DTU Library

Optimizing Wellfield Operation in a Variable Power Price Regime

Bauer-Gottwein, Peter; Schneider, Raphael; Davidsen, Claus

Published in:
Ground Water

Link to article, DOI:
10.1111/gwat.12341

Publication date:
2016

Document Version
Publisher's PDF, also known as Version of record

Citation (APA):
Bauer-Gottwein, P., Schneider, R., & Davidsen, C. (2016). Optimizing Wellfield Operation in a Variable Power Price Regime. *Ground Water*, *54*(1), 92-103. https://doi.org/10.1111/gwat.12341

Groundwater

Optimizing Wellfield Operation in a Variable Power Price Regime

by Peter Bauer-Gottwein[1], Raphael Schneider[2], and Claus Davidsen[2]

Abstract

Wellfield management is a multiobjective optimization problem. One important objective has been energy efficiency in terms of minimizing the energy footprint (EFP) of delivered water (MWh/m^3). However, power systems in most countries are moving in the direction of deregulated markets and price variability is increasing in many markets because of increased penetration of intermittent renewable power sources. In this context the relevant management objective becomes minimizing the cost of electric energy used for pumping and distribution of groundwater from wells rather than minimizing energy use itself. We estimated EFP of pumped water as a function of wellfield pumping rate (EFP-Q relationship) for a wellfield in Denmark using a coupled well and pipe network model. This EFP-Q relationship was subsequently used in a Stochastic Dynamic Programming (SDP) framework to minimize total cost of operating the combined wellfield-storage-demand system over the course of a 2-year planning period based on a time series of observed price on the Danish power market and a deterministic, time-varying hourly water demand. In the SDP setup, hourly pumping rates are the decision variables. Constraints include storage capacity and hourly water demand fulfilment. The SDP was solved for a baseline situation and for five scenario runs representing different EFP-Q relationships and different maximum wellfield pumping rates. Savings were quantified as differences in total cost between the scenario and a constant-rate pumping benchmark. Minor savings up to 10% were found in the baseline scenario, while the scenario with constant EFP and unlimited pumping rate resulted in savings up to 40%. Key factors determining potential cost savings obtained by flexible wellfield operation under a variable power price regime are the shape of the EFP-Q relationship, the maximum feasible pumping rate and the capacity of available storage facilities.

Introduction

Groundwater forms the backbone of water supply systems in many countries. Previous studies have addressed optimal management of groundwater resources and wellfields, combining models of the physical system with optimization and control strategies (Mayer et al. 2002; Fowler et al. 2008; Tsai et al. 2008; Hansen et al. 2013a). Typically, wellfield management is a dynamic and multiobjective optimization problem. Often, management objectives include maximizing supply reliability and water quality while minimizing contamination risk and energy consumption.

[1]Corresponding author: Department of Environmental Engineering, Technical University of Denmark, 2800 Kgs, Lyngby, Denmark; pbau@env.dtu.dk

[2]Department of Environmental Engineering, Technical University of Denmark, 2800 Kgs, Lyngby, Denmark.

Received November 2014, accepted April 2015.

doi: 10.1111/gwat.12341

Pumping and conveyance of groundwater consumes significant amounts of energy. For example, in Denmark, where 98% of drinking water is supplied from groundwater sources, electrical power consumption for pumping and conveyance of groundwater is estimated as ca. 80 GWh per year (Hansen et al. 2012). A focus of previous work has been to optimize wellfield operation in order to minimize the energy footprint (EFP) of delivered water (e.g., Ahlfeld and Laverty 2011; Hansen et al. 2013b, 2012). In this article, we define EFP as the amount of energy consumed per unit of water delivered, given in units of MWh/m^3.

In many countries, power systems are evolving in the direction of deregulated markets. Market organization varies from country to country and the outcomes of restructuring processes have been described in several comprehensive reviews (e.g., Al-Sunaidy and Green 2006; Sioshansi and Pfaffenberger 2006). A relatively recent and important trend in many power markets is increased penetration of highly intermittent power sources such as solar and wind (e.g., Jónsson et al. 2010). In the absence of large-scale storage options for power, intermittent renewable power sources lead to large price fluctuations on the wholesale electricity market. On the Nord Pool market, which Denmark is a part of (Flatabo et al. 2003), prices fluctuate widely and can even become negative, for instance in very windy conditions.

Against this background, the objective of wellfield management with respect to power should be minimizing costs arising from electric energy consumption for pumping and distribution, rather than minimizing EFP. This is the rational economic objective from the perspective of the individual wellfield operator and is also optimal from a societal perspective: When the opportunity cost of electric energy is high (i.e., power price is high) a large EFP of delivered water is undesirable because power could be used beneficially elsewhere and is produced by expensive and typically carbon-intensive power sources at the margin. However, when the power price is low, high energy consumption due to high EFP and large amounts of water pumped does not hurt because power cannot be put to beneficial use otherwise and is produced mainly by clean renewable sources. The rationale behind flexible wellfield pumping is to use power when it is abundant and clean, thereby accepting a higher total energy consumption for groundwater pumping and delivery. The decision how much to pump now and the given water demand together determine how much will have to be pumped later. If prices are low now, then it makes sense to fill up the storage, while at high prices, the optimal decision is to stop the pumps and supply the users from storage.

The objective of this study is to quantify the cost savings which can be obtained from flexible wellfield management in a variable power price regime. We consider a highly simplified water supply system as depicted in Figure 1: To keep things simple, we only consider usage charges for electric energy in this paper and assume zero demand charges. Costs are thus proportional to the amount of energy used (in MWh) and independent of the power draw (in kW). A wellfield pumps water into a storage facility and a deterministic time-variable water demand is served from the storage facility. The wellfield is characterized by a relationship between the EFP of delivered water and the pumping rate (EFP-Q relationship). The EFP-Q relationship is determined offline either from direct observations or using a model that represents the various system components. The EFP-Q relationship should take into account all relevant management constraints and objectives other than power, such as contamination risk, water quality considerations etc. Operators of the wellfield are assumed to buy electric power to drive the pumps on the wholesale power market, i.e., power price is variable. The management problem consists of finding the optimal sequence of pumping decisions that minimizes the total cost for electric energy over a given planning period, while guaranteeing water demand fulfilment and respecting the maximum capacity constraint of the storage facility. We solve this problem with a Stochastic Dynamic Programming (SDP) algorithm that exploits the auto-correlation in the price time series. The performance of SDP decision rules is benchmarked against a constant-rate pumping policy that represents current wellfield management practice and an optimization run that assumes perfect foresight of future price signals and quantifies the maximum savings that can be achieved with a hypothetical perfect price forecast.

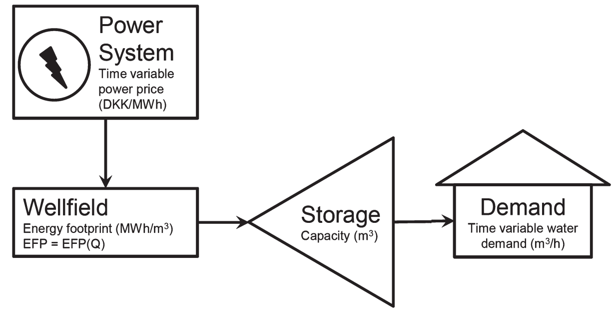

Figure 1. Conceptual system diagram. A wellfield is operated in a deregulated power system, where power price changes in time. The wellfield is characterized by an energy footprint function of the delivered water, which specifies the specific energy used for pumping and conveyance as dependent on the pumping rate. Pumped water can be stored in a storage facility (e.g., water tower), before being delivered to the public. Hourly water demand is deterministic and specified as a hard constraint.

