Journal of Hydrology 531 (2015) 1040–1053

Contents lists available at ScienceDirect

Journal of Hydrology

journal homepage: www.elsevier.com/locate/jhydrol

Many-objective optimization and visual analytics reveal key trade-offs for London's water supply

Evgenii S. Matrosov [a,b,1], Ivana Huskova [b,1], Joseph R. Kasprzyk [c], Julien J. Harou [a,b,*], Chris Lambert [d], Patrick M. Reed [d,e]

[a] School of Mechanical, Aerospace and Civil Engineering, The University of Manchester, Manchester M13 9PL, UK
[b] Department of Civil, Environmental and Geomatic Engineering, University College London, Chadwick Building, Gower Street, London WC1E 6BT, UK
[c] Department of Civil, Environmental and Architectural Engineering, University of Colorado Boulder, ECOT 441, UCB 428, Boulder, CO 80309, USA
[d] Thames Water, Clearwater Court, Vastern Road, Reading, RG1 8DB, UK
[e] School of Civil and Environmental Engineering, Cornell University, 211 Hollister Hall, Ithaca, NY 14853-3501, USA

ARTICLE INFO

Article history:
Received 21 January 2013
Received in revised form 9 October 2015
Accepted 3 November 2015
Available online 10 November 2015
This manuscript was handled by Andras Bardossy, Editor-in-Chief, with the assistance of Michael Brian Butts, Associate Editor

Keywords:
Trade-off analysis and visualization
Multi-criteria decision-making
Water resources planning
Infrastructure system design
Evolutionary multi-objective optimization

SUMMARY

In this study, we link a water resource management simulator to multi-objective search to reveal the key trade-offs inherent in planning a real-world water resource system. We consider new supplies and demand management (conservation) options while seeking to elucidate the trade-offs between the best portfolios of schemes to satisfy projected water demands. Alternative system designs are evaluated using performance measures that minimize capital and operating costs and energy use while maximizing resilience, engineering and environmental metrics, subject to supply reliability constraints. Our analysis shows many-objective evolutionary optimization coupled with state-of-the art visual analytics can help planners discover more diverse water supply system designs and better understand their inherent trade-offs. The approach is used to explore future water supply options for the Thames water resource system (including London's water supply). New supply options include a new reservoir, water transfers, artificial recharge, wastewater reuse and brackish groundwater desalination. Demand management options include leakage reduction, compulsory metering and seasonal tariffs. The Thames system's Pareto approximate portfolios cluster into distinct groups of water supply options; for example implementing a pipe refurbishment program leads to higher capital costs but greater reliability. This study highlights that traditional least-cost reliability constrained design of water supply systems masks asset combinations whose benefits only become apparent when more planning objectives are considered.

1. Introduction

Population and economic growth drive increased water demand (Vorosmarty et al., 2000) while climate change may further increase stress on the water supplies in some regions (Milly et al., 2008). Water supply capacity expansions are being considered in many areas, especially in fast growing cities (McDonald et al., 2011). In such cases, water resource system planners are faced with choosing the most appropriate mix of proposed new supply and demand management options for their system. This study investigates the trade-offs revealed by a many objective optimization approach to selecting future Thames basin (UK) infrastructure options. The term "many-objective" refers to optimizing systems with 4 or more design objectives as introduced by Fleming et al. (2005). Both supply and demand management options are considered to meet demands forecasted to 2035. Optimal water supply portfolios are evaluated according to their performance across a range of measures (economic, engineered, and ecological). We show how incorporating a broader suite of objectives into the planning exercise reveals information that is hidden when only one or two objectives are considered. The trade-offs generated by the many objective optimization reveal that ecological, engineered and economic performance can be improved with relatively modest investments. Trade-off visualization shows how similar schemes may cluster in certain areas of the trade-off, or Pareto, space. Pareto optimality is defined as those solutions whose performance cannot be improved in any single objective without degrading their performance in one or more remaining objectives (Coello Coello et al., 2005). The set of all Pareto points is referred

* Corresponding author at: School of Mechanical, Aerospace and Civil Engineering, The University of Manchester, Manchester M13 9PL, UK. Tel.: +44 (0) 161 306 4589.
E-mail address: julien.harou@manchester.ac.uk (J.J. Harou).
[1] These authors contributed equally.

http://dx.doi.org/10.1016/j.jhydrol.2015.11.003
0022-1694/© 2015 The Authors. Published by Elsevier B.V.

to as the Pareto optimal set and when plotted constitute the Pareto frontier. Visualizing the trade-offs and how portfolio mixes are distributed throughout the Pareto-space gives planners valuable information about the diverse array of planning options that are available for the Thames system.

Optimization algorithms such as mathematical programming and dynamic programming have long been used to solve the water resource capacity expansion problem (Loucks et al., 1981; Loucks and van Beek, 2005; Mays, 2005; Revelle, 1999). Such optimization methods have known success but also have limitations including the difficulty of representing the non-linearities of simulations, the diversity of discrete options and their potential to mask important performance trade-offs for real systems (Woodruff et al., 2013). Water resource systems often use non-linear rules and are likely subject to nonlinear cost and benefit functions. Their complexity may mean that aggregation and simplification of performance measures are often required when using classical optimization methods in water management models. Often minimizing costs or maximizing economic benefits has been the sole objective with non-commensurable objectives translated into costs or benefits (Harou et al., 2009; Lund et al., 2006). When classical optimization methods address multiple objectives, the relative weightings of each of the objectives must be pre-assigned, or changed iteratively. In real systems planners seek to simultaneously minimize expenditures while maximizing performance criteria such as resilience, reliability and ecological benefits. Given the historical consensus view that the water planning problem is inherently multi-objective (Cohon and Marks, 1975; Haimes and Hall, 1977), it is critical to move beyond classical commensuration approaches that require a single common unit of measure (typically monetary). The interaction of multiple objectives with investment opportunities has been long argued in many fields (Brill et al., 1982; Major, 1969) and water is no exception (Maass et al., 1962). Linking a multi-objective evolutionary algorithm to a water resource management simulator can overcome many traditional limitations of water management modelling. Water resource system simulators are able to incorporate non-linearities and explicitly calculate system performance using multiple criteria without the need to translate non-commensurable metrics into a single monetary metric.

The addition of multiple performance objectives in the infrastructure portfolio problem can reduce the possibility of human decision biases (Brill et al., 1982) such as 'cognitive myopia' or 'short-sightedness' (Hogarth, 1981) which can negatively bias planning decisions. This typically occurs in low-dimensional problems when managers feel they have sufficient knowledge about their system's behavior but may in fact lack a full understanding of innovative possibilities (Woodruff et al., 2013). A second decision bias was described by Gettys and Fisher (1979) as "cognitive hysteresis", where decision makers' pre-conceptions limit their incorporation of new ideas in their formulations of the future. These biases can lead to under- or over-estimation of reliability risks as shown by Kasprzyk et al. (2009) where adding additional objectives and decision variables led to alternative solutions that met reliability requirements at lower cost. Kollat et al. (2011) demonstrate how adding objectives can change the objective space and decision makers' preferences about the system's performance. They show that considering only two objectives can result in "extreme" solutions located at the edges of the objective space where they fail to satisfy other decision relevant concerns. Fogel (1997) argues that heuristic global optimization techniques such as evolutionary algorithms (EAs) can help overcome our biases by discovering new solutions to new problems.

Evolutionary algorithms imitate the process of natural evolution and have strongly contributed to the water resources literature as reviewed by Nicklow et al. (2010). Evolutionary algorithms are heuristic search algorithms that mimic the biological process of natural selection to produce an approximation of the Pareto optimal solution space. Classical optimization methods may require simplifications to the problem structure (e.g. removing non-linearities) in order to find global optima; in contrast, evolutionary algorithms link directly with simulation models and often do not require any simplification of the problem during the solution process. Evolutionary algorithms are suitable for solving real-world problems which often exhibit nonlinear, discrete, non-convex and high-dimensionality characteristics (Reed et al., 2013). A detailed review of evolutionary algorithms can be found in Coello Coello et al. (2005).

Evolutionary optimization has been shown particularly suitable for multi-objective water management applications (Nicklow et al., 2010; Reed et al., 2013; Maier et al., 2014) when linked to non-linear simulation models. Simulators are often developed over decades by water management agencies that include customized performance metrics which become trusted measures to evaluate management alternatives. In this approach, simulation models evaluate the optimization model's objective function, which means the full flexibility and descriptive ability of simulation models is harnessed (Labadie, 2004).

As reviewed by Reed et al. (2013) our study falls within a rapidly growing body of water resources literature focused on evolutionary multi-objective optimization. More formally, we are seeking an approximation to the set of Pareto optimal solutions. Rapid and highly interactive visualization of 'Pareto-optimal' solutions including their corresponding design components is critical for understanding complex trade-offs for applications with large numbers of objectives. Many-objective visual analytics ('visual analytics' for short, Woodruff et al., 2013) refers to emerging software packages that facilitate this process. The approximations of Pareto optimal sets usually contain a large number of solutions that increases rapidly with the number of objectives considered. It is important not only to visualize these solutions in multi-dimensional space but also to be able to isolate promising solutions with adequate justification. As the number of trade-off dimensions increases, the role of visual analytics becomes more central to the design and decision making process. Many water related problems are complex and formulating such problems appropriately usually requires a continuous learning process and the exploration of multiple problem formulations (Kasprzyk et al., 2012; Zeleny, 2005).

Current planning regulations in the UK require water companies to demonstrate that their water resource system designs can meet future demand at least-cost. This is done using a least-cost planning approach (Padula et al., 2013) that does not directly consider important environmental and engineering performance metrics. This study addresses this limitation of the current framework by investigating the performance trade-offs inherent in designing the future Thames basin (UK) water resource system as revealed by many-objective optimization. The Thames basin is the most populous river basin in the UK and includes major urban centers such as London. Both supply and demand management options (proposed by the basin's water companies), with a wide range of capacities and impacts, are considered to meet demands in 2035. Optimal portfolios (mixes) of different schemes are evaluated according to their performance across a range of measures (economic, engineered, and environmental). The proposed approach contributes to an improved supply-demand planning process for English water companies where there is demand to improve the current approach (Defra, 2011).

