the de- and reactivation as well as relocation of boreholes is presented in Section III-B.

The optimization task is to approximate the set of Pareto-optimal location-allocation strategies upon which future discussions on the common utilization of the groundwater resource under consideration can be carried out. In the following subsections, the essential aspects of the simulation-optimization model are discussed in more detail.

C. Groundwater Resource Model

The hydraulics of groundwater flow in a confined aquifer is governed by the following diffusion-type equation

$$\nabla \cdot (m_A K \nabla h(t)) + q(\mathbf{l}, t) = S \frac{\partial h(t)}{\partial t} \quad (3)$$

where $K = f(x, y, z)$, a function in space, is the hydraulic conductivity and m_A is the saturated aquifer depth with $m_A = h(t) - b$, b being the aquifer bottom. The product $T = K \cdot m_A$ is called transmissivity. $q(\mathbf{l}, t)$ are source and sink terms such as pumping rates at specific locations \mathbf{l} . S is the spatially variable storage coefficient and $h(t)$ the piezometric level at a certain time t and location within the aquifer domain (e.g. [6]). Equation 3 is a second order partial differential equation for the unknown head distribution as a function of time and location. It is defined in a domain Ω with boundary Γ and subject to the initial conditions

$$h(t_0) = h_0 \quad \text{in } \Omega \quad (4)$$

and boundary conditions

$$T \partial h(t) / \partial n = \lambda(H - h(t)) + Q \quad \text{on } \Gamma \quad (5)$$

where H is the prescribed head of an external source such as a stream or lake and λ a known function. Finally, Q is a prescribed flux on the boundary.

In order to find a solution of Equation 3 to a particular problem, the relevant physical coefficients (i.e. S and T), the domain geometry as well as initial and boundary conditions have to be provided. Except for very simple geometries, analytical solutions are difficult to obtain. Therefore, numerical methods are applied to approximate such solutions (see for example [20], [6] or [21]). One such approach is the finite-difference method wherein the continuous system is replaced by a finite set of discrete points in space and time and the partial derivatives are replaced by terms calculated from the differences in head values at these points. The process leads to a coupled system of linear algebraic equations and their solution yields the head values at specific locations and times.

D. Supply infrastructure

Each feasible realization r of the pumping and conveyance infrastructure consists of K given locations, where a specific rate $Q_\kappa(t)$ is consumed at time t , $0 < t < T$. This demand can be supplied by a random spatial configuration of pumping boreholes that are distributed at locations \mathbf{l} within the catchment domain Ω .

In order for the pumped water to be delivered to the final sink, the boreholes together with the corresponding demand centers are modeled as a directed network with a set of arcs connecting nodes. With that, any of the K directed networks of a particular realization r of a pumping configuration is fully described by the set $\aleph_\kappa(r) = \{\kappa, l_{\kappa,i}, q(l_{\kappa,i}, t), \mathbf{A}_\kappa\} \forall i \in I_\kappa, \kappa \in K$ where \mathbf{A}_κ describes the set of arcs between matching nodes.

A possible algorithm for the determination of such a directed network and its associated allocation strategy contains the following steps:

1. Directed network membership identification, $l \rightarrow l_{\kappa,i}$.
2. Creation of network topology \mathbf{A}_{κ} .
3. Determination of $q(l_{\kappa,i}, t)$ for each borehole location $l_{\kappa,i}$ given κ .

The creation of the network topology is as follows. Out of the set of boreholes assigned to a certain network cluster, we identify the location $l_{\kappa,*}$ for which $\overline{l_{\kappa,*}\kappa} < \overline{l_{\kappa,i}\kappa} \forall i$. It is assumed that a principal conveyer connects $l_{\kappa,*}$ with κ through which all of the pumped water is conveyed to the according demand center. This implies that the remaining $l_{\kappa,i} \neq l_{\kappa,*}$ are connected to the borehole at location $l_{\kappa,*}$. With that, the network topology is entirely defined.

E. Cost Function

The j th entry of the cost function \mathbf{C} can be expressed as

$$C_j = \sum_{\kappa_j,i} \sum_{t=0}^T \frac{1}{(1+\delta)^t} (g_{\kappa_j,i}^*(t=t') + g_{\kappa_j,i}(t)) \quad (6)$$

with $1 \leq j \leq m$ and denoting the summed present costs over all directed networks $n(j)$ that are within the control of the j th decision-maker.

The costs are composed of fixed costs $g_{\kappa_j,i}^*(t=t')$ and running costs $g_{\kappa_j,i}(t)$. The former occur at time t' when a particular borehole and its corresponding conveyer are built and are proportional to the conveying distance and the borehole drilling depth. Hence, if $t' = 0$ then that particular borehole is in operation from the beginning of the simulation. Arc costs can be related to boreholes $l_{\kappa,i}$ simply by adding them to the borehole costs of the respective installations which are located at the origin of each corresponding arc in conveyance flow direction.

Contrary to that, $g_{\kappa_j,i}(t)$ are running costs that account for the energy requirement so that the water can be pumped and conveyed to the final sinks. Over time, these costs vary linearly with regard to the pumping depth and non-linearly in dependence of the borehole pumping rates $q(l_{\kappa,i}, t)$ and the conveyed water volumes. Running costs of particular locations might be non-continuous due to possible constraint infringements over the course of the simulation and the subsequent intermittent switching on and off of pumps. Both cost components are current costs discounted by the discount factor δ that reflects the time value of money. The precise definition of the current costs can be found in [19].

III. OPTIMIZATION METHOD

The general algorithmic structure developed to solve the optimization problem specified above is shown in Figure 3. The *variation* creates new offspring individuals out of the set of parent individuals ξ that were chosen for mating during the process of mating selection. Subsequently and according to Equation 6, the objective vectors are calculated. Based on that, the environmental selection then determines whether offspring individuals will be stored in the archive or discarded. This process is repeated until a termination criterion is met.