EFP of Delivered Water

We use the Søndersø wellfield in Eastern Denmark as a case example to develop an EFP-Q relationship. This wellfield, which is located at approximately 55.773 latitude north, 12.362 longitude east, has been the subject of previous studies (Hansen et al. 2013b, 2012) which focused on the trade-offs between minimizing EFP of delivered water and minimizing contamination risk. In the absence of actual observations of EFP at different pumping rates, we use a combined well and pipe network model to simulate EFP over a range of pumping rates. Key data on the wellfield are summarized in Table 1 and an overview of wellfield layout is given in Figure 2. Along with a number of other wellfields, Søndersø wellfield conveys water to the Tinghøj reservoir, which serves the Copenhagen area. Tinghøj is a large reservoir with a storage capacity of about 80 times the average hourly demand (Københavns Kommune 2011). We use a similar wellfield simulation model as presented in the study by Hansen et al. (2013b); however, for reasons of computational efficiency, we replace the numerical groundwater model with a simplified analytical aquifer response function using the well-known Thiem solution (Kruseman and de Ridder 1990). The aquifer response function is needed to obtain the static head against which the pumps have to produce water. The reader is referred to Hansen et al. (2013b) and Hansen et al. (2012) for details on the wellfield and its components as well as details on the wellfield modelling approach. We used the pump setup A2 from the study by Hansen et al. (2013b) to develop the EFP-Q relationship. In this pump setup, all wells are equipped with variable-frequency pumps. Because the Tibberup wellfield is equipped with suction pumps, the Tibberup wellfield was assumed to be pumping at a constant rate (Table 1). Consequently, Tibberup pumping rates are not considered when computing EFP. However, Tibberup flows affect the EFP of water delivered by Søndersø wellfield, because pumped groundwater from Søndersø East, Søndersø West and Tibberup flows through the same pipes and the head loss in those pipes will thus depend on Tibberup flows.

To determine the EFP-Q relationship, we find the optimal set of pump frequencies at the 11 wells in the Sønderø East and Søndersø West wellfields, which guarantees delivery of a given amount of water while minimizing the EFP of the delivered groundwater:

$$\begin{gathered} \min EFP(\mathbf{f}) \\ s.t. \\ Q(\mathbf{f}) \geq Q_{req} \end{gathered} \tag{1}$$

In this equation, the vector **f** contains the 11 non-negative pump frequencies, which are the decision variables, Q is the wellfield pumping rate, which depends on the pump frequencies, EFP is the energy footprint of delivered water, which also depends on the pump frequencies and Q_{req} is the requested amount of water delivery. The wellfield model determines EFP and Q for a given set of pump frequencies, taking into account well drawdown, head losses in the pipe network and pump characteristic curves for the pumps deployed in the system (see Hansen et al. 2012, 2013a for details). The relationships between EFP and Q and the pump frequencies are both highly nonlinear.

This optimization problem is solved using a genetic algorithm (GA) implemented in the MATLAB software package for scientific computing. Alternatively, any global search algorithm for nonlinear and multidimensional search spaces can be used. Figure 3 presents the resulting EFP-Q relationship. Red dots in the left-hand panel indicate the optimal EFP solution found by the GA for a range of pumping rates. The black line shows a third order polynomial fitted to the data series, which is used as the EFP-Q relationship for the Søndersø wellfield in the SDP optimization. The right-hand panel of Figure 3 shows the optimal pump frequencies for each of the 11 individual wells as determined by the GA for each requested pumping rate. The frequencies are given relative to a baseline frequency of 60 Hz and are plotted on a relative scale between 0 and 1.

The only constraint enforced in the GA optimization is the constraint on minimum water delivery. To keep things simple, we did not consider potential additional constraints due to contamination risk, water quality considerations or ecological flow requirements. Such constraints and objectives can be important in real-world applications and can be incorporated in the wellfield model and/or the GA optimization. This will result in more complex wellfield simulation models and higher computational effort in the GA optimization. However, because the EFP-Q relationship is determined offline, this does not pose any problems for operational application.

Power Price and Water Demand

Hourly time series of spot market price of electricity in Eastern Denmark are publicly available from the Danish nonprofit grid operator energinet.dk. A time series for the years 2012 and 2013 was downloaded from their archive of market data and is shown in Figure 4. Power prices in this period are highly variable ranging from negative prices to more than 1000 Danish kroner (DKK) per MWh (1 DKK $\approx$ 0.134 Euro). The price time series was normalized by subtracting the mean price for each hour of the day and dividing by the standard deviation for the same hour of the day. The normalized price time series was modelled as a Markov chain. The time series was classified into five equally likely price classes, i.e., class breaks were set to the 0.2, 0.4, 0.6, and 0.8 quantiles of the empirical distribution of power price. Transition probabilities between price classes were determined separately for each hour of the day based on the observed market data. Synthetic power price time series were generated using the price classes and transition probabilities. The first two statistical moments of the synthetic power price distributions matched the moments of the empirical price distribution well.

Water demand was assumed to be deterministic. According to the data presented in Table 1, the average long-term abstraction from Søndersø wellfield is equivalent to the water consumption of about 120,000 people. Diurnal variation of water demand was obtained from the study by Miljøstyrelsen (2005) and is shown in Figure 5.

Stochastic Dynamic Programming

Traditionally, pumping rates at water supply wells are relatively constant. The distribution of pumping between different wellfields and between different wells within a wellfield is determined based on empirical decision rules and experience. Operators take into account criteria such as reliability, quality of the water pumped from different wells, and EFP. Pumping rates are adjusted every now and then (e.g., for maintenance and repair) but not on a daily or hourly time scale, because this requires advanced real-time control tools which are not commonly implemented at wellfields. Because wholesale power prices vary significantly on the hourly time scale (Figure 4), flexible management of pumping rates at this time scale can result in significant cost savings in wellfield operation and can contribute to balancing the power market, enabling higher penetration of renewable energy sources ("smart grid"). In our approach, we assume that pumping rates are adjusted at the hourly time scale based on economic decision rules derived from the results of the SDP optimization run, which minimizes the total cost. Total cost is defined as the sum of immediate and expected future cost. In every single hourly time step, the optimal pumping rate is determined based on the present

Table 1
System Characteristics

	Value	Source
Long-term average abstraction Søndersø West	153 m^3/h	Hofor Vand København (2013)
Long-term average abstraction Søndersø East	435 m^3/h	Hofor Vand København (2013)
Long-term average abstraction Tibberup	236 m^3/h	Hofor Vand København (2013)
Long-term average inflow to Tinghøj reservoir	2843 m^3/h	Københavns Kommune (2011)
Capacity Tinghøj reservoir	228,000 m^3	Københavns Kommune (2011)
Per capita water demand in the Copenhagen area (summed over all sectors)	164.1 l/day	Københavns Kommune (2011)