Section 2 presents the Thames basin portfolio selection case study. The methods and optimization formulation is presented in Section 3 followed by results in Section 4 and a discussion in Section 5.

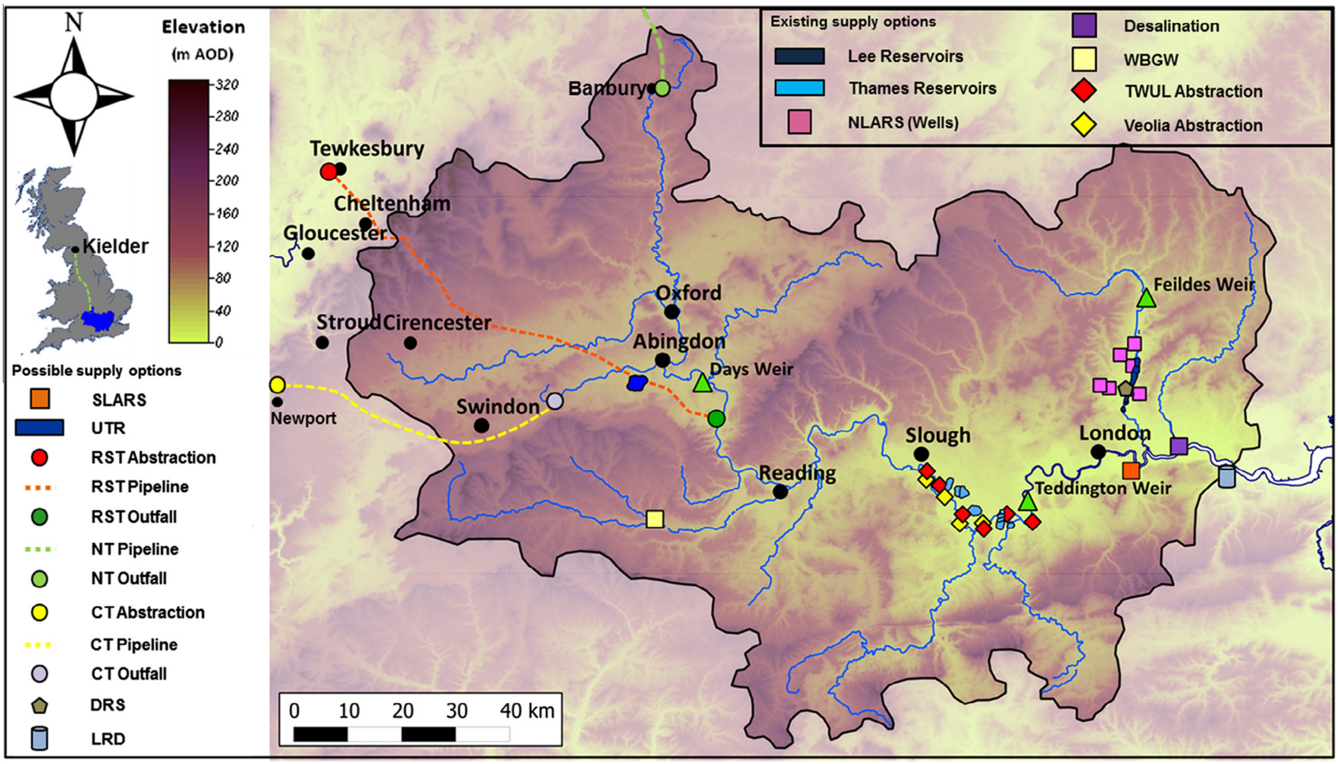

Fig. 1. Thames basin showing the river Thames and its major tributaries together with the existing and possible water supply options. Adapted from Matrosov et al. (2011).

2. Case study: London water supply planning

This study focuses on water supply source selection in the Thames basin of Southeast England (Fig. 1) considering demand levels projected for the year 2035. The basin (16,000 km^2) has over 12.5 million inhabitants and contains important cities including London. The United Kingdom's Environment Agency (EA) considers the region to be seriously water stressed. Parts of the basin are over-abstracted damaging the environment during low flow events (Environment Agency, 2008). The region has experienced six major droughts in the last 90 years (Marsh et al., 2007). Water stresses are expected to worsen resulting from an increase in demand and a decrease in supply resulting from climate change (Borgomeo et al., 2014; Walsh et al., 2015; Christierson et al., 2012; Sanderson et al., 2012). The population is estimated to increase by 2 million by 2026 (WWF-UK, 2008) while average baseline household water use is expected to remain the same (~150 l/person/day).

Surface water accounts for 60% of supply in the basin. London is the largest water user and relies on ten raised reservoirs fed from the Thames between Slough and Teddington Weir. Thirteen reservoirs supplement the storage in the Lee Valley with abstractions from the river Lee. Combined London-area storage is small (200 Mm^3). One surface water-groundwater conjunctive use scheme, the North London Artificial Recharge Scheme (NLARS), supplements supply during droughts and is recharged with treated water during wet periods. Additional groundwater is available from the West Berkshire Groundwater Scheme (WBGW) during dry conditions. A desalination plant along the Thames tidal estuary began operation in 2009.

These existing supply schemes are activated when thresholds on the Lower Thames Control Diagram (LTCD) (Fig. 2) are crossed. The LTCD is a function of real-time aggregate storage and the time of year and also controls when the minimum environmental flow on the Thames is reduced and demand restrictions come into effect. Demand restrictions are divided into four levels ranging from media campaigns to severe water rationing.

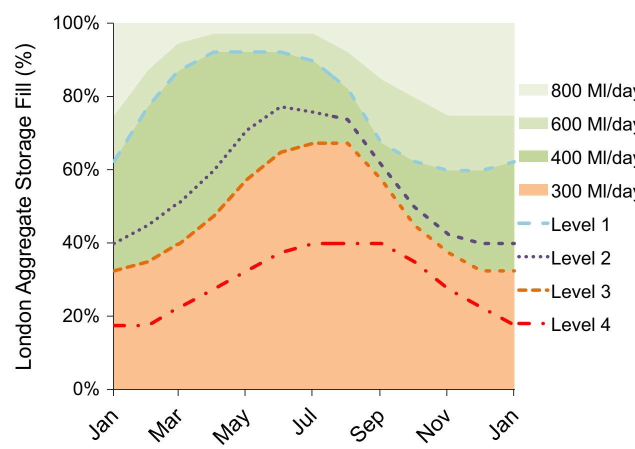

Fig. 2. Lower Thames Control Diagram (LTCD) showing the four demand restriction and minimum ecological flow thresholds as a function of London aggregate storage and the time of year. Adapted from Matrosov et al. (2011).

A combination of infrastructure expansion and demand management is likely necessary to maintain the supply-demand balance in the Thames basin. In their water resources management plan (Thames Water, 2010), Thames Water outlines many plausible supply and demand management options. In this study, we consider the seven main proposed supply options: three water transfers, a reservoir, a wastewater reuse scheme, a conjunctive use groundwater scheme and a brackish groundwater desalination plant. We also consider four demand management options: water efficiency improvements, increasing active leakage control, water mains pipes refurbishment and the installation of smart meters coupled with the introduction of seasonal tariffs. Please refer to Table 1 for a list of the supply and demand options considered.

Table 1
Possible future supply and demand options considered in this study.

Option	Description
Demand management options	
Active Leakage Control (ALC)	ALC refers to proactively seeking and fixing leaks in the water distribution system. Water companies consider levels of ALC implementation. Higher levels result in diminishing returns: higher costs are required for the same reduction in demand
Pipes refurbishment (Pipes)	Pipes includes replacement of water mains, communication pipes and supply pipes to reduce leakage in the distribution system
Enhanced efficiency improvements (EFI)	EFI includes water efficiency campaigns, retrofitting and household and commercial customer audit programmes
Installation of smart meters with seasonal tariffs (Meters)	Meters includes installing meters in properties. Seasonal tariffs can also be implemented and (Tariffs) are based on a summer/winter trail tariff implemented by Veolia Three Valleys Water (Veolia Water Central Limited, 2010). Tariffs effects on demand were calculated using the point expansion method (Griffin, 2006) to estimate the demand function at a known point on the demand curve assuming a constant price elasticity, ε, of −0.15 (Herrington, 2007)
Supply options	
Upper Thames Reservoir (UTR)	The UTR is a proposed reservoir which would release water into the Thames during times of low flow and provide constant supply to a neighboring area
River Severn Transfer (RST)	The RST is a proposed water transfer that would bring water from the River Severn to the Thames and a neighboring area during periods of low flow
Northern Transfer (NT)	The NT is a proposed water transfer that would bring water from Northern England to the Thames and a neighboring area during periods of low flow
South London Artificial Recharge Scheme (SLARS)	SLARS is a proposed conjunctive use groundwater recharge scheme what would function analogous to the existing NLARS
Deepham Reuse Scheme (DRS)	The DRS is a proposed planned indirect water reuse scheme in which a proportion of wastewater from Deepham's treatment plant would undergo additional treatment and be pumped into a surface storage reservoir during drought periods
Columbus Transfer (CT)	The CT is a proposed water transfer scheme that would bring water from the Dwr Cymru Welsh Water area to the Thames river and a neighboring area during periods of low flow
Long Reach Desalination (LRD)	LRD is a possible reverse osmosis treatment plant that would desalinate brackish groundwater leaking from the Thames Tideway and the Chalk aquifer underlying the Thames

3. Many-objective optimization formulation and implementation

3.1. IRAS-2010

We model the Thames water resource system with the open-source and computationally efficient Interactive River Aquifer Simulation IRAS-2010 (Matrosov et al., 2011) water resource system simulator. The model includes 51 nodes (reservoirs, aquifers, junctions, treatment and desalination plants, etc.) and 55 links (rivers, pipes, canals, water transfers) and has been shown to emulate the water infrastructure model maintained by the Environment Agency (Matrosov et al., 2011). The IRAS-2010 model focuses on the river Thames and major abstraction points and other supply and demand management schemes.