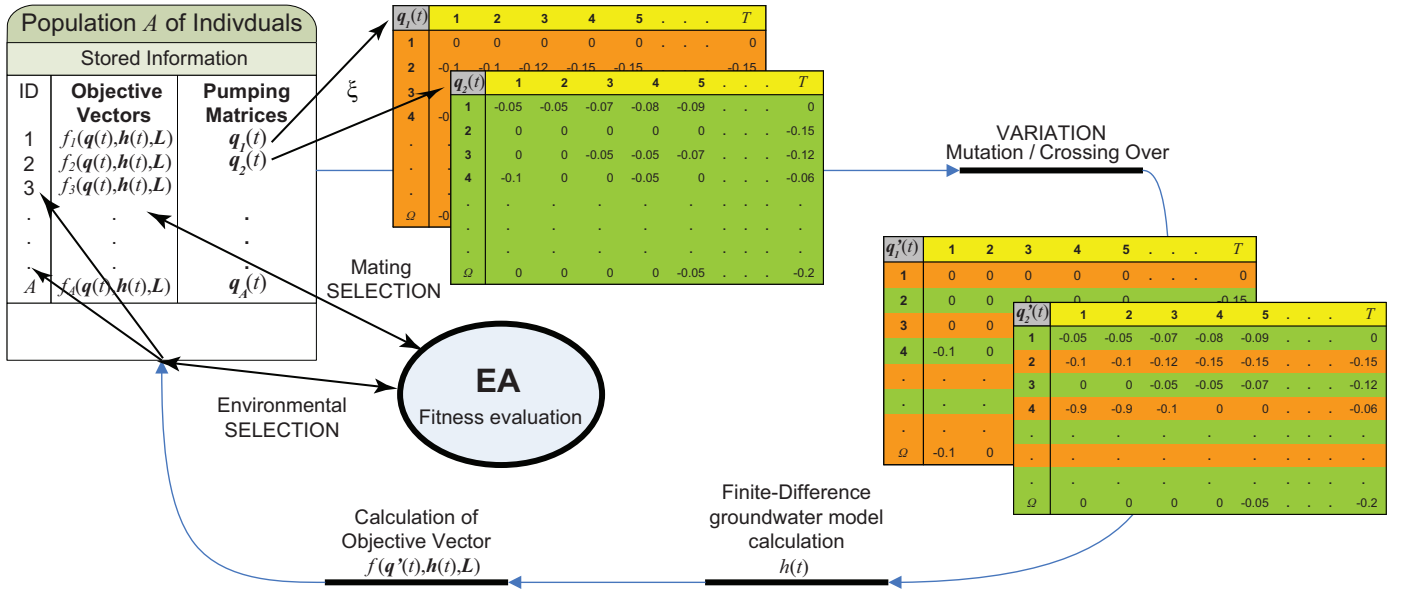


Fig. 3. Illustration of the problem representation in the context of a generational cycle. The blue arrows indicate a generational cycle of which one is shown. All of the relevant information is stored in the population-structure. From there and based on the objective vector $f(q(t), h(t), L)$, the evolutionary algorithm chooses a set ξ of individuals for the variation in a process called 'mating selection'. The variation operates on the decision-matrices $q(t)$ solely. Crossing over interchanges individual characteristics as shown by the color coding. Contrary to that, mutation randomly alters pumping time-series of a particular individual. After the variation, the states $h(t)$ that correspond to the individuals are updated. Based on that, the new objective values are computed.

A. Representation

All of the coding has been carried out in Matlab. The modularity of the programming approach allows to add new model characteristics quickly (see Figure 4). The underlying data such as temporal economic data and spatial aquifer characteristics are data made available in matrices. Like this and as more detailed information on price structures or aquifer parameters gets available, changes can be incorporated in a supposable simple manner.

Each individual is stored in a Matlab-structure that accommodates the corresponding data arrays. Similarly, the population A collects all individuals in such a structure. Where feasible, sparse matrices are utilized. Initially, random realizations of pumping time-series are defined for each individual in A . The pumping time-series $q(l_{\kappa,i}, t)$ are random variables in space as well as time and arbitrarily chosen, subject to the constraints defined by Equation 2.

Subsequently, the determination of network membership for each borehole, i.e. $l \rightarrow l_{\kappa,i}$, is accomplished by k-means clustering. The latter partitions the set of boreholes into K mutually exclusive supply clusters each one corresponding to one particular demand center κ . For our purpose, the `kmeans`-Matlab function implementation has been used for the cluster identification which solved the clustering problem by numerical methods [22]. Like in other problems of numerical minimization, the choice of the initial partition can greatly affect the final clusters that result. To improve optimization performance of the `kmeans`-function, the demand locations κ are chosen as initial guesses for the primary cluster centroids.

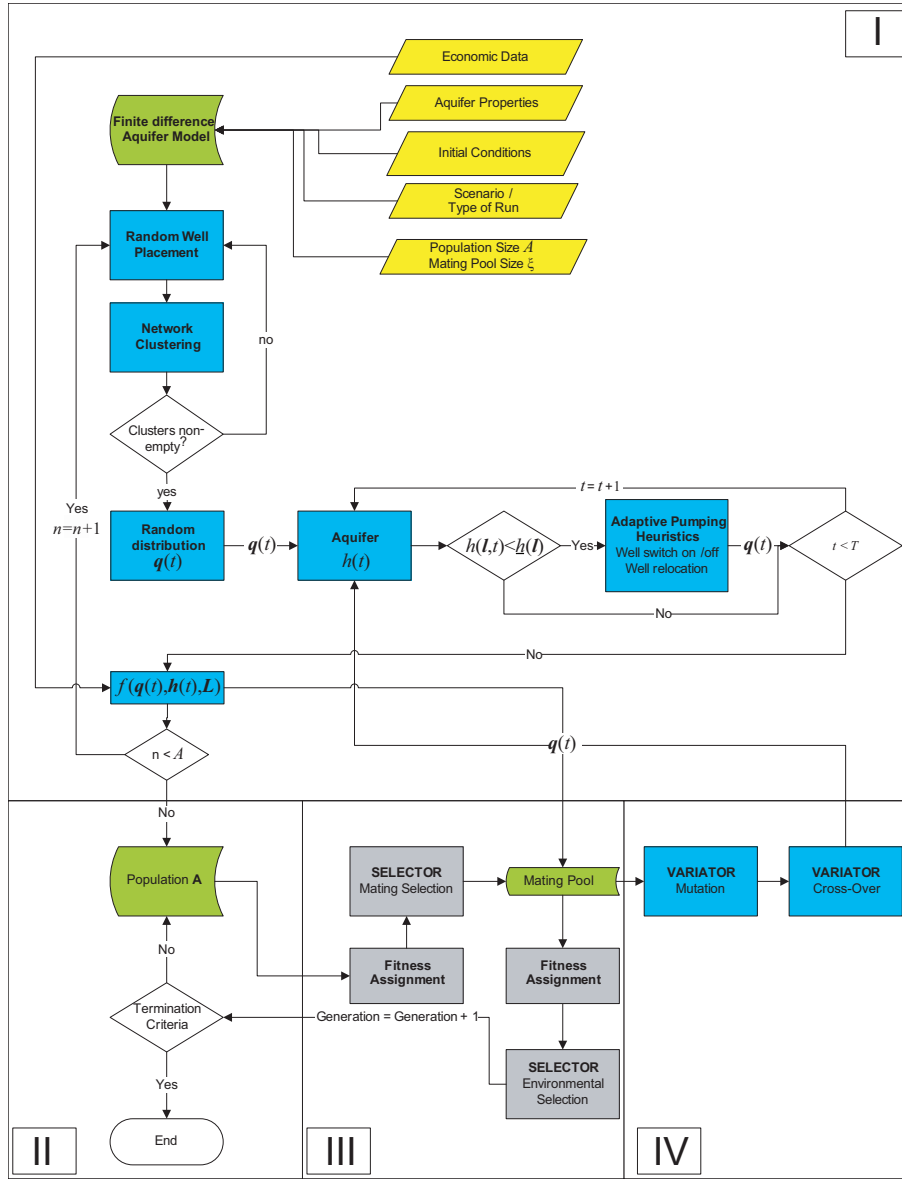


Fig. 4. Structure of the EA-based simulation-optimization framework. Module I: Creation of population A, calculation of corresponding states and objectives; Module II: Archive of individuals storage, maintenance and book keeping; Module III: Environmental and mating selection based on fitness assignment; Module IV: Random mutation of pumping and crossing over of conveyance networks.