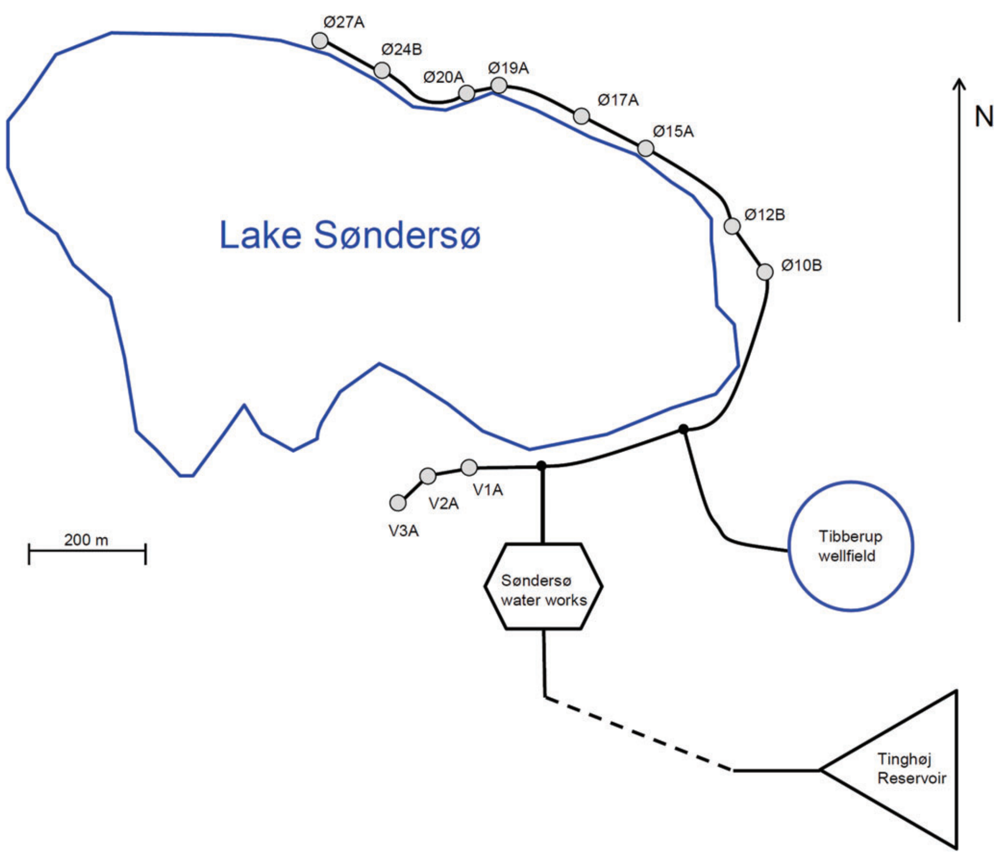

Figure 2. Map of Søndersø wellfield. Wells labelled Ø . . . belong to wellfield Søndersø East, wells labelled V . . . belong to wellfield Søndersø West.

wholesale power price, the amount of water in storage and the hour of the day. The immediate cost of pumping depends on the pumping rate and the present wholesale power price, while the expected future cost depends on the amount of water in storage at the end of the time step and, because of auto-correlation in the price time series, on the present wholesale electricity price. The auto-correlation in the price time series is modelled as a Markov Chain. Future cost is a stochastic variable because future power prices are unknown.

The wellfield management problem can be formalized as a stochastic dynamic program (e.g., Stedinger et al. 1984). SDP finds the series of hourly pumping rates that minimizes the total cost of groundwater pumping and delivery over a given planning period, subject to the constraint that deterministic hourly water demands must be fulfilled and respecting the finite capacity of the storage facility (Figure 1):

$$\min\left(Q \times EFP\left(Q\right) \times P \times \Delta t + EFC\left(S_{t+1}, P\right)\right)$$
$$s.t.$$
$$S_{t+1} = S_t + Q \times \Delta t - D_t \times \Delta t$$
$$S_t \leq K \qquad (2)$$

where Δt is the time step used in the SDP scheme (one hour in our case), Q is the pumping rate (m^3/h), EFP(Q) is the energy footprint (MWh/m^3), P is the power price on the market (DKK/MWh), S is the storage (m^3), D_t is the water demand (m^3/h), K is the storage capacity (m^3),

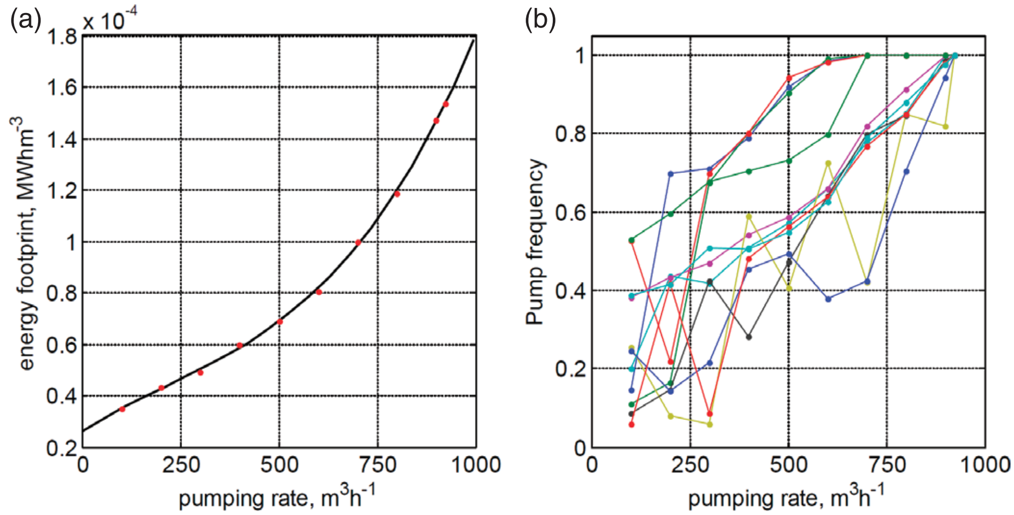

Figure 3. Energy footprint of delivered water as a function of total pumping rate for the Søndersø wellfield. (a) Red dots: Minimum energy footprint of delivered water for different pumping rates as determined by the genetic search algorithm. Black line: third-order polynomial fitted to the red dots. (b) Pump frequencies for the 11 individual wells determined by the genetic search algorithm in the optimal solution. Frequencies are given relative to a baseline frequency of 60 Hz.

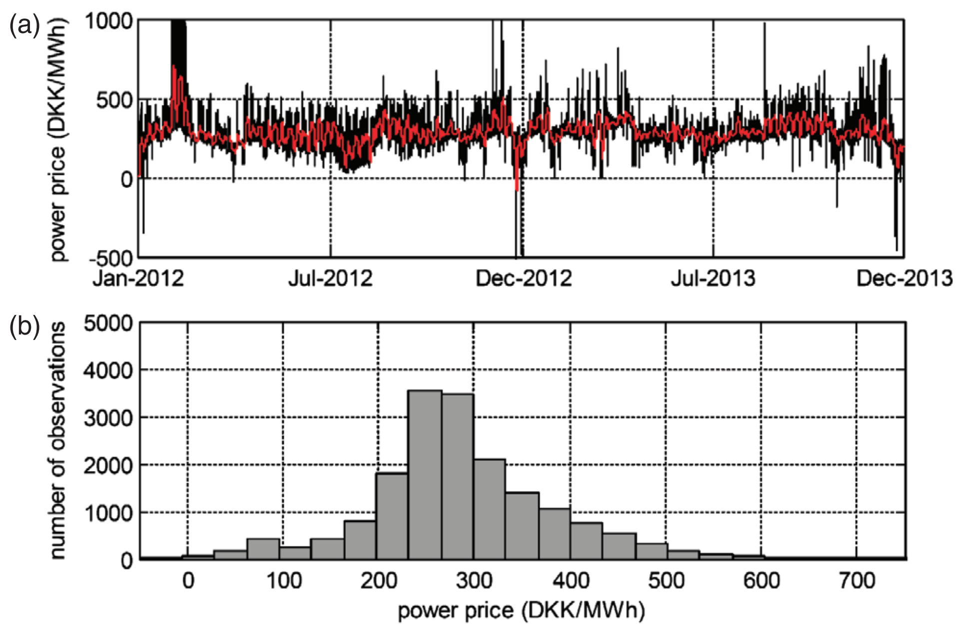

Figure 4. Power price on the Danish wholesale market (DK-East segment). (a) Time series of hourly power price for the years 2012 and 2013 (black). 48-h moving average of price time series (red). (b) Histogram of hourly power prices in the period 2012–2013.