The IRAS-2010 Thames model incorporates the Lower Thames Control Diagram (LTCD) (Fig. 2), system operating rules and generates multiple performance measures. We model 2035 demands using 85 years of historical hydrology using a weekly time-step; by this we mean that we use 85 years of historical hydrology to represent hydrological conditions that could occur in the year 2035. We assume any one of these hydrological years could occur in 2035. This historical flow sequence is the same that water companies are required to use under the current accepted planning framework (Environment Agency, 2012). We use estimated demands for the year 2035 projected by the water companies and based on expected population growth and household water use (Essex and Suffolk Water, 2010; Thames Water, 2010; Veolia Water Central Limited, 2010). Important nodes in the IRAS-2010 Thames model are summarized in Table 2; refer to Matrosov et al. (2011) for further detail.

Table 2
Important nodes in the IRAS-2010 model.

Node	Description
LAS	London Aggregate Storage (surface reservoirs)
Teddington	Gauging station downstream of all the Thames abstractions
London	Demand node representing Thames Water's distribution input for London (2377 Ml/day)
Central_Abs	Demand node representing aggregate raw water abstractions from the Thames for the Central WRZ (239 Ml/day)
Southern_Abs	Demand node representing aggregate raw water abstractions from the Thames for the Southern WRZ (160 Ml/day)
Essex_BS	Demand node representing a bulk supply transfer of 105 Ml/day
SWOX	Demand node representing supply to the SWOX Water Resource Zone (WRZ) (24 or 48 Ml/day depending on capacity of a future reservoir or water transfer scheme)

3.2. Epsilon-Dominance Non-dominated Sorting Genetic Algorithm II (ε-NSGAII)

The IRAS-2010 simulator is linked to the Epsilon-Dominance Non-dominated Sorting Genetic Algorithm II (ε-NSGAII) (Kollat and Reed, 2006) evolutionary algorithm. ε-NSGAII was chosen for its search effectiveness and efficient parallel performance (Reed et al., 2013; Hadka and Reed, 2012; Kollat and Reed, 2006; Tang et al., 2006). The algorithm employs non-dominated sorting, ε-dominance archiving (Laumanns et al., 2002) and adaptive population sizing tournament selection. The ε-dominance archive sorts solutions based on the user specified levels of significant precision for the objectives (i.e., the minimum magnitude of change in the objectives that the user cares about). ε-NSGAII uses a series of connected runs between which the population size is adjusted with the introduction of new random solutions. Initially, the algorithm starts the search with a small number of candidate solutions. Over successive generations of each connected run, high quality solutions are passed into the epsilon-dominance archive. The archived solutions are injected into the population at the beginning of the next run and used to automatically adjust the search population size. A quarter of this population size is comprised of the archived solutions while the remaining three quarters are randomly generated solutions (Kollat and Reed, 2006).

3.3. Optimization formulation

The full problem formulation is described as a seven-objective optimization problem with 34 decision variables:

$$\mathbf{F}(\mathbf{x}) = (f_{\mathrm{LondSR}}, f_{\mathrm{StoTSRel3}}, f_{\mathrm{StoRes3}}, f_{\mathrm{EnvSI}}, f_{\mathrm{ResMin}}, f_{\mathrm{CAPEX}}, f_{\mathrm{OPEX}}) \tag{1}$$

$$\mathbf{x} = (x_i, Cap_i,)$$
$$\forall \mathbf{x} \epsilon \Omega$$
$$x_i \epsilon \{0, 1\} \quad \forall\, i \in N$$
$$\left(\sum_{i \in TransRes} x_i\right) \leqslant 1$$

subject to:

$$c_{level2} : \mathrm{AnnRel}, \quad 2 \geqslant 0.9 \tag{2}$$
$$c_{level3} : \mathrm{AnnRel}, \quad 3 \geqslant 0.95 \tag{3}$$
$$c_{level4} : \mathrm{AnnRel}, \quad 4 = 1.0 \tag{4}$$

The decisions include a binary variable x_i that represents the decision to activate supply or demand management option i at capacity *Cap*. Objectives include maximizing the engineering (reliability, resilience and minimum storage) and environmental (ecological flow) performance of the system while minimizing economic costs (capital and operating costs). Table 3 summarizes the formulation's decision variables: the supply and demand options and their capacities or capacity ranges. The *TransRes* decision includes the Upper Thames Reservoir (UTR), Northern Transfer (NT) and the River Severn Transfer (RST) supply options. The total sum of these mutually exclusive options must be lower than or equal to 1 to ensure that maximum of only one of these options can be selected at a time.

Table 3
Supply and demand options and their possible capacities included as decision variables in the study.

Option (i)	Capacity, Cap_i (Ml/day)	Exclusivity or dependence
Demand management options		
London ALC	0–50	None
London Pipes	165.1	None
London EFI	11.6	None
London Meters	88.7	None
Essex ALC	0–1.36	None
Essex Pipes	3.0	None
Essex EFI	0.7	None
Essex Meters	2.4	None
Central ALC	0–3.29	None
Central Pipes	4.8	None
Central EFI	2.8	None
Central Meters	10.2	None
Southern ALC	0–2.08	None
Southern Pipes	1.4	None
Southern EFI	1.7	None
Southern Meters	6.3	None
Supply options		
Upper Thames Reservoir (UTR)	133.5–267 to London (75–150 Mm^3) 20 or 40 to SWOX	Mutually exclusive to RST and NT
River Severn Transfer (RST)	267 to London 40 to SWOX	Mutually exclusive to UTR and NT
Northern Transfer (NT)	74 to London 8 to SWOX	Mutually exclusive to UTR and RST
South London Artificial Recharge Scheme	5–19	None
Deepham Reuse Scheme	25–95	None
Columbus Transfer	39 to London 14.8 to SWOX	None
Long Reach Desalination	15	None

3.3.1. Cost objectives

Costs include capital (CAPEX), fixed operating costs (FOPEX) and variable operating costs (VAREX) and are incurred by the supply and demand management options (SDOs). This study considers forecasted demand conditions for 2035 and uses a historical hydrological sequence to help evaluate supply and demand management options. Because there is no passage of time in this approach, and it does not propose a schedule of option implementation, discounting is not performed. This approach simply tests future options by simulating them over the entire historical period. The different supply options have different lifespans (*LS*). To take into account the varying lifespan of each option, the capital costs of each option are normalized with respect to its lifespan (80 years for reservoirs; 60 for transfers, reuse, pipe refurbishment and meters; and 25 for all other options) so that their capital costs are directly comparable to each other. This normalization involves dividing the capital expenditure of each option by the number of years it is meant to be active, thus providing an expression for an undiscounted annual capital cost. Because the investments are not scheduled and cannot be discounted, this cannot be considered equivalent to capital costs which will need to be financed, but still is a pragmatic financial metric essential to the optimization formulation. The normalized but undiscounted annual capital costs are aggregated via summation into a capital cost metric that is minimized:

$$\mathrm{Minimize}: f_{CAPEX} = \sum_{SDO} \frac{CAPEX_{SDO}}{LS_{SDO}} \tag{5}$$

New and existing infrastructure as well as some demand management options incur fixed and variable operating costs. Variable operating costs are correlated with energy requirements of the supply options; i.e., options with high energy use such as desalination have high operating costs. Fixed and variable operating costs are aggregated into a single (similarly undiscounted) operating cost metric (OPEX) that is averaged over the 85-years of historical flows representing plausible conditions in 2035 (N_s) and minimized:

$$\mathrm{Minimize}: f_{OPEX} = \frac{\sum_{SDO} VAREX_{SDO} + \sum_{SDO} FOPEX_{SDO}}{N_s} \tag{6}$$

3.3.2. Engineering objectives

The set of engineering performance objectives includes the reliability, resilience and minimum storage of the London aggregate storage node, LAS. We focused on this node because it provides the majority of surface water supplies to the London Water Resource Zone (WRZ) and when certain LAS storage volume thresholds are breached, water use restrictions in the London demand node (explained below) are imposed to reduce demand. The London reservoir reliability objective, $f_{\mathrm{StoTSRel3}}$, is a temporal reliability indicator (Kiritskiy and Menkel, 1982; Klemeš, 1969) and gives the ratio of the number of weeks the LAS was above failure level 3, S_t (Fig. 2) to the total number of simulated weeks, N_{ts}:

$$\mathrm{Maximize}: f_{\mathrm{StoTSRel3}} = \left(\frac{S_t}{N_{ts}}\right) \tag{7}$$

In addition to indicating storage performance, the London reservoir reliability objective gives a measure of how often LTCD level 3 restrictions (hosepipe and non-essential use bans) were implemented in the basin (see Fig. 2). Level 3 restrictions were chosen to be minimized because the non-essential use ban that corresponds to the level 3 restrictions is likely to cause severe disruption to the public. Hashimoto et al. (1982) base their resilience metric on the average duration of failure. We minimize the average duration of failure $\overline{FD_3}$ for LTCD level 3 failure events:

$$\mathrm{Minimize}: f_{\mathrm{StoRes3}} = \overline{FD_3} \tag{8}$$

We define a minimum London reservoir storage objective as the lowest storage level reached (in % of total capacity) by LAS over the historic time-series representing possible conditions in 2035. This value is maximized:

$$\text{Maximize}: f_{\text{ResMin}} = \text{MinLASVol} \tag{9}$$

When LAS drops below 22.5% of capacity pressure-related distribution problems occur in the network (Cookson and Weston, 2008).

Finally, the volumetric supply reliability gives an idea of how well the London demand was met. The London supply reliability objective is an annual volumetric reliability metric (Kiritskiy and Menkel, 1982; Klemeš, 1969) that gives the ratio of the total volumetric shortage to the total demanded over each year. The year with the lowest annual volumetric supply reliability is recorded:

$$\text{Maximize}: f_{\text{LondSR}} = \min\left(1 - \frac{\sum_{t=1}^{N_{\text{ts}}} \text{WS}_{\text{ts}}}{\sum_{t=1}^{N_{\text{ts}}} \text{WD}_{\text{ts}}}\right) \times 100 \tag{10}$$

where WS_{ts} is the time-step (weekly) (ts) flow shortage and WD_{ts} is the weekly flow demand target. The supply reliability gives an indication of the average deficit over the whole simulation run. The supply reliability is not calculated for the other demand nodes listed in Table 2. These nodes represent exports to other water companies and are not reduced during droughts.