The physical problem of calculating the spatio-temporal development of $h(t)$ on Ω is solved in Modflow-2000 which is a state-of-the-art finite-difference groundwater flow model solver [23]. Details with regard to the coupling of the EA and Modflow-2000 can be found in [19]. In order to avoid excessive memory requirements, state information is not permanently stored in the archive. It can be restored at any later state conveniently by the calculation of the finite difference groundwater model utilizing the corresponding pumping time-series.

B. Constraint Handling

Infrastructural or operational change in existing borehole configurations at a certain point in time may be necessary due to the violation of constraints. One possible way to deal with constraint violation are methods based on penalty

functions. Approaches incorporating these penalties suffer from two problems. First, the optimal set of solutions depends on the magnitude of the penalty. Users normally have to try different penalties to find a value which steers the search towards the feasible region. Second, the inclusion of a penalty term may critically distort the objective function [9]. We therefore propose a different implementation that has a further advantage of being capable to reproduce the spatial dynamics of resource allocation strategies as discussed above.

In order to prevent successive switching off of boreholes and a resulting concentration of an initially distributed pumping on the domain, we implement the option of moving boreholes. In other words, if local pumping and quality constraints can no longer be met, relocation of boreholes will have to happen. The dynamics is represented by a heuristics that relocates pumping at a particular location $l_{\kappa,i}$ once a head or gradient constraint is no longer fulfilled there. Boreholes get relocated in the direction of the steepest ascent, i.e. $\nabla h(l_{\kappa,i}, t)$. If the direction of steepest ascent is ambiguous, the decision with regard to direction is randomly taken. The displacement distance is assumed to be proportional to $|\nabla h(l_{\kappa,i}, t)|$.

C. Variation - Mutation

After the process of mating selection, the selected candidate individuals are passed on for variation. There, the mutation operator arbitrarily alters one or more components of a selected individual so as to increase the structural variability of the population A . The role of mutation is that of restoring lost or unexplored genetic material in A to prevent the premature convergence of the algorithm to suboptimal solutions and it ensures that the probability of reaching any point in the search space is never zero [9].

The mutation operator for the problems discussed can be implemented in many different ways (see [10] and references therein). Mutation, as implemented here, has 2 aspects – mutation of the spatial location of selected boreholes and mutation in time of the pumping rates of selected boreholes. Boreholes are randomly chosen for mutation if $p_{\kappa,i}(t) < p_{mu}$, where $p_{\kappa,i}(t) \in [0, 1]$ is a random number assigned to each active borehole $l_{\kappa,i}$ at time t and p_{mu} is a user defined mutation rate. If for all boreholes and all times $p_{\kappa,i}(t) > p_{mu}$ and in the absence of crossing over, then variation becomes a simple reproduction.

The principle of location mutation is depicted in Figure 5. Basically, the easting and northing (x, y) of a set of randomly chosen boreholes get mutated according to $(x', y') = (x, y) + (\Delta x, \Delta y)$ if $(x', y') \in \Omega$ with Δx and Δy being independent uniformly distributed random numbers that are drawn from a range $[-\sigma, \sigma]$. The latter is measured over a certain number of finite difference cells and represents a user specified measure of the range of mutation. It can be dynamically decreased over the course of optimization by an annealing factor until a user specified lower limit σ_{min} is reached. Like this, critical distortions by large mutation steps at a late optimization stage can be prevented.

Mutation in the pumping rate adds random disturbance to the distribution of pumping $q(l_{\kappa,i}, t)$ within a certain supply network κ at a certain time t . The mutated pumping rate at a particular borehole needs to be counterbalanced by a corresponding change in the remaining locations chosen for mutation so as not to violate the demand constraint specified in Equation 2. Therefore, if the list of borehole locations to experience pumping rate mutation contains less




Fig. 5. Illustration of the spatial mutation of chosen boreholes within the aquifer domain Ω . The red square is a switched off borehole whereas the blue squares indicate active boreholes of which 2 are relocated by the process of mutation. The windows within which mutation is feasible, i.e. the user specified mutation ranges, are indicated by the bold black frames with width 2σ . After relocation, the network infrastructure is updated correspondingly (before: bright blue lines; after: dark blue connecting lines).

than two entries, no such mutation takes place.

D. Variation - Crossing Over

The crossover operator is a method for sharing information between individuals; it combines the features of two parent individuals to form two offsprings with the possibility that good individuals may generate better ones, assuming a certain degree of decomposability of the problem. The type of the problem suggests the implementation of a simple crossover algorithm which has been proposed by [24].

If two realizations, i.e. parent individuals r_u and r_v have been selected for crossing over, they swap randomly chosen directed network clusters. In other words, they interchange selected supply infrastructure $l_{\kappa,i}^u \forall i, i \in I_{\kappa}^u$ and $l_{\kappa,j}^v \forall j, j \in I_{\kappa}^v$ for similar demand centers $\kappa^u = \kappa^v$. Conveyance clusters of particular demand centers are swapped if $p(\aleph(r_u)) \geq p_{xo} \wedge p(\aleph(r_v)) \geq p_{xo}$ where $p(\aleph(r_u))$ and $p(\aleph(r_v))$ are 2 vectors of random numbers between 0 and 1 assigned to the given set of the K demand center and p_{xo} a user specified crossing over probability.

Imagine for example, the following two parent individuals r_u and r_v swap cluster

$$r_u : \{\aleph_1(r_u), \aleph_2(r_u), \aleph_3(r_u), \dots, \aleph_K(r_u)\}$$

$$r_v : \{\aleph_1(r_v), \aleph_2(r_v), \aleph_3(r_v), \dots, \aleph_K(r_v)\}$$

After crossing over, the offspring's r'_u and r'_v attributes are changed correspondingly. So for example, r'_u and r'_v could become

$$r'_u : \{\aleph_1(r_v), \aleph_2(r_u), \aleph_3(r_v), \dots, \aleph_K(r_u)\}$$

$$r'_v : \{\aleph_1(r_u), \aleph_2(r_v), \aleph_3(r_u), \dots, \aleph_K(r_v)\}$$

according to a particular realization of the random cross-over probabilities.

E. Implementation

The present optimization problem is a typical example of a complex real-world application for which MOEA's have been developed. However, such real-world problems are rarely used to test and compare MOEA's. Often expert knowledge about the application domain and a prohibitively large effort are required for a re-implementation of the complex simulations involved. PISA [25] provides a simple way to reuse existing implementations of both optimization problems as well as optimizers. More specifically, PISA provides an interface between the problem specific part of an EA (subsequently referred to as **variator**), like representation, objective function evaluation and variation on the one hand and the problem independent part consisting mainly of selection on the other hand (called **selector**). Modules on both sides can be freely combined and a variety of modules is available on www.tik.ee.ethz.ch/pisa, including the one presented in this study.

On the **variator**-side, each individual in the population A gets assigned a unique identification number. During optimization, only this number together with each corresponding objective vector $f(\mathbf{q}(t), \mathbf{h}(t), \mathbf{L})$ is passed between the **variator** and the **selector**. This implementation is very efficient since it reduces communication overhead to a minimum (see also Figure 3).