and EFC is the expected future cost, which depends on the storage at the end of the time step and, because of autocorrelation in the price time series, on the present power price. All variables are constrained to be non-negative. The decision variables in this problem are the hourly pumping rates.

Because the EFC is a strictly convex function of storage, we can use a semidiscrete variant of SDP known as the water value method (Stage and Larsson 1961; Wolfgang et al. 2009). The water value method uses a series of linear constraints (or "cuts") to represent the EFC. The water value method operates with finite discrete time steps (hourly time steps in our case) and finite discrete storage volumes. The outer loop through the time steps proceeds backwards in time. For each initial storage volume in the inner loop (which runs through all discrete storage volumes), an optimization subproblem is solved. In this subproblem, the end storage for the time step is a continuous free decision variable. Figure 6 provides a graphical illustration of the procedure. Choosing the end storage fixes both the immediate cost (via the water balance equation and the EFP-Q relationship) and the expected future cost (via the cuts). The slopes and intercepts of the linear constraints on the EFC ("cuts") are determined recursively in

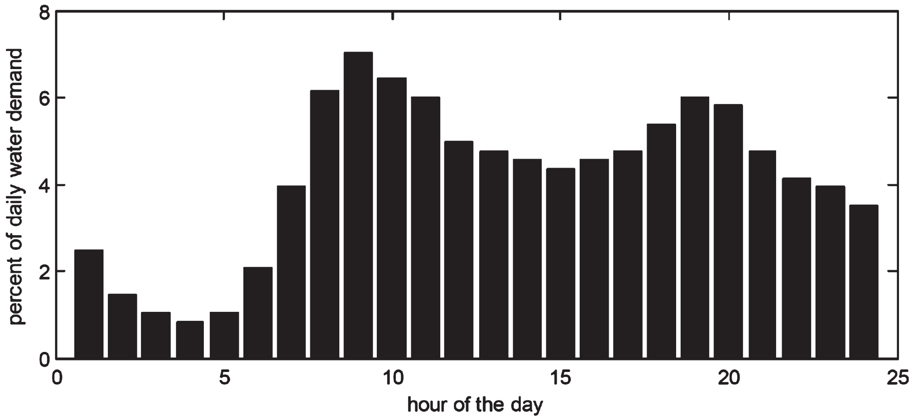

Figure 5. Variation of water demand over the course of the day in Denmark. The bars indicate the fractions of total daily water demand that fall into each hourly interval of the day.

the backward-in-time procedure from the objective function value and the shadow price on the water balance constraint resulting from the solution of Equation (2). The EFC for time step t is equal to the weighted average of the total costs obtained for time step $t + 1$ (note that these are already available because we proceed backwards in time), weighted by the transition probabilities of the Markov chain representing the power price variability. The problem (2) is transformed into a constrained optimization problem with two additional decision variables (expected future cost, EFC and end storage, S_{t+1}), which are linked by a series of linear constraints (Figure 6):

$$
\begin{gathered}
\min\left(Q \times EFP\left(Q\right) \times P \times \Delta t + EFC\right) \\
s.t. \\
S_{t+1} = S_t + Q \times \Delta t - D_t \times \Delta t \\
S_t \leq K \\
EFC \geq \alpha_1\left(S_{t+1} - S_1\right) + \beta_1 \\
EFC \geq \alpha_2\left(S_{t+1} - S_2\right) + \beta_2 \\
\ldots \\
EFC \geq \alpha_i\left(S_{t+1} - S_i\right) + \beta_i
\end{gathered} \tag{3}
$$

The storages S_i are the points on the storage axis evaluated at the previous stage of the recursive backward moving loop. The number of storage points (storage discretization) is determined by the trade-off between accuracy and computational load. Typically, future cost at the end of the planning period is set to zero for all storage levels, but the backward-in-time recursion can be initiated with any convex future cost function. The coefficients of the linear cuts, α and β, are weighted averages of shadow prices and objective values in the different price classes of the Markov chain that represents the power price variability. The weight of each price class is given by the corresponding transition probability of the Markov chain. Problem (3) is solved using an interior-point optimization algorithm for constrained nonlinear problems implemented in MATLAB and described by Byrd et al. (1999). Because problem (3) has to be solved for a number of points on the storage axis, for each power price class and for each hour of the day, the computational effort is significant and high-performance computing (HPC) facilities are required. The SDP algorithm is straightforward to parallelize because problem (3) can be solved independently for each point on the storage axis within one stage of the recursive loop. We used the HPC cluster of the Technical University of Denmark to solve the SDP. Each SDP run was solved in parallel on 12 individual cores. This resulted in run times ranging from a few hours to three days for the different SDP scenarios. We do not consider the significant computational load as a problem for operational applicability of the methodology, because the SDP runs are performed offline and produce decisions rules which are then used in operational management. Application of predetermined SDP decision rules in operational management requires negligible computational resources.

In the SDP optimization, the recursive backward loop was repeated until the shadow prices of the water balance constraint (the so-called water values) became approximately constant for every hour of the day, every price class and every storage level. Convergence was completed after less than 100 days in all scenarios. These equilibrium water values were subsequently used as decision rules in a forward-moving simulation run for the entire planning period, forced by the true observed time series of power prices. In the forward simulation run, equilibrium shadow prices and intercepts are used to solve Equation (3) repeatedly for every time step of the simulation period. We start from an arbitrary initial storage (which we chose as one half of the maximum storage) and then simulate sequential pumping decisions, considering in each time step, the present power price, storage level and hour of the day. The sum of all real costs occurring in this simulation run was compared to the total cost resulting from a constant-rate pumping scenario. Savings were computed as the relative difference between the total cost in the constant-rate pumping scenario and