3.3.3. Environmental objective

The environmental performance objective is a measure of how well the ecological flow of the Thames is maintained. An ecological shortage measure is adapted from the Shortage Index (SI) (Fredrich, 1975; Hsu et al., 2008):

$$\text{Minimize}: f_{\text{EnvSI}} = \frac{100}{N_{\text{ts}}} \sum_{\text{ts}=1}^{N_{\text{ts}}} \left(\frac{\text{WS}_t}{\text{WD}_t}\right)^2 \tag{11}$$

Higher SI values signal worse performance. Because of the square in the term, larger and longer shortages will have more effect on the SI index than a sequence of smaller and shorter shortages. SI increases when the residual flow at Teddington goes below 800 Ml/day (shaded zones in Fig. 2). For reference, the probability of exceedance of 800 Ml/day on the Thames at Kingston is 92% while for 300 Ml/day, the lowest allowable ecological flow at Teddington, it is 99% (Q92 = 800 Ml/day and Q99 = 300 Ml/day).[2]

3.3.4. Constraints

In their plan, Thames Water state that level 2 failures should not occur more often than once every 10 years. Level 3 failures should not occur more often than once every 20 years and level 4 failures should never occur (Fig. 2) (Thames Water, 2010). An occurrence reliability (Kiritskiy and Menkel, 1982; Klemeš, 1969) metric is used to impose these constraints which gives the ratio of the number of years that LAS did not experience a failure of level i, S_y, to the number or years in the time horizon, N_y.

$$\text{AnnRel}, i = \left(\frac{S_y}{N_y}\right) \tag{12}$$

The algorithm implements a constraint based tournament selection operator over all solutions within the generation where solutions are successively compared in pairs with respect to their objective function performance. Feasible solutions are always preferred to infeasible solutions. In general, simulations that do not meet these constraints are considered infeasible and are not passed into the archive of the MOEA. However, if all solutions are infeasible, the constrained tournament selection promotes solutions with the smallest aggregate constraint violations (Deb, 2001; Kasprzyk et al., 2009).

3.4. Computational experiment

The ε-NSGAII generates its initial random population of candidate solutions composed of combinations of decision variables by exploiting uniform random sampling within the user specified ranges given in Table 3. These variables are then passed as input variables to the IRAS-2010 simulator. The simulation evaluates performance using 85 years of historical hydrology. The performance information is passed back to ε-NSGAII for computing objectives and constraints upon which the algorithm evaluates the fitness of the decision variables and applies its selection operator to select "better" individuals to reproduce. The variation operators – crossover and mutation – are subsequently applied to these individuals, where the former combines genetic information of two individuals (parents) while the latter perturbs a genetic code of a single individual (parent) to create new individuals (children) for the next generation of decision variables. This represents one generation of the heuristic search process. The operator parameters such as the probability of crossover and mutation are user defined. Both of these operators when applied within real-coded genetic algorithms to multi-objective continuous problems have been proven to perform well (e.g. Deb and Kumar, 1995; Kollat and Reed, 2006). The parameter values for these operators were chosen based on recommendations of previous work that applied the ε-NSGAII algorithm to a multi-objective problem (Kollat and Reed, 2006; Kasprzyk et al., 2009). Even though the selection decision variables are in the final results binary (Eq. (1)), the algorithm treats them as real variables, i.e. generates a real value between 0 and 1. This value is then rounded for further analysis to determine if an option was selected or not (i.e. if the value is greater than or equal to 0.5, it is rounded up to 1, otherwise it is rounded down to 0). Often consistent rounding or truncation of real operators has been shown to outperform classical binary crossover (Nicklow et al., 2010) and many water resource applications demonstrate the importance of real-valued mating/mutation operators when focusing on continuous or mixed integer optimization (Bayer and Finkel, 2004; Kollat and Reed, 2006; Yoon and Shoemaker, 2001). We ran the algorithm for 70,000 function evaluations based on a visual assessment of the convergence and time-varying diversity of the evolving solutions. The initial population size was set to 72 and the algorithm operator parameters were chosen according to previous study recommendations (Kasprzyk et al., 2009; Kollat and Reed, 2007; Kollat et al., 2008). The algorithm parameters and objective epsilon values are summarized in Table 4. The epsilon values were set to capture the minimum level of precision to be used in distinguishing an alternative's performance in each objective. The population scaling factor directs the adaptive population sizing and represents the proportion of the population size at the beginning of each new run which consists of the ε-archived individuals. For instance, the population scaling factor of 0.25 means that if there are 50

Table 4
Algorithm parameters and objective epsilon values used in the case study.

Algorithm parameters	Value	Objective	Epsilon
Initial population size	72	f_{LondSR}	0.05
Population scaling factor (for injection)	0.25	$f_{StoTSRel3}$	0.01
Number of generations per run	250	$f_{StoRes3}$	0.10 (weeks)
Probability of crossover	1.0	f_{ResMin}	1.00 (Mm^3)
Probability of mutation	0.5	f_{EnvSI}	0.01
Distribution index for SBX crossover	15	f_{CAPEX}	0.5 M£
Distribution index for polynomial mutation	20	f_{OPEX}	0.15 M£

[2] Q92 and Q99 were calculated using daily gauging records from the National River Flow Archive (1883–2010).

archived solutions at the end of one run, the following run will begin with the population size of 200, where one quarter will consist of the archived solutions and the remaining 150 individuals will be generated randomly. This improves the search by directing it with previously evolved solutions and by adding new solutions to further explore the search space (Kollat and Reed, 2006).

Random number generation can strongly impact evolutionary search, particularly the randomly generated initial search population. To minimize random seed effects we ran the algorithm 10 times with different seed values. The results from each run are then sorted together to provide the best overall reference set based on the approach of Kollat et al. (2008). It should be noted that this reference set was found to be nearly identical to our original run results which indicates the search solutions are replicable and likely highly representative of the true Pareto optimal set.

4. Results

This section presents the results of the many-objective optimization formulation discussed in Section 3 for the Thames infrastructure portfolio design problem. Fig. 3 shows the approximation of the Pareto front generated by the multi-objective search process projected onto the two dimensional capital cost vs. London reservoir reliability (Eq. (5)) trade-off space. Each point represents a non-dominated solution, in our case, a portfolio of new supply and demand management measures. Many solutions show 100% London reservoir reliability. The upper left side and lower center of the figure is characterized by a steep cost to reliability gradient (i.e., small financial investments result in large reliability improvements). The cost vs. reliability subspace represents a classic lower dimensional view that has been the dominant focus of prior water resources systems design (Kundzewicz and Kindler, 1995; Lund and Israel, 1995; Rani and Moreira, 2009; Watkins and McKinney, 1997; Wurbs, 1993). This lower dimensional view shows the often discussed “flat surface” nature of the water supply

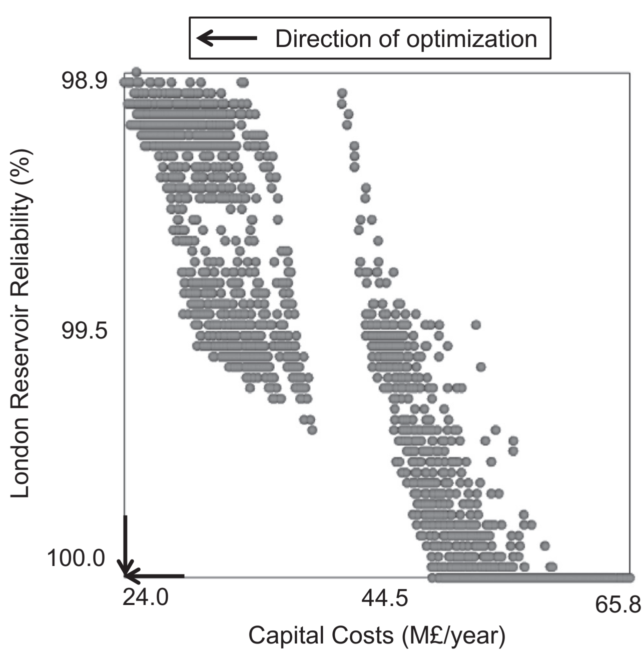

Fig. 3. Two-dimensional plot showing the trade-off between capital costs and London reservoir reliability. Each sphere on the plot represents a unique portfolio of supply and demand management measures. A steep trade-off exists between capital costs and reservoir reliability, i.e. relatively low investments can achieve large increases in reliability. Many solutions have perfect reliability. The arrow points towards the optimzation direction (optimal value of the objective).