IV. CASE STUDY

Table IV summarizes the default parameter setting utilized in the present case study. We assume that 3 decision-makers are in charge of covering water demand over the period of $T = 5$ years. Each one has to supply water to one corresponding demand center, i.e. $m = K = 3$, as described in Section I. The location of the demand centers is shown in Figure 7. Similar and constant per period water demand is given. We suppose that the decision-makers share cooperatively a generic aquifer which has impervious boundaries and extends over 30 km x 30 km. Cooperation in this sense means that no domain restrictions in the form of political boundaries exist. In other words, each decision-maker can pump at each location within the domain Ω . The aquifer is modeled as a confined aquifer with a finite-difference cell discretization of 1 km x 1 km. Groundwater recharge is nil. The transmissivity field has been synthetically generated using the cross-correlated random field generator FGEN (see [26]) and is shown in 7.

All subsequent simulation runs were performed on an Intel Xeon 3.06 GHz CPU with 2 GB RAM. An average optimization run (10'000 objective function evaluations) took approx. 4 hours.

General simulation parameters		Unit
simulation periods	5	[-]
sim. period length	31536000	[s]
Aquifer properties		
S	2.00E-04	[-]
h_0	1	[m]
Political / Economic parameters		
m	3	[-]
K	3	[-]
$Q_1(t) = Q$	9	[m^3/s]
$Q_2(t) = Q$	9	[m^3/s]
$Q_3(t) = Q$	9	[m^3/s]
C_{IP}	10000	[\$]
C_{IC}	1	[\$/m]
c_T	8.11	[-]
c_L	9000	[-]
Energy costs	0.1	[\$/kWh]
Discount factor	0.03	[-]
Penalty depth	200	[m]
Recovery level	10	[m]
Infrastructure properties		
Borehole diameter	0.3	[m]
Pipe diameter	1	[m]

TABLE I

SELECTED SIMULATION PARAMETER SETTINGS UTILIZED IN THE COMPARATIVE STUDY. Q_1 – Q_3 ARE CONSTANT PER-PERIOD DEMANDS. C_{IP} AND C_{IC} ARE TOTAL PRESENT INSTALLATION COSTS OF BOREHOLE INFRASTRUCTURE, PRESENT PER CONVEYOR METER INSTALLATION COSTS OF THE CONVEYANCE SYSTEM. c_L AND c_T ARE FACTORS OF PROPORTIONALITY THAT ACCOUNT FOR PUMPING PLANT EFFICIENCY AND HYDRAULIC LOSSES IN THE CONVEYANCE PIPE (SEE [27] FOR MORE INFORMATION).