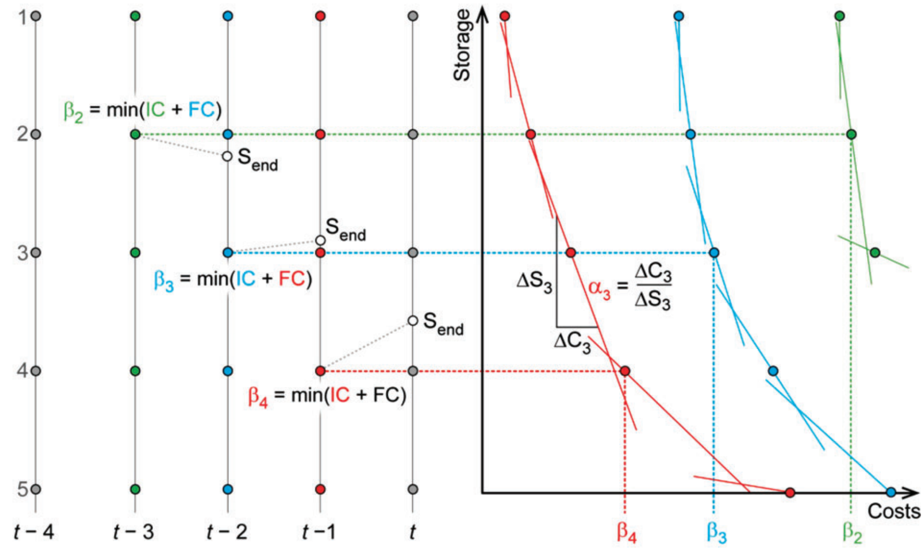

Figure 6. Illustration of the water value method. The figure illustrates the discretization of time and of the storage axis, the backward moving procedure (total cost of time step *t* becomes future cost of time step *t* − 1), and the construction of the future cost cuts. Note that a case with only one price class is shown, in order to keep the figure simple and clear.

the simulation run using equilibrium water values from SDP as decision rules, and are reported in percent.

Total operational costs resulting from the application of SDP decision rules were compared to two benchmarks: Constant pumping and perfect foresight. In the constant pumping benchmark, the assumption is that the wellfield pumps at a constant rate, equal to the average water demand (588 m^3/h). In the perfect foresight benchmark, the backward recursive loop was run for the entire price time series, resulting in a set of water values for each individual hour of the 2-year time series. System operation was subsequently simulated i.e., managers were assumed to exactly know future price signals and be able to adjust their decisions to those.

A complete listing of the MATLAB code used to perform the SDP optimization may be found in the additional Supporting Information of the online version of this article.

Results of SDP Optimization Runs

Prior to running scenario simulations, a convergence analysis was carried out for the baseline conditions in order to determine the maximum storage discretization in the SDP scheme that would still result in acceptable accuracy. Figure 7 presents the results of the convergence analysis. Savings are plotted as a function of storage discretization for two different storage capacities. Savings converge for storage discretization intervals below about half the average hourly water demand for both storage capacity scenarios. From this result, we conclude that a storage discretization equal to half the average hourly water demand results in the best compromise between accuracy and computational load and this discretization is subsequently used in all scenario runs. Figure 7 also compares the performance of SDP runs with 1, 3, 5, 7, and 9 price classes as well as the perfect foresight scenarios. Obviously, the perfect foresight run outperforms all SDP runs by a significant margin. The 7- and 9-price class runs do not result in significantly better performance compared to the run with 5 price classes. We concluded from this result that a set-up with 5 price classes offers the best compromise between performance and computational load and used 5 price classes in all subsequent scenario runs.

Figure 8 presents a set of equilibrium water values for a scenario with the baseline EFP-Q relationship, 5 price classes and a total storage capacity equivalent to 20 times average hourly water demand. Equilibrium water values range from zero to about 0.1 DKK per m^3. When deciding on the hourly pumping rate the equilibrium water value is compared to the present cost of pumping water into storage. If that cost is higher than the equilibrium water value, no water will be pumped and vice versa. Equilibrium water values are significantly different for the different price classes and track the diurnal variation of water demand. Figure 9 presents selected policy results for the baseline scenario. Storage variations in the storage facility are much more pronounced in both the SDP and perfect foresight runs as compared to the constant-rate pumping benchmark. In the flexible management runs (perfect foresight and SDP), the storage capacity is used actively to bridge periods of high power prices while the storage is re-filled during low-price periods. Significant variation of pumping rates and EFP occur over the course of the planning period. While the flexible management runs result in reduced cost compared to the constant-rate benchmark, the total amount of energy used for pumping and distribution of water is higher in the flexible pumping scenarios as compared to the constant-rate pumping scenario.

Table 2 presents an overview of the scenarios considered in this analysis. In all scenarios, the same Markov

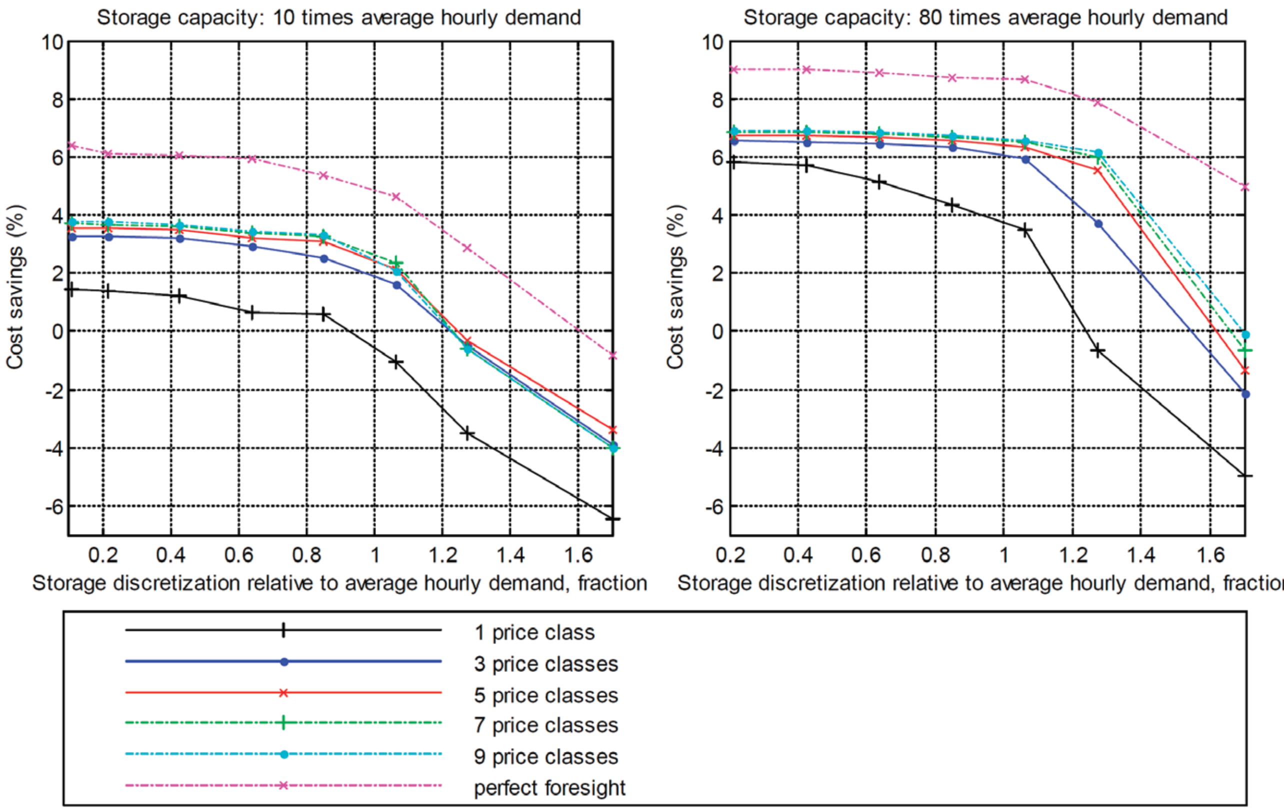

Figure 7. Results of convergence analysis. Cost savings (expressed as percent of costs at constant pumping rate) plotted against storage discretization (expressed relative to average hourly demand) for SDP runs with 1, 3, 5, 7, and 9 price scenarios as well as for the perfect foresight run. Upper panel: Storage capacity equal to 10 times average hourly water demand. Lower panel: Storage capacity equal to 80 times average hourly water demand.