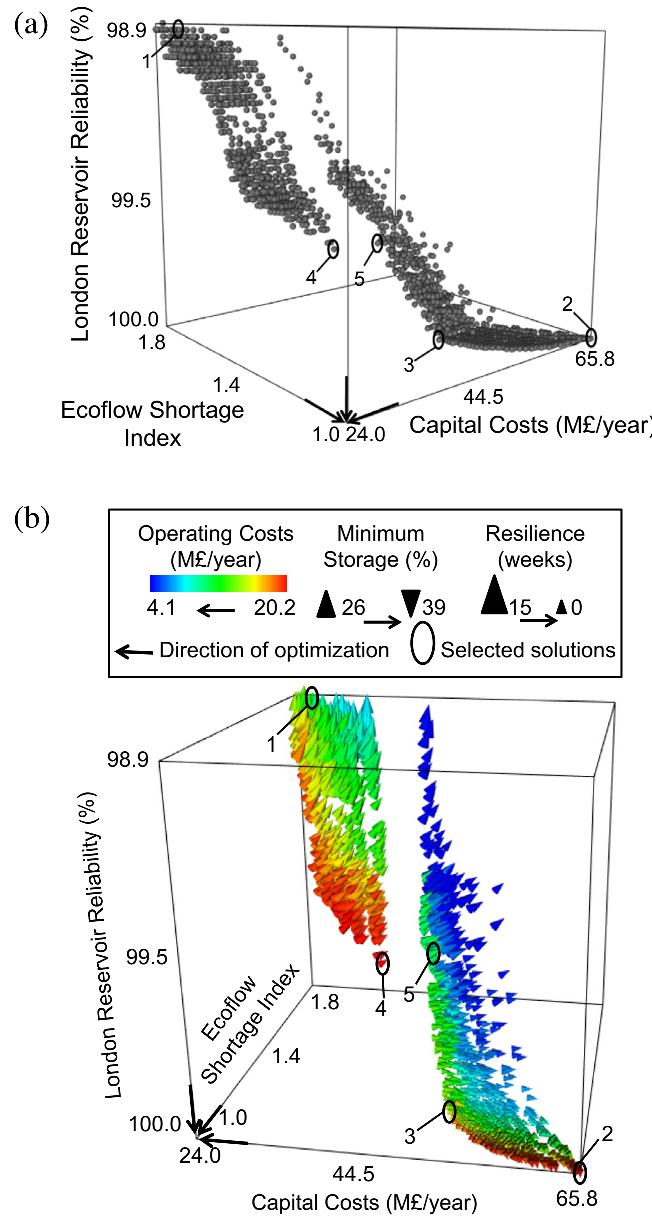

Fig. 4. Many-dimensional interactive plot showing the performance of the Pareto-approximate solutions. Plot A shows the London reservoir reliability, ecological flow and the capital cost performance metrics in the cardinal axes. B shows the same surface seen from a different angle. Additionally the color of the cones (glyphs) in plot B represents the operating costs metric while their orientations and size represent the minimum storage and resilience metrics, respectively. The arrows point towards the optimization direction (optimal value of the objective).

cost and reliability performance measures yielding many solutions with seemingly identical levels of performance (Loucks et al., 1981; Loucks and van Beek, 2005). As was discussed in our introduction, this lower dimensional view can negatively bias decision making hiding the broad array of water supply options discoverable in a many-objective formulation.

Fig. 4A shows the same Pareto-approximate solution set in three dimensions by adding the ecological flow index dimension. A benefit of many-objective visual analytics is that it facilitates rapid and interactive exploration of multiple views of the same high dimensional Pareto-approximate set. This view strongly distinguishes the performance of the solutions that appear as being analogous to one another in Fig. 3. Fig. 4B shows a rotated view of Fig. 4A and visualizes three additional dimensions by using cone color, orientation and size to represent the average operating cost, minimum reservoir storage and resilience metrics. These visualizations show the complex multi-dimensionality of the problem. The bottom right corner of Fig. 4B shows that the most capital cost intensive solutions also incur higher operating costs while achieving better performance in the engineering and environmental metrics. The Pareto-approximate space is divided into two 'fronts' where one front, at the top left of Fig. 4A and B, is characterized by green, yellow and red[3] cones (i.e. high operating costs) but is located on the lower end of the capital cost spectrum. The most reliable solution on this front achieves 99.7% London reservoir reliability. The reliability of the solutions that comprise the front on the right side of the plot ranges from 98.9% to 100%. If a decision maker wants to achieve 100% storage reliability, he/she must choose a solution from the right front. Up to 99.7% reliability, the two fronts also imply a trade-off between capital and operating costs; a similar level of reliability can be achieved by either increasing capital costs and reducing operating costs or vice-versa (e.g. by picking solutions with similar reliabilities from either the right or left fronts).

Five solutions are singled out for further analysis in Fig. 4A and B. Solution 1, the 'Lowest cost' solution, has the lowest undiscounted sum of annualized capital and operating costs in the Pareto-approximate set which satisfied the minimum service constraints (annual reliability constraints, Section 3.3.4). This portfolio is most similar to what would be recommended by the least-cost asset selection approach currently in use in England (Padula et al., 2013). Solution 2, the 'Highest cost' solution (undiscounted sum of annualized capital and operating costs) is Pareto-approximate because of its excellent performance in all but the two cost performance measures. Solution 3 corresponds to the lowest cost solution that resulted in perfect London reservoir reliability. Decision makers could pick solution 3, which we call the 'Cost efficient reliability' solution if they were only concerned about London reservoir reliability and cost. Solutions 4 and 5 show similar performance in all objectives except for the two cost objectives. Solution 4 is located on the left front and is more operating cost intensive than solution 5 (which is located on the right front) but has lower capital costs. Solutions 4 and 5 are named the 'Low capex compromise' and 'Low opex compromise' solutions, respectively.

Fig. 5 explores the portfolio composition of the Pareto-approximate solutions based on Fig. 4B. Color represents the activation of the Long Reach Desalination scheme; red represents solutions that activated the scheme while blue represents solutions that did not. Cones facing up represent solutions that include the Pipe refurbishment program in the London WRZ, a demand management option, while cones pointing down are solutions that do not. Transparency adds an additional dimension; opaque solutions build the Deepham reuse scheme while translucent solutions exclude it. Not all decisions are included in Fig. 5. Differently composed figures could show how other option choices lead to different areas of the efficient trade-off. Such figures can be used to show that certain options appear less favorably overall than others, e.g. in our analysis every solution in the Pareto-approximate space includes the Upper Thames Reservoir and none of the solutions include the River Severn Transfer.

Fig. 5 displays a mixture of both decisions (the option choices) and a subset of performance objectives. As noted by Tsoukias (2008), decision makers find the strict mathematical separation of decisions and objectives to be a false construct that can limit decision relevant insights. By visualizing decisions and objectives simultaneously Fig. 5 allows decision makers to discover how different mixes (portfolios) of supply and demand options can quantitatively affect performance. For example, the inclusion of the Pipe refurbishment program results in increased performance in the engineering and environmental metrics and an increase in capital costs without increasing operating costs. This relationship results in the two fronts described earlier; all solutions in the capital cost intensive front on the right of Fig. 4A and B include the Pipe refurbishment program while those in the left front do not. The solutions on the bottom right of the right front in the figures include the Pipe program, Long Reach Desalination and the Deepham reuse scheme resulting in large capital and operating costs. To contrast, the solutions at the top of the left front in the figures have lower capital costs and do not perform as well in the London reservoir reliability, resilience and environmental performance metrics. These solutions are less infrastructure intensive and implement fewer demand management schemes (e.g. they do not implement the Pipe refurbishment scheme). However as noted previously, despite lower capital costs, the solutions on the left front still incur high operating costs because they build the Deepham Reuse scheme. Fig. 5 shows that the discrete water supply decisions yield different clusters of portfolio options (i.e. the groupings created by London Pipes, Deepham Reuse, and Long Reach Desalination). The figure further shows that while solutions 4 ('Low capex compromise') and 5 ('Low opex compromise') achieve similar London reservoir reliability, minimum storage and environmental perfor-

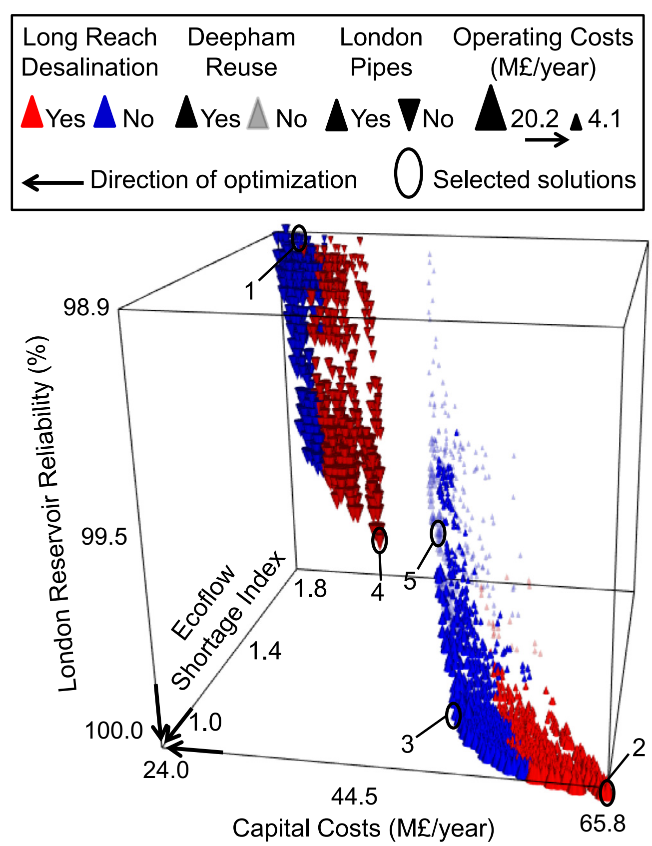

Fig. 5. Many-dimensional interactive plot showing groupings of different portfolio options and their performance. Opaque cones show solutions that include the Deepham reuse scheme, while translucent cones do not include the scheme. Red solutions include the Long Reach Desalination plant while blue solutions do not. Cones pointing up show solutions that implement the Pipes refurbishment program while solutions represented by cones pointing down do not implement it. The size of the cones represents the operating costs metric. The arrows point towards the optimization direction (optimal value of the objective). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

[3] For interpretation of color in Fig. 4, the reader is referred to the web version of this article.

Table 5
Decisions and objective values characterized by the five selected solutions.