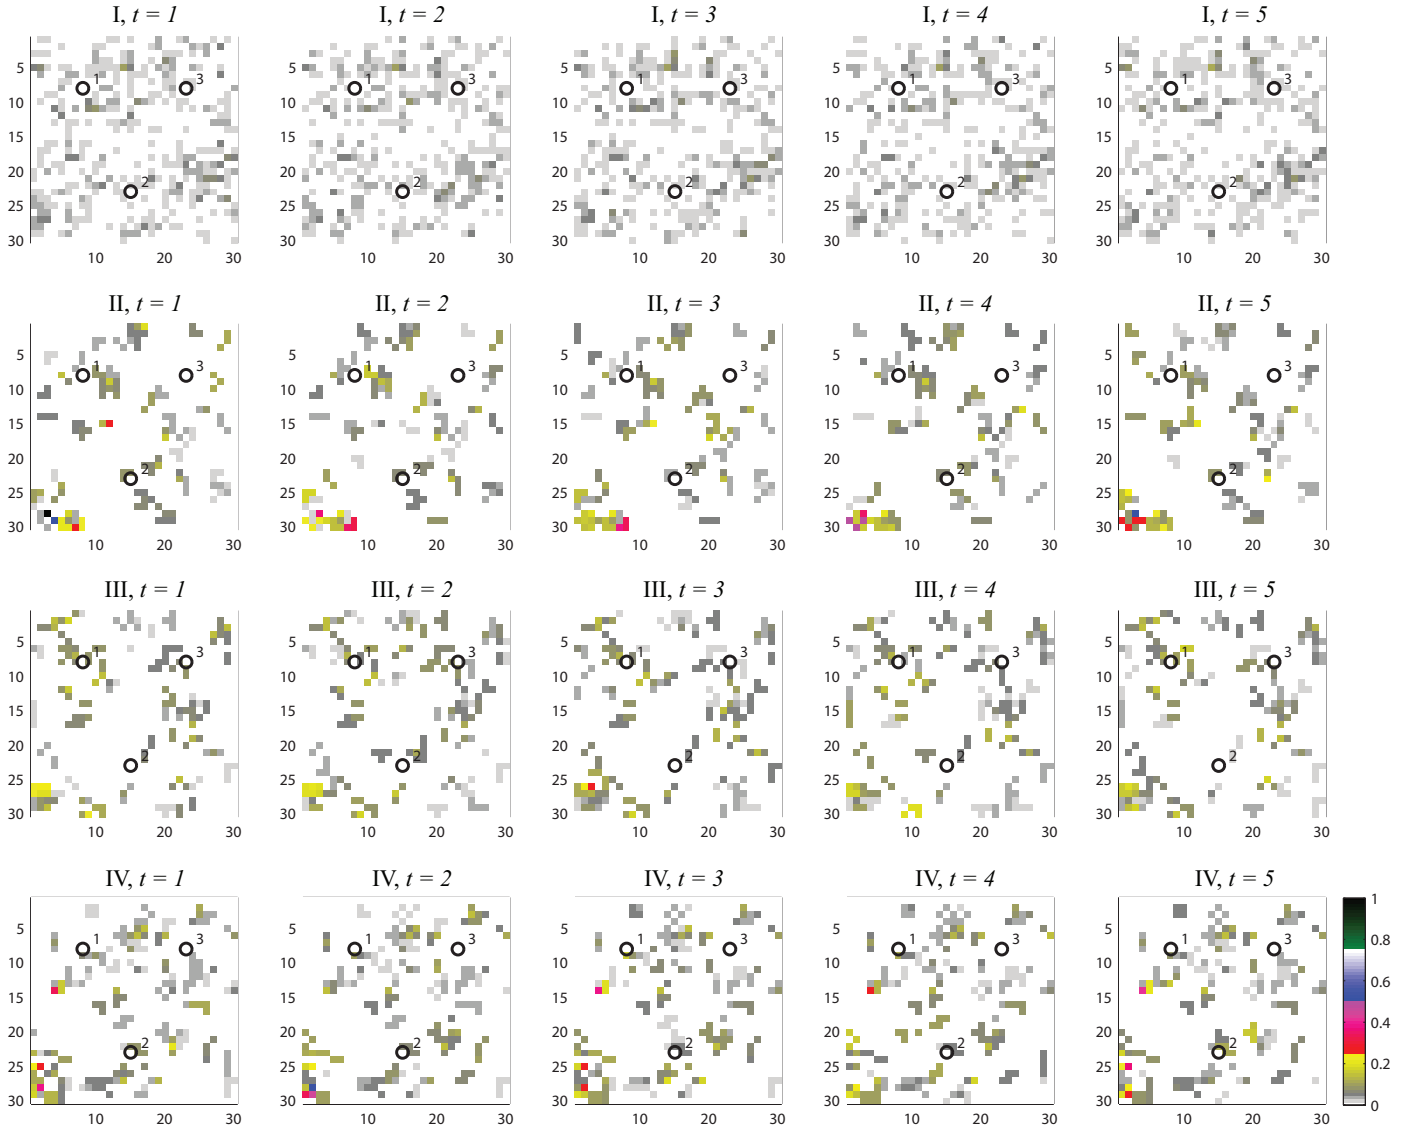


Fig. 6. Location and intensity of pumping of related last generation individuals. Plates in the first row show the ensemble averaged pumping of allocation solutions that are within $\psi = 0.1$ with $\bar{C}_{max} = 2.5028e + 007$ from $\min(\Sigma)$. Rows 2 to 4 depict solutions that are similarly related to each other for $\max(f_1)$, $\max(f_2)$ and $\max(f_3)$ respectively. Values are normalized so that individual Plates are comparable. x- and y-scales of all plates are in kilometers.

A. Simulation Results

Due to the heterogeneity of the spatial transmissivity field we expect recognizable management patterns to emerge in the characteristics of optimal solutions. Figure 8 shows the last generation Pareto-surface of individuals of the run under investigation. Generally, a good spread of the nondominated solutions in the objective space is visible. In order to understand the implications in the field, i.e., the decision space, Figure 9 shows individual supply network configurations of selected individuals together with the depth to the groundwater at the end of the last simulation period. Small per meter conveyer costs C_{IC} were chosen (see Table IV). As a result and as already discussed conceptually in Section II, the pumped quantities are generally spread over a large number of boreholes. This explains the complex spatial geometry of the conveyance infrastructure of optimal solutions.

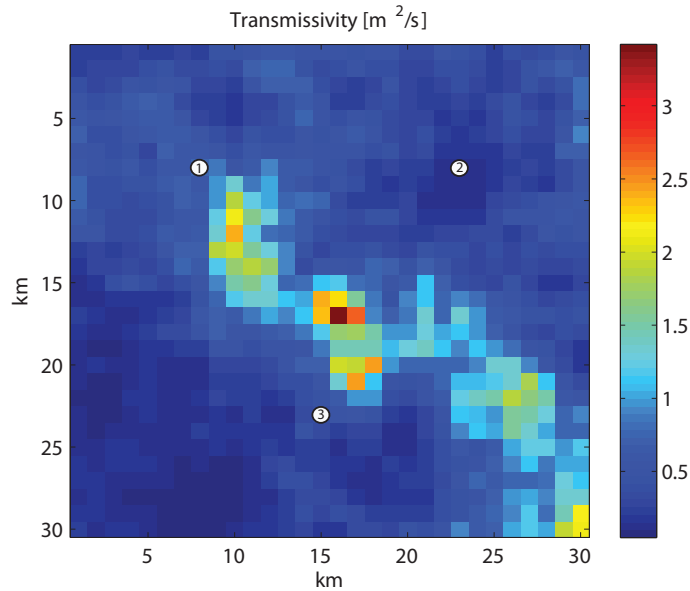


Fig. 7. Transmissivity field utilized in the experimental setup. Yellow to blue colors in the SE–NW diagonal mean moderate to high transmissivity values, while the SW corner is characterized by very low transmissivity values.

The overall cheapest individual $\min(\Sigma)$ is shown in Figure 9 Plate I and characterized by a regionally almost uniform drawdown (standard deviation of heads $\sigma(h(T)) = 1.321$ m over the whole domain). It thus exemplifies a uniform solution to allocation. Plate II shows the overall most expensive individual $\max(\Sigma)$. In this allocation solution, none of the demand centers are supplied by nearby supply clusters. This causes high overland transportation costs. Furthermore, demand center 1 covers its demand by pumping entirely in the south–western low–transmissivity zone where a pronounced drawdown cone is visible which causes high lift costs ($\sigma(h(T)) = 8.96$ m). Compared to this, in the uniform allocation, pumping rates in this low transmissivity region are small.

Plates III–V show how in this common–property regime, prosperity of one implies sacrifice of others. For example, by looking at Plate III, it is visible that the cheapest solution for demand center 1, i.e. $\min(f_1)$, is characterized by nearby supply for the corresponding decision–maker. At the same time, demand center 3 covers a sizeable fraction of its demand in the distant low transmissivity region. This minimizes stock interferences due to the limited spatial extent of the drawdown cone. Similarly, this *gain at the expense of others* phenomena can also be seen in Plate IV and V.

More generally, the characteristics of individuals located in specific regions of the objective space are related. Figure 6 shows normalized average ensemble pumping rates of the individuals v that are located within a certain radius $\psi\bar{C}_{max}$ of a specific parent individual u as defined by $|\mathcal{C}_v - \mathcal{C}_u| \leq \psi\bar{C}_{max}$. ψ is an arbitrarily chosen fraction of the biggest Cartesian distance \bar{C}_{max} between a particular set of two cost objectives of the optimization run under consideration.

On the ensemble average level, individuals in the vicinity of the cheapest solutions are equitable (see top row in Figure 6). In other words, pumping is fairly homogenously spread over the whole domain as well as in time. None of the decision–makers face an over–proportional cumulative share of drawdown. In comparison, maximum individual allocation costs occur when pumping is a) clustered and b) concentrated in the low transmissivity domains over the




Fig. 8. Last generation Pareto surface in the objective space. Selected individual solutions are colored. Green dot: $\min(\Sigma)$; blue dot: $\max(\Sigma)$; red dots: $\min(f_1)$, $\min(f_2)$ and $\min(f_3)$. The 3 tail ends show $\max(f_1)$, $\max(f_2)$ and $\max(f_3)$ correspondingly with $\max(f_1) = \max(\Sigma)$.

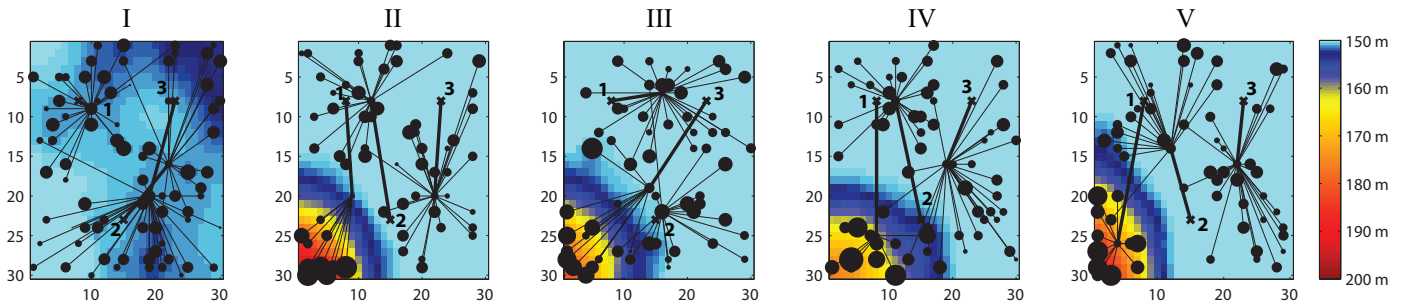


Fig. 9. The plates show the spatial pumping configurations of selected individuals. Demand centers are shown as open circles with the corresponding identification number. Small dots mean low pumping rates, large ones high pumping quantities respectively. Lines connecting the pumping locations show the conveyor system. Dot sizes scale linearly between 0.04 and $10.58 \text{ m}^3/\text{s}$. Furthermore, depth to groundwater is shown. Plate I: $\min(\Sigma)$; Plate II: $\max(\Sigma)$; Plate III: $\min(f_1)$; Plate IV: $\min(f_2)$; Plate V: $\min(f_3)$ (see also Figure 8). x- and y-scales are in kilometers.

simulation time horizon as can be seen by looking at the Plates II–IV in Figure 6. For example, and according to intuition, it turns out that decision-maker $m = 1$ faces the highest present allocation costs if he is supplying the majority of pumped water from the SW corner in the aquifer domain.