chain was used to represent power price variability and the same historical price time series was used to evaluate performance of decision rules. The "Baseline" scenario represents conditions at Søndersø wellfield according to pump set-up A2 in the study by Hansen et al. (2013b). The EFP-Q relationship is shown in Figure 3 and the maximum pumping rate that can be delivered by the wellfield is 995 m^3/h. The "Larger Wellfield" scenario represents a situation where the wellfield has been expanded with new wells and pumps so that a larger amount of water can be pumped for a given EFP. The scenario assumption is that the wellfield has been scaled up by 50%, i.e., for any given EFP, 50% more water can be delivered than in the Baseline scenario. Consequently the maximum pumping rate is also 50% higher than in the Baseline scenario. The "Stronger Pumps" scenario represents a situation where the existing wells in the wellfield have been equipped with stronger pumps, i.e., the wellfield can deliver more water but only at a higher EFP. In this scenario, it is assumed that the maximum pumping rate from the wellfield is increased by 50% and the EFP-Q relationship is extrapolated to the higher pumping rates using the 3rd order polynomial from Figure 3. The scenario "Flat EFP" is a purely hypothetical scenario. A constant EFP is assumed, equal to the EFP at the average hourly pumping rate, independent of the actual pumping rate. The maximum pumping rate is still constrained to the same value as in the Baseline scenario. This scenario serves as a benchmark to show what cost savings could be achieved, if it was possible to reduce the slope of the EFP-Q relationship to close to zero, for instance by increasing wellfield capacity. The "Ideal World" scenario assumes a flat EFP-Q relationship and an unlimited pumping rate from the wellfield. In this hypothetical scenario, there are no physical constraints on the pumping and pumping rates can adjust entirely flexibly to the price signals coming from the power market. Finally, the "EFP offset" scenario represents a situation where the total EFP "from aquifer to storage" is considered in the SDP by increasing the baseline EFP by a constant and rate-independent offset. The offset was chosen such that the EFP at the average pumping rate matches the 0.2 kWh/m^3 from the study by Hansen et al. (2012). This resulted in an offset of 0.12 kWh/m^3.

Figure 10 presents savings as a function of storage capacity for all scenarios and both SDP and perfect foresight runs. The considered range of storage capacity is from 5 times average hourly pumping rate to 120 times average hourly demand (for comparison: Tinghøj capacity is 80 times average hourly demand). For all scenarios except the Ideal World scenario, the marginal value of additional storage capacity is decreasing and the savings curves flatten out for higher storage capacities. The lion's share of possible savings can be realized with a storage capacity of 20 times the average hourly demand in all

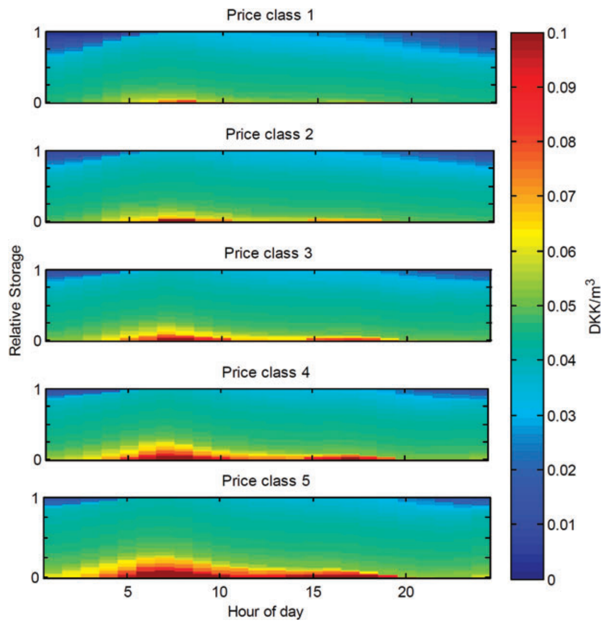

Figure 8. Equilibrium water values as dependent on relative storage level, hour of the day and present power price class. Price classes are sorted from lowest (class 1) to highest (class 5). Class breaks are located at the 20-, 40-, 60- and 80-percentiles of the empirical distribution of power price. The example shown here is for a total storage capacity equal to 20 times the average hourly water demand.

scenarios except the Ideal World scenario. For the Ideal World scenario, savings keep increasing roughly linearly as a function of the logarithm of the storage capacity. This is expected because in this scenario the pumping rate is unlimited, i.e., the storage can be completely refilled in the periods when power price is low or negative, irrespective of the size of the storage facility. However, this results in extremely high and unrealistic pumping rates during low-price periods. Relative savings in the "EFP offset" scenario are very close to relative savings in the baseline scenario. While the absolute power costs and the absolute magnitude of the savings are very different in the two scenarios, it appears that the fraction of the costs that can be saved by switching to a flexible pumping strategy are more or less equal for the two scenarios. For all scenarios, SDP performance is significantly worse than performance under perfect foresight. However, a large share of the perfect foresight savings can be realized using the SDP decision rules. Savings in the Larger Wellfield and Stronger Pumps scenarios are comparable in magnitude and significantly higher than in the baseline scenario.

Discussion

This study presented an approach to quantify potential cost savings that can be obtained by flexible wellfield management at the hourly time scale and provided first-order quantitative estimates of actual savings for a wellfield in Eastern Denmark. While we believe that these estimates are robust, a number of limitations and simplifications need to be highlighted and discussed.

The EFP-Q relationship for Søndersø wellfield was derived using a number of simplifying assumptions. The most significant simplification is the representation of the aquifer response to pumping with the stationary Thiem solution. In reality well drawdowns will adjust dynamically to time-variable pumping rates, while the Thiem solution assumes steady state between the well and a given radius of influence. This allows us to formulate the EFP as a function of the present pumping rate only, while in reality EFP will be a function of the present pumping rate and the recent history of pumping rates. In real-world applications, the EFP-Q relationship should preferably be determined from actual observations of EFP or from wellfield models that integrate a distributed groundwater model and do not rely on simplified aquifer response functions. In such an approach, other dynamic effects in the groundwater system (e.g., seasonal variations of recharge etc.) could also be taken into account and could be considered in management through seasonally varying water values.

In real-world applications it is also important to consider additional operational constraints when developing the EFP-Q relationship. These constraints could include limits on mixing ratios for water from specific wells due to groundwater quality variations, limits on pumping rates for individual wells or wellfields due to contamination risk or limits due to surface water-groundwater interaction and ecosystem flow requirements.

The EFP-Q relationship used in this study represents the EFP of the water from the aquifer to the Søndersø waterworks. However, ideally, the EFP from aquifer to storage should be used in the SDP runs. Because of lack of data, we were unable to quantify the EFP of water transport form Søndersø waterworks to Tinghøj reservoir. However, the dominant factor for pumping rate dependence of EFP is well drawdown. Water transport from Søndersø waterworks to Tinghøj by pressurized pipe is therefore expected to shift the EFP-Q relationship upwards by a more or less constant, rate-independent offset. Comparison of scenario results for the "Baseline" and "EFP offset" scenarios indicates that this will change the absolute magnitudes of costs and savings but savings relative to costs in a constant pumping strategy remain more or less unchanged.