Solution	1 (Lowest cost)	2 (Highest cost)	3 (Cost efficient reliability)	4 (Low capital costs compromise)	5 (Low operating costs compromise)
Objectives					
London supply reliability (%)	96.3	99.3	98.0	97.4	97.3
Capital costs (M£/year)	26,933	65,763	49,410	39,494	46,244
Operating costs (M£/year)	9500	19,919	13,780	20,227	10,745
Resilience (time-steps)	11.75	0	0	3.5	3.5
Reliability (%)	98.9	100.0	100.0	99.7	99.7
Ecological flow shortage index	1.79	1	1.25	1.45	1.43
Minimum reservoir storage (%)	26.5	38.7	32.5	30.2	30.3
Supply decisions objectives					
UTR capacity (Mm^3)/RST/NT	149	150	149	150	148
SLARS capacity (Ml/day)	23.7	24.0	22.8	24.0	20.1
Deepham reuse capacity (Ml/day)	62.1	93.0	41.4	94.8	0
Columbus transfer	No	Yes	Yes	Yes	Yes
Long reach desalination	No	Yes	No	Yes	No
Demand management decisions					
London ALC (Ml/day)	49.2	49.6	49.9	49.9	45.9
London Pipes	No	Yes	Yes	No	Yes
London EFI	Yes	Yes	No	Yes	No
London Meters (M)/Seasonal Tariffs (T)	MT	MT	MT	MT	MT
Central ALC (Ml/day)	2.9	3.0	3.0	3.2	0.7
Central Pipes	No	Yes	No	Yes	No
Central EFI	Yes	Yes	No	Yes	No
Central Meters (M)/Seasonal Tariffs (T)	M	M	No	M	M
Southern ALC (Ml/day)	1.5	1.9	2.0	1.6	1.1
Southern Pipes	No	Yes	No	Yes	No
Southern EFI	No	Yes	No	Yes	Yes
Southern Meters (M)/Seasonal Tariffs (T)	MT	MT	MT	MT	MT
Essex ALC (Ml/day)	0.8	1.2	0.9	1.3	0.7
Essex Pipes	No	Yes	No	Yes	No
Essex EFI	Yes	Yes	No	Yes	No
Essex Meters (M)/Seasonal Tariffs (T)	No	MT	No	MT	No

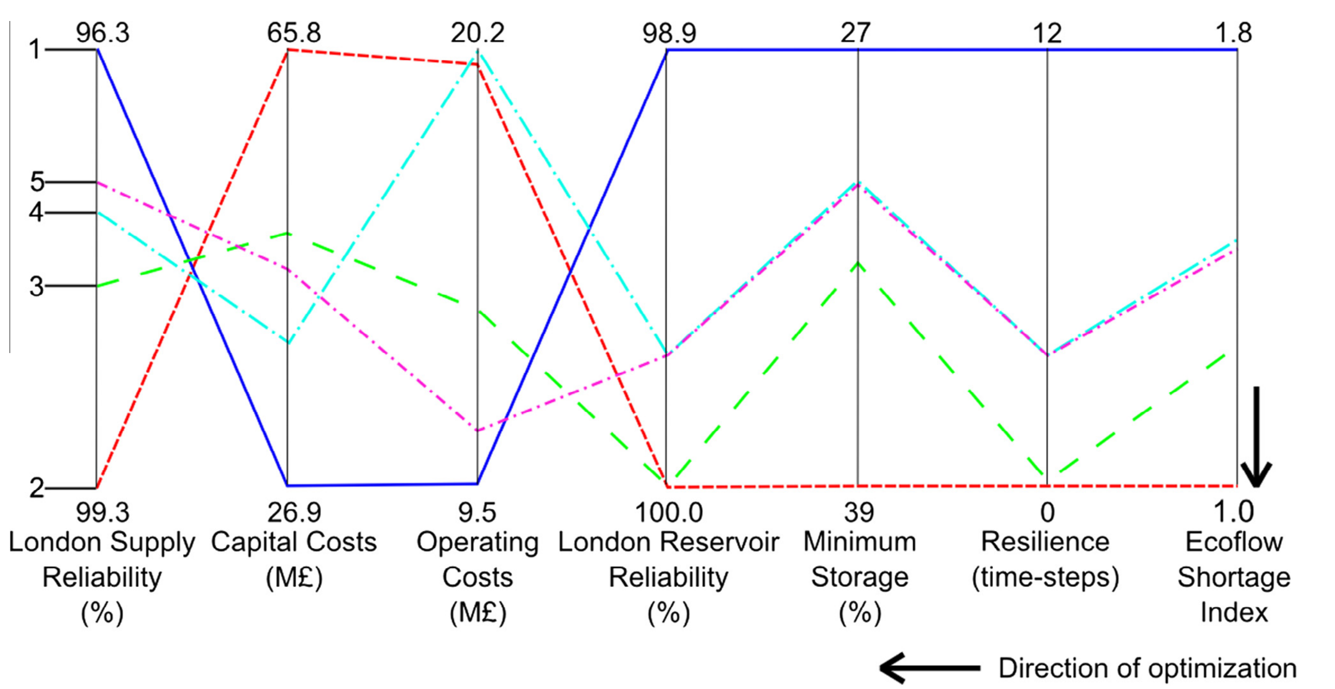

Fig. 6. Parallel plot of the five selected solutions. Each line on the figure represents the performance of one candidate Pareto approximate solution. The intersections of the lines with the vertical axes represent the objective performance. The arrow points towards the optimization direction (optimal value of the objective). Ideal performance would be a horizontal line at the bottom of the axes. Diagonal lines represent objective trade-offs.

mance they do so by implementing different options. Solution 4 builds the Long Reach Desalination and Deepham Reuse schemes resulting in high operating costs while solution 5 refurbishes the pipes in the London WRZ resulting in higher capital expenditure.

Table 5 summarizes the decisions and performance objectives of the five selected solutions.

The defining attribute of visual analytics (Keim et al., 2010) is the exploitation of multiple, linked views of high dimensional data.

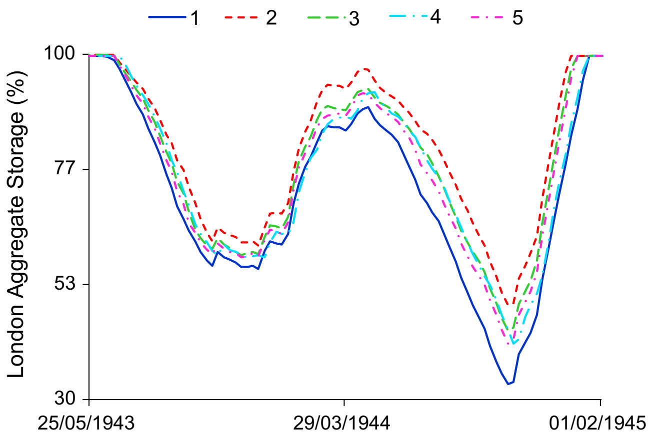

Fig. 7. Minimum storage performance of the five selected solutions during a major drought event.

The parallel axis plot (Inselberg, 2009) provides a highly scalable tool for exploring the trade-offs and performance differences of the five highlighted solutions illustrated in Figs. 4 and 5. The volumetric glyph plots discussed above show the strong geometrical context of the alternatives and when supplemented with the parallel axis plot in Fig. 6, the full suite of the Thames system's trade-offs come into perspective.

When interpreting Fig. 6 each vertical axis represents objective performance. Each line represents the many-objective performance of the five highlighted solutions. An ideal solution would be a horizontal line intersecting the bottom of every axis. Conflicts in the objectives are represented by diagonal lines between the respective objectives' vertical axes. The figure shows that trade-offs (and correlations) exist between the seven performance metrics. Solutions 1 and 5 perform well in the operating cost metric, but solution 1 has reduced performance with respect to the London reservoir reliability objective. Solution 2 has strong performance in all but the cost objectives. Relationships can be seen between other pairings and across all objectives. This helps decision makers visualize the consequences of only considering one or two objectives. As depicted in Figs. 4 and 5 solutions 2 and 3 have 100% London reservoir reliability (and therefore resilience). However, the parallel plot demonstrates that these solutions exhibit significantly different performance in the other performance metrics.

Solution 3 has good performance in most of the objectives. Despite the major infrastructure differences between solutions 4 and 5 (solution 4 includes the Long Reach Desalination and Deepham Reuse schemes while 5 includes the Pipes refurbishment program) they have similar performance in all metrics except for the two cost objectives as was seen in Fig. 4. The least infrastructure intensive portfolio, solution 1, displays the worst performance in the London supply reliability, resilience, London reservoir reliability, minimum storage and ecological flow metrics but has the best performance in the two cost measures. Conversely, the most infrastructure-intensive portfolio, solution 2, exhibits the best performance in all but the two cost metrics.

Visual analytic plots enable planners to 'browse' the Pareto-approximate front and introduce preferences for individual options based on non-optimized factors (ease of construction, land use, public opinion, etc.). The refurbishment of leaking water pipes in solution 5 vs. adding new desalination and reuse plants in solution 4 needs to be decided by strategic thinking on how these options help meet other less tangible goals (e.g. the relationship with regulators and client base).

The volumetric glyph and parallel axis plots show the performance objectives of each solution evaluated over the whole of a simulation run. The use of a simulation model in this optimization approach allows for direct performance comparison between any of the Pareto-approximate plans. Fig. 7 shows the simulated results for the London aggregate storage node during a major drought for each of the five selected portfolios. The plot serves as a reminder that each cone or point in the Pareto-approximate plots is backed up by a detailed and realistic system simulation. The 'Highest cost' portfolio (solution 2) sees the least drawdown of the London Aggregate Storage (LAS) node. The 'Lowest cost' solution 1 performs most poorly. In this early 1930s drought scenario, the existing storage becomes stressed signaling that solution 1 may suffer from under-investment. In the 1920s drought (not shown) the minimum storage for solution 1 goes down to 27% storage, which is near the 22.5% critical threshold. Solutions 3, 4 and 5 have similar minimum storage performance and may be good candidates for droughts such as the one seen in Fig. 7 as they achieve good performance without incurring high costs.

5. Discussion

This study demonstrates linking a water system simulator to multi-objective search to generate a diverse set of Pareto-approximate optimal water supply portfolios. The set is assessed using interactive visual analytic plots that reveal the trade-offs in the performance space and that map the composition of portfolios to the trade-offs. Below we discuss the innovations, limitations and implications of the water infrastructure selection approach presented here.