At this point, it is important to mention that no generic policy recommendation can be drawn from these findings since these allocation solutions are sensitive to the physical properties of the aquifer, the energy and installation cost parameters as well as the maximum drawdown restrictions $\underline{h}(l)$. Yet, the fact that the overall cheapest solutions are equitable is a strong case for the joint, cooperative management of common property groundwater resources. The proposed simulation–optimization model can help to achieve this as well as facilitate implementation by delivering a quantitative base for decision-making.

B. Comparative Study

The application studied in this paper is not only of interest to the water resource management domain, but can also serve as a real-world benchmark problem in the evolutionary multi-objective optimization community. The high-dimensional search space, the complicated procedure to decode and evaluate the solutions, and the particular shape of the Pareto front (see Figure 8) represent a combination of challenges that most artificial test problems cannot offer. The purpose of this section is, therefore, to use this application problem as a benchmark for an empirical performance evaluation of different multi-objective optimization algorithms and their parameterization. In particular we address the questions

(1) What are reasonable parameter settings for the application scenario under consideration?

(2) Are significant performance differences observable in this problem when applying multi-objective evolutionary algorithms with different selection strategies?

We compare the different algorithms and parameter settings based on the Pareto-set approximation they obtain, that is, the set of non-dominated individuals they maintain after the given number of iterations. These final approximation sets are evaluated using the following quality indicators.

The unary hypervolume-indicator $I_H(A)$ measures how much volume of the objective space is dominated by the obtained approximation set A [28]. It is given by the volume of the union over all $a \in A$ of the polytopes $\{y \in \mathbb{R}^m : a \leq y \leq y_H\}$, where $y_H \in \mathbb{R}^m$ is a reference point. In our case all objective values were normalized to $[1, 2]$, and the reference point was chosen as $y_H = (2.1, 2.1, 2.1)$. A larger hypervolume corresponds to a better approximation set.

The binary ϵ -indicator $I_\epsilon(A, B)$ measures how far approximation set A is from entirely dominating another approximation set B . It is defined as the largest scalar $\epsilon > 0$ which when multiplied to each element of A , produces a scaled set A' such that each $b \in B$ is weakly dominated by some $a' \in A'$. Here, smaller ϵ values are preferable.

Both quality indicators are widely used as they comply with the Pareto dominance relation and are among those indicators with the largest inferential power (see [29]). Nevertheless, each quality indicator represents a certain preference structure for comparing approximation sets that should be kept in mind when interpreting the results.

B.1 Comparison of Mutation / Crossover Intensities and Population Sizes

In a first set of simulation runs we studied the influence of different mutation step sizes and crossover rates on the resulting non-dominated fronts. SPEA2 was used for the following comparison as it performed competitively on this optimization problem, cf. SectionIV-B.2. For each of four different crossover rates $p_{xo} = \{0, 0.3, 0.7, 1\}$ as well as four mutation ranges $\sigma = \{1, 2, 4, 20\}$ 11 runs with different random number generator seeds were performed. The unary hypervolume-indicator $I_H(A)$ was calculated for each resulting front so as to evaluate the quality of the optimization results.

Figure 10 shows that small mutation step sizes yield better indicator values. For large step sizes, mutation introduces larger changes in an individual and thus hinders a steady progress to good solutions. Additionally, for each setting of

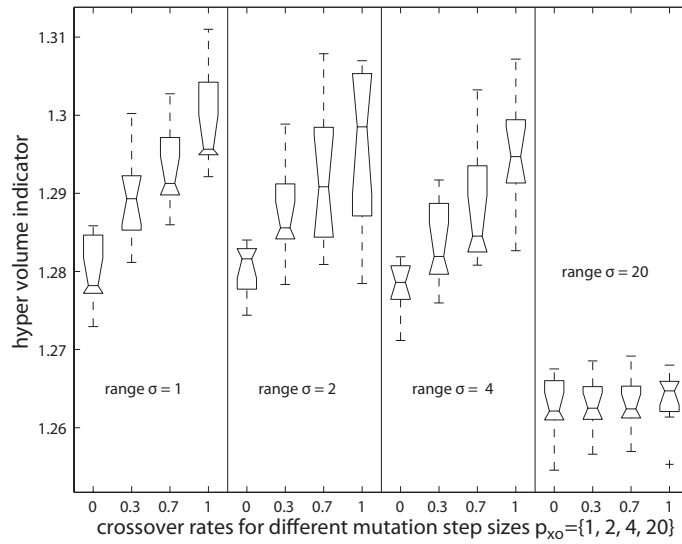


Fig. 10. Comparing different crossover rates and mutation step sizes. Boxplots show the distribution of the 11 runs with different seeds.




Fig. 11. Comparing different population sizes. Boxplots show the distribution of the 11 runs with different seeds.

the mutation step size large crossover rates are clearly preferable, i. e., recombination seems to play an important role in the progress towards good solutions. A possible reason is that the crossover operator exchanges whole network clusters and therefore never disrupts a solution.

In a second set of runs we investigated model sensitivity to population size. For these optimization runs all parameter settings as well as the total number of objective function evaluations was held constant. In other words, the number of generations changed according to the population size. Figure 11 shows that the chosen population size is sensitive to the quality of the solutions. For population sizes smaller than 100 individuals the quality of the approximation set A deteriorates. Above that, model sensitivity to variable population sizes is limited.

B.2 Comparison of Multi-objective Selection Schemes

We compare four different multi-objective evolutionary algorithms, SPEA2 (the improved version of the Strength Pareto EA [30]), NSGA-II (the Nondominated Sorting GA II [31]), IBEA (the Indicator-Based EA [28]), and ABEA (the Adaptive Boxes EA [32]). As we are interested in performance differences due to the selection schemes only, the entire `variator` part was kept identical throughout the following tests.