In order for wellfield operators to take advantage of the cost savings generated by flexible wellfield pumping, power price signals from the wholesale electricity market must be passed on to wellfield operators. This implies that wellfield operators either act as buyers on the wholesale electricity market or that they have special agreements with the power utility. The first scenario would generate significant transaction costs for the wellfield operators, which might be larger than the savings that can be achieved. Therefore, a plausible application scenario would be that the power utility obtains some level of control of wellfield pumping rates in return for lower power prices charged to wellfield operators. Another key

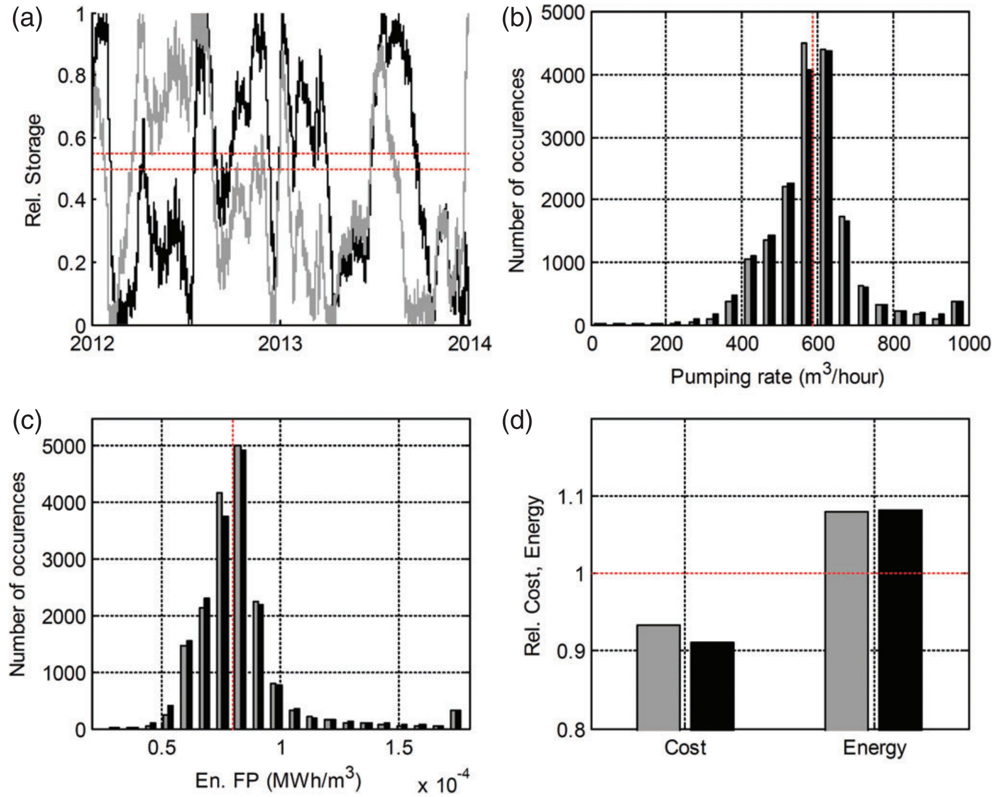

Figure 9. Policy results for the five price classes SDP run (gray) and the perfect foresight run (black). Total storage capacity: 80 times average hourly water demand. (a) Time series of total storage. The two red lines bracket the storage variation for constant pumping. (b) Histogram of hourly pumping rates. The red line indicates the average pumping rate. (c) Histogram of energy footprints of delivered water. The red line indicates the energy footprint of delivered water at the average pumping rate. (d) Cost and energy consumption relative to the run with constant pumping rate. The red line indicates performance of the constant-rate pumping policy

assumption in the presented approach concerns the amount of information available at the time the decision is taken. We assume that hourly decisions are taken based on present power price and that no additional information on future prices is available to managers except the statistical auto-correlation of price signals. In real-world applications, managers may have access to price forecasts or, depending on market organization, may even know prices for some time in advance. In such a situation, larger savings can probably be obtained compared to the SDP model. However, this is not expected to change the order of magnitude of possible savings, because savings will always be bounded by what is achievable under perfect foresight of future prices and the SDP model already achieves about two thirds of perfect foresight savings for most scenarios (Figure 10).

Besides the financial advantages for individual wellfield operators flexible wellfield management also has a desirable stabilizing effect on the power market. This is in line with the ideas of the smart grid, which are extensively discussed in the power systems literature. The basic concept is that power demands should be made adaptive

Table 2
Overview of Scenarios

Scenario	EFP-Q Relationship	Maximum Pumping Rate (m^3/h)
Baseline	Søndersø, see Figure 3	995
Larger wellfield	Baseline relationship scaled by 1.5. The amount of water that can be extracted for each EFP is 1.5 times the amount in the baseline scenario	1493
Stronger pumps	Baseline relationship extrapolated to 1.5 times the baseline maximum rate	1493
Flat EFP	Constant EFP equal to EFP at average pumping rate in the baseline scenario	995
Ideal world	Same as Flat EFP	No limit
EFP offset	Baseline EFP + 0.12 kWh/m^3	995

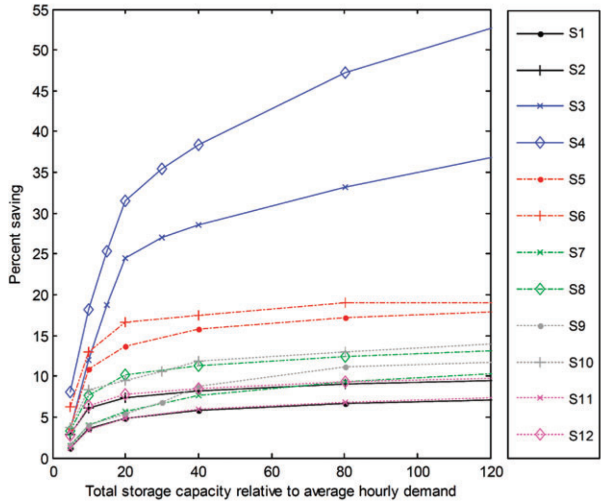

Figure 10. Cost savings (in percent of cost at constant pumping) vs. total storage capacity (plotted relative to average hourly water demand) for different scenarios. S1, baseline scenario, SDP; S2, baseline scenario, perfect foresight; S3, ideal world scenario, SDP; S4, ideal world scenario, perfect foresight; S5, flat EFP scenario, SDP; S6, flat EFP scenario, perfect foresight; S7, larger wellfield scenario, SDP; S8, larger wellfield scenario, perfect foresight; S9, stronger pumps scenario, SDP; S10, stronger pumps scenario, perfect foresight; S11, EFP offset scenario, SDP; S12, EFP offset scenario, perfect foresight.

or elastic to the extent possible so that power is consumed when it is abundant (i.e., prices are low) and power demand is low when power is scarce (i.e., prices are high). Along with other initiatives (e.g., smart household appliances, fridges, dishwashers), flexible wellfield scheduling can contribute to increased power demand elasticity and support increased penetration of clean renewable power sources in the power market. While the total power consumption of wellfields is relatively small, wellfield management is centralized at a few waterworks and adaptive power demand management is therefore much easier to implement than, for instance, for household appliances.