5.1. Improving system designs by increasing problem dimensions

Increasing the number of dimensions considered in scheme selection gives decision makers information they would not have if the problem were solved considering fewer factors of performance (e.g., a cost-only optimization). In our example the 'lowest cost' solution (selection 1) meets minimum service reliability requirements set by government regulators (i.e. model constraints in Section 3.3.4). Had decision makers only considered cost they would have likely chosen this portfolio. The parallel plot (Fig. 6) reveals this portfolio (solution 1) performs poorly in the London supply reliability, resilience, London reservoir reliability, minimum storage and ecological flow metrics. Many-objective visual analysis shows relatively small increases in cost can result in better performance across all of these metrics as is seen for the 'cost efficient reliability' portfolio, solution 3. Similarly, if London reservoir reliability were the sole selection criteria (with perhaps a maximum cost constraint), single objective optimization would lead to multiple optima as many solutions in the Pareto-approximate front display perfect London reservoir reliability (Fig. 3). Of all the solutions with 100% London reservoir reliability, all are non-dominated and have a range of varying performance in other objectives. Without the possibility to visualize these other dimensions, valuable gains could be missed and even the presence of multiple optima along one metric could be easily missed. Considering multiple objectives allows these solutions to be differentiated (Fig. 4). Figs. 3 and 4 show how solutions can have similar performance in some metrics but have diverging performance when seen using other dimensions.

Work by Woodruff et al. (2013) corroborates our findings and suggests how aggregated analyses of complex engineered systems can suffer from myopia and mathematical biases that lead to opportunity costs by ignoring key tradeoff alternatives between otherwise aggregated metrics. These aggregations occur for example in traditional cost minimization-only approaches (Padula et al., 2013) and cost-benefit analysis (Banzhaf, 2009).

5.2. Visualization

Many-objective visual analytics allows decision makers to survey the trade-offs between objectives and to distinguish the effects of individual supply or demand options within the Pareto-approximate set (Lotov and Miettinen, 2008). Interactive visualization of trade-offs in multiple dimensions is well suited for situations where stakeholders have diverse interests. For instance, an environmental regulator could be interested in how different portfolios impact the environmental flows downstream of abstraction sites while water companies could be interested in seeing how well portfolios meet service reliability requirements.

We demonstrate that it is important to exploit visual analytics to promote linked views of both performance objectives and investment decision variables simultaneously. We show how the Thames system's Pareto-approximate portfolios 'cluster' into distinct suites of water supply options. Visualizing these diverse groups of water supply plans in the performance space provides water managers with a rich perspective on key decision trade-offs and significant flexibility when choosing alternatives for further consideration. The many-dimensional visualization allows decision makers to consider the quantifiable performance metrics and navigate through them directed by factors not considered in the optimization (such as easiness of construction permits, land rights, etc.). Decision makers can quickly build a mental map of the consequences of including certain water supply schemes.

5.3. Uncertainty of future supply and demand

A limitation of the application described here is its consideration of one set of future conditions: it assumes historical inflows are representative of future plausible ones and that future demands are known. The historical record used in this study contains several severe droughts and therefore provides a useful stress test for future system designs. A 30-year historical hydrological record is used in the current planning framework English water companies use, and is substantially shorter than the ones used in this study. This deterministic study provides a baseline against which results from a future stochastic or multi-scenario optimization seeking robustness could be compared. An implementation accommodating multiple plausible futures (Matrosov et al., 2013a,b; Borgomeo et al., 2014) could incorporate the uncertainty of exogenous and endogenous factors into the planning approach.

5.4. Use of proposed approach for water utility system design

Current modeling to assist water supply-demand planning in England uses single objective least-cost optimization subject to reliability constraints (Padula et al., 2013). In the approach proposed here, the use of a water resource simulator allows performance metrics to be measured in diverse units familiar to stakeholders who may not agree on how or whether metrics should be monetized. As such the approach documented here is a contribution towards improved water planning for water utilities. Matrosov et al. (2013a,b) apply simulation-based water planning approaches on a UK case-study (Robust Decision Making and Info-Gap Analysis) and contrast them to the current regulator approved least-cost optimization approach. Borgomeo et al. (2014) present a risk-based framework that uses simulation to incorporate climate change projections into water resource planning. These simulation-based system design approaches allow considering engineering, economic and environmental performance in greater detail, but they do not consider all combinations of proposed interventions as economic optimization does. A few options with ranges of possible capacities (even if coarsely discretized) quickly lead to an exponential number of possible scheme portfolios to try. Simulation based approaches such as those applied by Matrosov et al. (2013a,b) can be criticized for choosing to evaluate in depth portfolios of options that are to some extent arbitrary defined. The approach presented in this study frees planners from having to choose *a priori* which portfolios of options (at fixed capacities) to evaluate; instead here the search for the most promising groupings of options and their capacities is automated.

If trusted simulators are used in the proposed analysis, and performance metrics used in the optimization have been defined with stakeholders (Herman et al., 2015), the Pareto-approximate solutions will likely be of interest to decision makers. The IRAS-2010 Thames basin simulator used in this study was shown to accurately emulate a similar simulation model used by the Environment Agency of England (Matrosov et al., 2011).

A detailed proposal and assessment of how this approach could be integrated into utility decision-making, collaborative river-basin planning and regulatory practice is beyond the scope of this article. The multi-objective infrastructure system portfolio design approach proposed here is more complex to implement in a regulated industry (as exists in England) than the current least-cost approach because the relevant performance metrics and relevant stakeholders vary somewhat by region and system. For each application a concerted effort would need to be made to define the most regionally relevant system goals, iteratively working with stakeholders to develop appropriate performance measures. A stakeholder-driven planning approach (Groves et al., 2013) using the proposed methods would benefit from stakeholders (a) defining system goals (metrics) to be optimized, (b) interactively using trade-off visualizations and (c) interacting in a deliberative forum to negotiate down to one or to a reduced number of preferred plans. Task (b) could for example include stakeholders adding *a posteriori* minimal acceptable performance thresholds and "brushing" (Kasprzyk et al., 2013) out solutions (erasing them from trade-off plots) that do not meet these negotiated preferences from task (c), thus reducing the number of solutions to consider. Stakeholder use of efficient trade-off analysis for collaborative decision-making and negotiation will benefit from further research. Readers can note that multi-criteria methods (Dodgson et al., 2009; Tanyimboh and Kalungi, 2009) to choose a candidate solution could supplement the proposed approach.

6. Conclusions

Water resource system and water supply planning are inherently multi-objective problems where decision makers must balance complex priorities such as costs, resilience and reliability, ecosystem services, etc. Single-objective planning such as least cost optimization gives planners only part of the picture when designing real systems where many aspects of system performance are relevant. Even if all system goals can and have been translated to one commensurate unit system (typically monetary), planners would lack the ability to understand the trade-offs embodied by different system designs. This study presented a water resources and supply system design optimization model with 7 simultaneous objectives: minimizing capital and operating costs while maximizing environmental performance and engineering performance metrics such as storage, resilience and reliability. The objectives were subject to regulatory supply reliability and environmental flow constraints. The optimization problem was solved by linking a water resource system management simulation model and a many-objective evolutionary optimization algorithm. The multi-objective search engine used the system simulator as the optimization function evaluator.

The approach was applied to identify promising designs for London's future water supply system assuming projected demands for 2035. Seven supply and four demand management options were considered in a many-objective capacity expansion optimization formulation. The output of the optimization was a set of Pareto-approximate (non-dominated) portfolios of supply and demand management schemes. Results showed that, out of the new options tested in our study (from the 2009 price review), the Upper Thames Reservoir (UTR) is always selected in all of the Pareto-approximate portfolios. We also showed that implementing demand management through pipe refurbishment in the London WRZ can reduce the need for a new desalination plant and reuse scheme. Visualizing time-series of detailed simulated results that underlie each point ('glyph') on the trade-off plots helps planners assess system responses to specific extreme events and helps prevent over- and under-investment.

State-of-the art many-objective visual analytics was used to explore the Pareto-approximate solution space which manifests as a multi-dimensional trade-off surface. These multi-dimensional interactive visual aids help analysts and decision makers see how individual supply and demand management options affect performance in each dimension. Portfolios which share certain schemes were seen in some cases to cluster in some parts of the decision space showing that choosing certain options leads to certain types of performance. Conversely, other parts of the Pareto-approximate front revealed that quite different portfolios had similar performance. Together the graphics underline the complexity of selecting interventions in complex human-natural systems when many metrics of performance are relevant, and also the richness of information communicable through a multi-objective search-based approach. The visual analytics graphics allow stakeholders and decision makers to assess trade-offs between objectives and show how different options and portfolios of options map to those trade-offs. The study showed that in cases where multiple optima are present in one dimension, other objectives can be used to differentiate between these solutions.

Acknowledgements

Evgenii Matrosov and Julien Harou were supported by the ARCC-Water project funded by the UK EPSRC (EP/G060460/1). Ivana Huskova is funded jointly by the UK EPSRC and Thames Water Utilities Limited. Patrick Reed was partially supported by U.S. NSF CAREER Grant CBET-0640443. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not reflect the views of the NSF.