All four algorithms employ truncation selection for updating the population A (environmental selection) and binary tournament selection (with tournament size 2) for choosing the individuals that undergo variation (mating selection). As both these schemes are rank-based, and all algorithms assign ranks in compliance with Pareto dominance, the major differences arise when selecting from a set of mutually non-dominating individuals. In this case a secondary criterion besides Pareto dominance has to be applied, which is typically some measure of the diversity, but could be any proxy for measuring the contribution of each individual to the quality of the whole population. For instance, nearest-neighbor density estimators are used in NSGA-II (the average coordinate-wise distances to the first neighbors) and SPEA2 (the lexicographic ordering of all Euclidean distances). ABEA used an adaptive coarse-graining of the objective space into boxes where individuals in more densely populated boxes are assigned a worse rank, hence implementing a histogram-type density estimator. The recursive update of the box boundaries precludes any deterioration of the maintained set of nondominated solutions. In contrast to the other three algorithms, IBEA determines the rank of each individual directly, based on its contribution to the population's value of one of the quality indicators discussed above, in this case the ϵ -indicator.

Our experiment consists of 11 runs for each of the four selection schemes discussed above, where for each algorithm the same sequence of random seeds for the `variator` module, and hence the same initial population, were used. The four selection schemes can thus be considered as different treatments in a paired (or quadrupled) sample. We want to use non-parametric testing for the hypotheses that, when trying to improve the given initial population, algorithm A obtains better values than algorithm B for the performance indicators described above, and for all ordered pairs of the four algorithms. This gives $3 \cdot 12 = 36$ tests in total, all of which are carried out with the same samples so that the significance levels need to be adjusted to account for multiple testing.

The hypothesis of the first test is that the output of algorithm A dominates a larger volume of the objective space than algorithm B . The test statistic counts the number of paired samples out of the total of 11 where $I_H(A) > I_H(B)$. Under the null hypothesis that both algorithms achieve identical hypervolume values, the distribution of test statistic will be the binomial distribution with $n = 11$ and $p = 0.5$. The p -values of the one-sided test are 0.05 % for a value of 11, 0.59 % for a value of 10, 3.27 % for a value of 9, and $> 5\%$ for all values less than 9. Table IV-B.2 (top) displays the values of the test statistic for all pairwise comparisons. All values of 11 are significant under a level of $\alpha = 5\%$ even after a simple Bonferroni correction to account for the total of 36 tests. Both SPEA2 and NSGA-II reach significantly higher hypervolume values than both IBEA and ABEA, and IBEA in turn obtains significantly better values than ABEA, while between SPEA2 and NSGA-II no significant differences could be detected.

	$I_H(A) > I_H(B)$			
SPEA2	-	5	11	11
NSGA-II	6	-	11	11
IBEA	0	0	-	11
ABEA	0	0	0	-
	$I_\epsilon(A, B) < I_\epsilon(B, A)$			
SPEA2	-	7	3	9
NSGA-II	4	-	3	7
IBEA	8	8	-	10
ABEA	2	4	1	-
	$I_\epsilon(A) < I_\epsilon(B)$			
SPEA2	-	9	8	10
NSGA-II	2	-	7	7
IBEA	3	4	-	6
ABEA	1	4	5	-
	SPEA2	NSGA-II	IBEA	ABEA

TABLE II

VALUES OF THE TEST STATISTICS *number of runs where $I_H(A) > I_H(B)$* (UNARY HYPERVOLUME INDICATOR, TOP), *number of runs where $I_\epsilon(A, B) < I_\epsilon(B, A)$* (BINARY HYPERVOLUME INDICATOR, MIDDLE), AND *number of runs where $I_\epsilon(A) < I_\epsilon(B)$* (UNARY ϵ INDICATOR, BOTTOM), WHERE A STANDS FOR THE OUTPUT OF THE ROW ALGORITHM AND B FOR THE OUTPUT OF THE COLUMN ALGORITHM.

The second test considers the hypothesis that the output algorithm A is closer to entirely dominating the output of algorithm B than vice versa. This can be cast into a test statistic based on the binary ϵ -indicator counting the number of runs where $I_\epsilon(A, B) < I_\epsilon(B, A)$, whose distribution under the null hypothesis is again the same binomial distribution as above. The values in Table IV-B.2 (middle) show that IBEA is particularly good at these pairwise comparisons, but none of the results is statistically significant under the given $\alpha = 5\%$ after the Bonferroni correction.

In the third test we investigate how far each algorithm is from a reference set, as measured by the ϵ -indicator. The reference set is the union of the output of all four algorithms as the best-known approximation achievable when starting from the given initial population. As the second argument of the ϵ -indicator is now the reference set, which is the same for each algorithm, this reduces to the unary ϵ -indicator $I_\epsilon(A) \equiv I_\epsilon(A, R)$ where R is the reference set. The test statistic is given by the number of runs where $I_\epsilon(A) < I_\epsilon(B)$, and the hypothesis is that algorithm A is closer to the reference set in terms of the unary ϵ -indicator than B . The values of the test statistic in Table IV-B.2 reveal that SPEA2 achieves better indicator values than all other algorithms in most of the runs, but again, none of them is significant such that the null hypothesis could be rejected.




Fig. 12. Empirical distribution of the values of the binary ϵ -indicator for the 11 runs for each pair of algorithms.

While the above tests have been performed using paired samples, Figure IV-B.2 now shows the empirical distribution \hat{F} of the values of the ϵ -indicator for all pairs of algorithms without the additional pairing with respect to the same initial populations. While the empirical distributions look very similar for certain pairs of algorithms, such as NSGA-II compared to IBEA or ABEA, others seem to indicate significant differences. Most prominently, IBEA seems likely to be less distant from being better than ABEA than vice versa — the empirical distribution function $\hat{F}[I_\epsilon(\text{IBEA}, \text{ABEA})]$ stochastically dominates $\hat{F}[I_\epsilon(\text{ABEA}, \text{IBEA})]$ by a wide margin and is by itself relatively close to one. A similar, albeit much smaller stochastic dominance is visible for SPEA2 versus ABEA. All these observations correspond well to the results of the second test above, where the different initial population was taken into account. This suggests that the initial population does not have a significant effect on the distribution of indicator values obtained.

In summary, SPEA2 and NSGA-II are able to obtain significantly better hypervolume values than IBEA and ABEA. IBEA is particularly strong with respect to the binary ϵ -indicator, but the results are not significant given the small sample size and large number of tests. The unary ϵ -indicator can be considered as an indication of the robustness of an algorithm, as the value is given by the maximum of the indicator values of the individual binary comparisons. Here, SPEA2 again seems to have a slight advantage over all three other algorithms, which is remarkable given that this wasn't the case for the binary comparisons, especially when compared to IBEA. Finally, ABEA's particular monotonicity property with respect to the development of the approximation set does not seem to lead to better approximation set quality, which might be due to the relatively low number of objective function evaluations so that the algorithms haven't reached the state yet where convergence problems because of cycling are an issue.

V. CONCLUSIONS

As this paper has demonstrated, evolutionary algorithms are well suited for optimization tasks in the context of groundwater resource management: mainly due to the model complexity, but also because of the inherently multi-objective structure of the corresponding problems. Furthermore, a key advantage of evolutionary approaches is the flexibility regarding model extensions. Here, several future research directions arise:

- The proposed approach combines a robust global search strategy with some numerical approximation methods as well as heuristics; an important question is how to optimally exploit the available computational resources, e.g., by adapting the accuracy of the numerical approximation methods over time;
- Uncertainty is a fundamental property of the natural environment; taking uncertainty into account with both model and optimization algorithm would be the next steps towards more realistic applications;
- As the computational requirements of the application, especially for real scenarios, are demanding, a parallelization of the algorithm using, for example, a master-worker model for the creation and evaluation of new solutions, could be a valuable extension with respect to employment in practice.