A last important limitation of the presented approach is the so-called curse of dimensionality: SDP methods are limited to less than a few state dimensions because the computational effort increases exponentially with the number of states. In our problem, the relevant state dimensions are the storage facilities, so problems with more than 2–3 coupled storage facilities in the water supply system will test the limits of what is computationally feasible today. It is clear that the water supply system shown in Figure 1 is highly simplified. Most real-world systems will have several inter-connected storages. A promising technique to handle more complex systems with a higher number of inter-connected storages is stochastic dual dynamic programming (SDDP; Pereira and Pinto 1991), which has been successfully applied to multireservoir problems in river basin planning. Alternatively heuristic search algorithms such as evolutionary algorithms (Nicklow et al. 2010) may provide effective solution strategies for complex water supply systems.

Conclusions

This study presented an approach to quantify cost savings obtainable from flexible wellfield management at the hourly time scale in a variable power price regime. Quantitative estimates of potential savings indicate that with present infrastructure, electric energy costs could be reduced by about 7%. To put this into perspective, we can estimate annual absolute savings for all of Denmark. The total groundwater abstraction in Denmark is about 392 million m^3 per year (Thorling 2009). With an average EFP of 0.2 kWh/m^3 (Hansen et al. 2012) and an average power price of 287.5 DKK/MWh, the total cost is 23 million DKK per year and a 7% saving would be equivalent to about 1.61 million DKK per year. Infrastructure upgrades could increase the possible savings as shown in the scenario runs presented in this study. In the hypothetical "Ideal World" scenario, where no infrastructure constraints limit flexibility, savings could be increased to about 35%, which is equivalent to 7 million DKK per year for all of Denmark. The key factors determining the magnitude of the savings which can be achieved by managing wellfields flexibly are the shape of the EFP-Q relationship, the maximum wellfield pumping rate and the available storage capacity. An equally important aspect of this study is the potential contribution of flexible wellfield pumping to the smart grid. Wellfield power consumption, albeit small, can contribute to increased penetration of renewable power sources in the power system because it is managed centrally and can be flexibly scheduled to make total power demand more elastic.

Acknowledgments

The authors would like to thank three anonymous reviewers for their insightful comments, which greatly improved this manuscript.

Authors' Note: The authors do not have any conflicts of interest or financial disclosures to report.

Article Impact Statement: Groundwater pumping consumes energy and flexible management of wellfields can help balance the power market under increasing penetration of renewables.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. MATLAB code listing for the water value method.

References

Ahlfeld, D.P., and M.M. Laverty. 2011. Analytical solutions for minimization of energy use for groundwater pumping. *Water Resources Research* 47: W06508. DOI:10.1029/2010WR009752.

Al-Sunaidy, A., and R. Green. 2006. Electricity deregulation in OECD (Organization for Economic Cooperation

and Development) countries. *Energy* 31: 769–787. DOI:10.1016/j.energy.2005.02.017.
Byrd, R., M. Hribar, and J. Nocedal. 1999. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization* 9: 877–900.
Flatabo, N., G. Doorman, O.S. Grande, H. Randen, and I. Wangensteen. 2003. Experience with the Nord Pool design and implementation. *IEEE Transactions on Power System* 18: 541–547. DOI:10.1109/TPWRS.2003.810694.
Fowler, K.R., J.P. Reese, C.E. Kees, J.E. Dennis, C.T. Kelley, C.T. Miller, C. Audet, A.J. Booker, G. Couture, R.W. Darwin, M.W. Farthing, D.E. Finkel, J.M. Gablonsky, G. Gray, and T.G. Kolda. 2008. Comparison of derivative-free optimization methods for groundwater supply and hydraulic capture community problems. *Advances in Water Resources* 31: 743–757. DOI:10.1016/j.advwatres.2008.01.010.
Hansen, A.K., H.-J. Hendricks Franssen, P. Bauer-Gottwein, H. Madsen, D. Rosbjerg, and H.-P. Kaiser. 2013a. Well field management using multi-objective optimization. *Water Resources Management* 27: 629–648. DOI:10.1007/s11269-012-0125-y.
Hansen, A.K., H. Madsen, P. Bauer-Gottwein, A.K.V. Falk, and D. Rosbjerg. 2012. Multi-objective optimization of the management of a waterworks using an integrated well field model. *Hydrology Research* 43: 430. DOI:10.2166/nh.2012.142.
Hansen, A.K., H. Madsen, P. Bauer-Gottwein, D. Rosbjerg, and A.K.V. Falk. 2013b. Optimization of well field operation: case study of Søndersø waterworks, Denmark. *Journal of Water Resources Planning and Management* 139: 109–116. DOI:10.1061/(ASCE)WR.1943-5452.0000219.
Hofor Vand København. 2013. Teknisk baggrundnotat til VVM for Hofors regionale vandindvinding—Eksisterende forhold for værket ved Søndersø. Copenhagen: Hofor Vand Købvenhavn.
Jónsson, T., P. Pinson, and H. Madsen. 2010. On the market impact of wind energy forecasts. *Energy Economics* 32: 313–320. DOI:10.1016/j.eneco.2009.10.018.
Kruseman, G.P., and N.A. de Ridder. 1990. *Analysis and evaluation of pumping test data*. Wageningen, The Netherlands: International Institute for Land Reclamation and Improvement.
Københavns Kommune. 2011. Vandforsyningsplan 2012. Høringsudkast. Copenhagen: Københavns Kommune.
Mayer, A.S., C.T. Kelley, and C.T. Miller. 2002. Optimal design for problems involving flow and transport phenomena in saturated subsurface systems. *Advances in Water Resources* 25: 1233–1256. DOI:10.1016/S0309-1708(02)00054-4.
Miljøstyrelsen. 2005. Vandforbrug og forbrugsvariationer. Copenhagen: Miljøministeriet.
Nicklow, J., P. Reed, D. Savic, T. Dessalegne, L. Harrell, A. Chan-Hilton, M. Karamouz, B. Minsker, A. Ostfeld, A. Singh, and E. Zechman. 2010. State of the art for genetic algorithms and beyond in water resources planning and management. *Journal of Water Resources Planning and Management* 136: 412–432. DOI:10.1061/(ASCE)WR.1943-5452.0000053.
Pereira, M.V.F., and L. Pinto. 1991. Multistage stochastic optimization applied to energy planning. *Mathematical Programming* 52: 359–375. DOI:10.1007/BF01582895.
Sioshansi, F., and W. Pfaffenberger. 2006. *Electricity Market Reform: An International Perspective*. Amsterdam, The Netherlands: Elsevier.
Stage, S., and Y. Larsson. 1961. Incremental cost of water power. *Transactions of the American Institute of Electrical Engineers* 80: 361–364.
Stedinger, J., B. Sule, and D. Loucks. 1984. Stochastic dynamic programming models for reservoir operation optimization. *Water Resources Research* 20: 359–375.
Thorling, L. 2009. Grundvandsovervågning 2009—Status og udvikling 1989–2008. Technical Report Chapter 8. Copenhagen, Denmark.
Tsai, F.T.-C., V. Katiyar, D. Toy, and R.A. Goff. 2008. Conjunctive management of large-scale pressurized water distribution and groundwater systems in semi-arid area with parallel genetic algorithm. *Water Resources Management* 23: 1497–1517. DOI:10.1007/s11269-008-9338-5.
Wolfgang, O., A. Haugstad, B. Mo, A. Gjelsvik, I. Wangensteen, and G. Doorman. 2009. Hydro reservoir handling in Norway before and after deregulation. *Energy* 34: 1642–1651. DOI:10.1016/j.energy.2009.07.025.