ACKNOWLEDGMENTS

Tobias Siegfried was partially financed by the Alliance for Global Sustainability. Stefan Bleuler has been supported by the SEP program at ETH Zurich under the Project TH-8/02-2.

REFERENCES

- [1] T. Shah, D. Molden, R. Sakthivadivel, and D. Seckler, "The global groundwater situation: Overview of opportunities and challenges," International Water Management Institute, Sri Lanka, Tech. Rep., 2000.
- [2] M. S. Daskin, *Network and discrete location models, algorithms and applications*. New York [etc.]: Wiley, 1995.
- [3] P. R. Krugman, *The Self-Organizing Economy*. Cambridge, MA [etc.]: Blackwell Publishers, 1996.
- [4] B. Provencher and O. Burt, "The externalities associated with the common property exploitation of groundwater," *Journal of Environmental Economics and Management*, vol. 24, no. 1, pp. 139–158, 1993.
- [5] S. Bleuler, M. Laumanns, L. Thiele, and E. Zitzler, "PISA - a platform and programming language independent interface for search algorithms," *Evolutionary Multi-Criterion Optimization, Proceedings*, vol. 2632, pp. 494–508, 2003.
- [6] J. Bear, *Hydraulics of groundwater*. New York a.o.: McGraw-Hill, 1979, mcGraw-Hill series in water resources and environmental engineering.
- [7] K. Binmore, *Game theory and the social contract*. Cambridge, Massachusetts [etc.]: MIT Press, 1994.
- [8] E. Zitzler, "Evolutionary algorithms for multiobjective optimization: Methods and applications," Ph.D. dissertation, Swiss Federal Institute of Technology, Zurich, 1999, tIK-Schriftenreihe, NR. 30.
- [9] K. Deb, *Multi-objective optimization using evolutionary algorithms*. Chichester: John Wiley & Sons, 2001.
- [10] D. E. Dougherty and R. A. Marryott, "Optimal groundwater-management by simulated annealing," *Water Resources Research*, vol. 27, no. 10, pp. 2493–2508, 1991.
- [11] D. C. McKinney and M. D. Lin, "Genetic algorithm solution of groundwater-management models," *Water Resources Research*, vol. 30, no. 6, pp. 1897–1906, 1994.

- [12] L. L. Rogers and F. U. Dowla, "Optimization of groundwater remediation using artificial neural networks with parallel solute transport modeling," *Water Resources Research*, vol. 30, no. 2, pp. 457–481, 1994.
- [13] W. Wang and D. P. Ahlfeld, "Optimal groundwater remediation with well location as a decision variable - model development," *Water Resources Research*, vol. 30, no. 5, pp. 1605–1618, 1994.
- [14] M. Wang and C. Zheng, "Optimal remediation policy selection under general conditions," *Ground Water*, vol. 35, no. 5, pp. 757–764, 1997.
- [15] A. Harbaugh and M. McDonald, "Programmer's documentation for MODFLOW-96, an update to the U.S. Geological Survey modular finite-difference ground-water flow model," Tech. Rep., 1996, U.S. Geological Survey Open-File Report 96-486.
- [16] C. Zheng and G. D. Bennett, *Applied contaminant transport modeling*, 2nd ed. New York: Wiley, 2002, Chunmiao Zheng, Gordon D. Bennett.
- [17] M. Erickson, A. Mayer, and J. Horn, "The niched pareto genetic algorithm 2 applied to the design of groundwater remediation systems," in *Evolutionary multi-criterion optimization first international conference proceedings*, E. Zitzler, Ed., Zurich, Switzerland, 2001.
- [18] E. Zitzler and S. Künzli, "Indicator-based selection in multiobjective search." in *PPSN*, 2004, pp. 832–842.
- [19] T. Siegfried, "Optimal utilization of a non-renewable transboundary groundwater resource - methodology, case study and policy implications," Ph.D. dissertation, Swiss Federal Institute of Technology, ETHZ, 2004.
- [20] W.-H. Chiang and W. Kinzelbach, *3D-groundwater modeling with PMWIN a simulation system for modeling groundwater flow and pollution*. Berlin: Springer, 2001.
- [21] G. De Marsily, *Quantitative Hydrogeology: Groundwater Hydrology for Engineers*. Academic Press, 1986.
- [22] MathworksInc., *Statistics Toolbox Version 4.1*, Natick, MA, 2003.
- [23] A. W. Harbaugh, E. R. Banta, M. C. Hill, and M. G. McDonald, "Modflow-2000, the u.s. geological survey modular ground-water model – user guide to modularization concepts and the ground-water flow process," U.S. Geological Survey, Tech. Rep., u.S. Geological Survey Open-File Report 00-92.
- [24] A. Wright, *Foundations of genetic algorithms*. San Mateo, California: Morgan Kaufmann Publishers, 1991, vol. 1, ch. Genetic Algorithms for Real Parameter Optimization, pp. 205–218.
- [25] S. Bleuler, M. Laumanns, L. Thiele, and E. Zitzler, "Pisa – a platform and programming language independent interface for search algorithms," in *Evolutionary Multi-Criterion Optimization (EMO 2003)*, ser. LNCS, C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, Eds., vol. 2632/2003. Springer-Verlag Heidelberg, 2003, pp. 494 – 508.
- [26] M. J. L. Robin, A. L. Gutjahr, E. A. Sudicky, and J. L. Wilson, "Cross-correlated random field generation with the direct fourier transform method," *Water Resources Research*, vol. 29, pp. 2385–2398, 1993.
- [27] T. Siegfried and W. Kinzelbach, "A multiobjective discrete stochastic optimization approach to shared aquifer management: Methodology and application," *Water Resources Research*, vol. Vol. 42, no. 2, February 2006.
- [28] E. Zitzler and S. Künzli, "Indicator-based selection in multiobjective search," in *Parallel Problem Solving from Nature (PPSN VIII)*, X. Yao *et al.*, Eds. Berlin, Germany: Springer-Verlag, 2004, pp. 832–842.
- [29] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonesca, and V. G. da Fonseca, "Performance assessment of multiobjective optimizers: An analysis and review," *IEEE Transactions on Evolutionary Computation*, vol. 7, no. 2, pp. 117–132, 2003.
- [30] E. Zitzler, M. Laumanns, and L. Thiele, "SPEA2: Improving the Strength Pareto Evolutionary Algorithm for Multiobjective Optimization," in *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)*, K. Giannakoglou *et al.*, Eds. International Center for Numerical Methods in Engineering (CIMNE), 2002, pp. 95–100.
- [31] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II," in *Parallel Problem Solving from Nature (PPSN VI)*, ser. Lecture Notes in Computer Science Vol. 1917, M. Schoenauer *et al.*, Eds. Springer, 2000, pp. 849–858.
- [32] M. Laumanns, L. Thiele, K. Deb, and E. Zitzler, "Combining convergence and diversity in evolutionary multiobjective optimization," *Evolutionary Computation*, vol. 10, no. 3, pp. 263–282, 2002.