Appendix A. Key to variables and abbreviations

Acronym	Definition
ALC	Active Leakage Control
AnnRel,*i*	Annual reliability occurrence metric – ratio of the number of years that LAS did not experience a failure of level *i*
Cap	Capacity
CAPEX	Capital costs
Central_Abs	Demand node representing aggregate raw water abstractions from the Thames for the Central WRZ
$C_{\text{level}i}$	Constraint, minimum annual reliability for LTCD level *i*
CT	Columbus Transfer
DRS	Deepham reuse scheme
EFI	Enhanced efficiency improvements
ε-NSGAII	Epsilon-Dominance Non-dominated Sorting Genetic Algorithm II

Appendix A (*continued*)

Acronym	Definition
f_{CAPEX}	Capital costs objective
$\overline{FD_3}$	Average duration of failure of LTCD level 3 failure events
f_{EnvSI}	Environmental objective
f_{LondSR}	London volumetric supply reliability objective
FOPEX	Fixed operating costs
f_{OPEX}	Operating costs objective
f_{ResMin}	Minimum reservoir storage objective (see MinLASVol)
$f_{StoRes3}$	Resilience objective (see $\overline{FD_3}$)
$f_{StoTSRel3}$	London reservoir reliability objective – the ratio of the number of time-steps (weeks) the LAS was above LTCD failure level 3
IRAS-2010	Interactive River Aquifer Simulation 2010
LAS	London Aggregate Storage
LRD	Long Reach Desalination
LTCD	Lower Thames Control Diagram
Pipes	Mains pipes refurbishment
Meters	Installation of smart meters
MinLASVol	Lowest storage level reached (in % of total capacity) by LAS over the entire modelled time horizon
MOEA	Many-objective evolutionary algorithm
NLARS	North London Artificial Recharge Scheme
NT	Northern Transfer
N_{ts}	Total time-steps (weeks) in the simulation time horizon
N_y	Number of years in the simulation time horizon
RST	River Severn Transfer
SDO	Supply or demand management option
SI	Shortage Index
SLARS	South London Artificial Recharge Scheme
Southern_Abs	Demand node representing aggregate raw water abstractions from the Thames for the Southern WRZ
S_t	Number of time-steps (weeks) the LAS was above LTCD failure level 3
S_y	Number of years that did not experience a supply failure
Tariffs	Activation of seasonal tariffs
ts	Time-step (week)
UTR	Upper Thames Reservoir
VAREX	Variable operating costs
WBGW	West Berkshire Groundwater Scheme
WD_{ts}	Time-step (weekly) flow demand target
WRZ	Water Resource Zone
WS_{ts}	Time-step (weekly) flow shortage
X_i	The decision to activate supply or demand management option *i*
$X_{TransRes}$	The $X_{TransRes}$ supply option includes the mutually exclusive UTR, NT and RST supply options which are represented by decision values of 1, 2 and 3 respectively

References

Banzhaf, H.S., 2009. Objective or multi-objective? Two historically competing visions for benefit-cost analysis. Land Econ. 85 (1), 3–23.
Bayer, P., Finkel, M., 2004. Evolutionary algorithms for the optimization of advective control of contaminated aquifer zones. Water Resour. Res. 40 (6), W06506.

Borgomeo, E., Hall, J.W., Fung, F., Watts, G., Colquhoun, K., Lambert, C., 2014. Risk-based water resources planning: incorporating probabilistic nonstationary climate uncertainties. Water Resour. Res. 508, 6850–6873.

Brill, E.D., Chang, S.Y., Hopkins, L.D., 1982. Modeling to generate alternatives – the HSJ approach and an illustration using a problem in land-use planning. Manage. Sci. 28 (3), 221–235.

Christierson, B.V., Vidal, J.-P., Wade, S.D., 2012. Using UKCP09 probabilistic climate information for UK water resource planning. J. Hydrol. 424, 48–67.

Coello Coello, C.A., Lamont, G.B., Van Veldhuizen, D.A., 2005. Evolutionary Algorithms for Solving Multi-Objective Problems Genetic and Evolutionary Computation Series Springer Science+Business Media. LLC, New York.

Cohon, J.L., Marks, D.H., 1975. Review and evaluation of multiobjective programming techniques. Water Resour. Res. 11 (2), 208–220.

Cookson, A., Weston, A., 2008. Thames Water Modelling Project: Model 1a Final Report, Environment Agency.

Deb, K., Kumar, A., 1995. Real-coded genetic algorithms with simulated binary crossover: studies on multi-modal and multi-objective problems. Complex Syst. 9, 431–454.

Deb, K., 2001. Multi-objective Optimization Using Evolutionary Algorithms. John Wiley, Hoboken, NJ.

Defra, 2011. Water for Life. Defra.

Dodgson, J.S., Spackman, M., Pearman, A., Phillips, L., 2009. Multi-criteria Analysis: A Manual. Department for Communities and Local Government. London.

Environment Agency, 2008. Water Resources in England and Wales – Current State and Future Pressures.

Environment Agency, 2012. Water Resource Planning Guidelines, Bristol.

Essex and Suffolk Water, 2010. Final Water Resources Management Plan 2010–2035, Durham, England.

Fleming, P., Purshouse, R., Lygoe, R., 2005. Many-objective optimization: an engineering design perspective. In: Coello Coello, C., Hernández Aguirre, A., Zitzler, E. (Eds.), Evolutionary Multi-Criterion Optimization. Lecture Notes in Computer Science. Springer, Berlin Heidelberg, pp. 14–32.

Fogel, D.F., 1997. Introduction. In: De Jong, K., Fogel, L., Schwefel, H.P. (Eds.), Handbook of Evolutionary Computation. IOP Publishing Ltd. and Oxford University Press.

Fredrich, A.J., 1975. Hydrologic Engineering Methods for Water Resources Development. Davis, California.

Gettys, C.F., Fisher, S.D., 1979. Hypothesis plausibility and hypothesis generation. Organizational Behavior and Human Performance 24 (1), 93–110.

Griffin, R., 2006. Water Resource Economics: the Analysis of Scarcity, Policies, and Projects. MIT Press, Cambridge, MA.

Groves, David G., Fischbach, Jordan R., Bloom, Evan, Knopman, Debra, Keefe, Ryan, 2013. Adapting to a Changing Colorado River: Making Future Water Deliveries More Reliable Through Robust Management Strategies. RAND Corporation, Santa Monica, CA, <http://www.rand.org/pubs/research_reports/RR242> Also available in print form.

Hadka, D., Reed, P., 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. Evol. Comput. 20 (3), 423–452.

Haimes, Y.Y., Hall, W.A., 1977. Sensitivity, responsivity, stability and irreversibility as multiple objectives in civil systems. Adv. Water Resour. 1 (2), 71–81.

Harou, J.J., Pulido-Velazquez, M., Rosenberg, D.E., Medellin-Azuara, J., Lund, J.R., Howitt, R., 2009. Hydro-economic modeling: concepts, design, applications and future prospects. J. Hydrol. 375, 334–350.

Hashimoto, T., Stedinger, J.R., Loucks, D.P., 1982. Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. Water Resour. Res. 18 (1), 14–20.

Herman, J., Reed, P., Zeff, H., Characklis, G., 2015. How should robustness be defined for water systems planning under change? J. Water Resour. Plan. Manage. 141 (10), 04015012.

Herrington, P., 2007. Waste Not, Want Not? Water Tariffs for Sustainability. WWF, Bristol.

Hogarth, R.M., 1981. Beyond discrete biases – functional and dysfunctional aspects of judgmental heuristics. Psychol. Bull. 90 (2), 197–217.

Hsu, N.S., Cheng, W.C., Cheng, W.M., Wei, C.C., Yeh, W.W.G., 2008. Optimization and capacity expansion of a water distribution system. Adv. Water Resour. 31 (5), 776–786.

Inselberg, A., 2009. Parallel Coordinates: Visual Multidimensional Geometry and Its Applications. Springer, New York, New York.

Kasprzyk, J.R., Nataraj, S., Reed, P.M., Lempert, R.J., 2013. Many objective robust decision making for complex environmental systems undergoing change. Environ. Modell. Software 42, 55–71.

Kasprzyk, J.R., Reed, P.M., Characklis, G.W., Kirsch, B.R., 2012. Many-objective de Novo water supply portfolio planning under deep uncertainty. Environ. Modell. Software 34, 87–104.

Kasprzyk, J.R., Reed, P.M., Kirsch, B.R., Characklis, G.W., 2009. Managing population and drought risks using many-objective water portfolio planning under uncertainty. Water Resour. Res. 45.

Keim, D.J., Kohlhammer, G.E., Mansmann, F., 2010. Mastering the Information Age – Solving Problems with Visual Analytics, Eurographics.

Kiritskiy, S.N., Menkel, M.F., 1982. Water Management Computations (in Russian), GIMIZ, Leningrad, USSR.

Klemeš, V., 1969. Reliability estimates for a storage reservoir with seasonal input. J. Hydrol. 7 (2), 198–216.

Kollat, J.B., Reed, P.M., 2006. Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design. Adv. Water Resour. 29 (6), 792–807.

Kollat, J.B., Reed, P.M., 2007. A computational scaling analysis of multiobjective evolutionary algorithms in long-term groundwater monitoring applications. Adv. Water Resour. 30 (3), 408–419.

Kollat, J.B., Reed, P.M., Kasprzyk, J.R., 2008. A new epsilon-dominance hierarchical Bayesian optimization algorithm for large multiobjective monitoring network design problems. Adv. Water Resour. 31 (5), 828–845.

Kollat, J.B., Reed, P.M., Maxwell, R.M., 2011. Many-objective groundwater monitoring network design using bias-aware ensemble Kalman filtering, evolutionary optimization, and visual analytics. Water Resour. Res. 47.

Kundzewicz, Z.W., Kindler, J., 1995. Multiple criteria for evaluation of reliability aspects of water resource systems. In: Simonovic, S.P.K.Z.R.D.T.K. (Ed.), Modelling and Management of Sustainable Basin-Scale Water Resource Systems. Iahs Publications, pp. 217–224.

Labadie, J.W., 2004. Optimal operation of multireservoir systems: state-of-the-art review. J. Water Resour. Plan. Manage. 130 (2), 93–111.

Laumanns, M., Thiele, L., Deb, K., Zitzler, E., 2002. Combining convergence and diversity in evolutionary multiobjective optimization. Evol. Comput. 10 (3), 263–282.

Lotov, A.V., Miettinen, K., 2008. Visualizing the Pareto Frontier. In: Branke, J., Deb, K., Miettinen, K., Slowinski, R. (Eds.), Multiobjective Optimization: Interactive and Evolutionary Approaches. Lecture Notes in Computer Science, pp. 213–243.

Loucks, D.P., Stedinger, J.R., Haith, D.A., 1981. Water Resources Systems Planning and Analysis. Prentice-Hal, Englewood Cliffs, N.J.

Loucks, D.P., van Beek, E., 2005. Water Resources Systems Planning and Management – An Introduction to Methods, Models and Applications. UNESCO, Paris.

Lund, J., Cai, X., Characklis, G., 2006. Economic engineering of environmental and water resource systems. J. Water Resour. Plan. Manage. 132 (6), 399–402. http://dx.doi.org/10.1061/(ASCE)0733-9496(2006)132:6(399).

Lund, J.R., Israel, M., 1995. Optimization of transfers in urban water-supply planning. J. Water Resour. Plan. Manage. – ASCE 121 (1), 41–48.

Maass, A., Hufschmidt, M., Dorfman, R., Thomas, H., Marglin, S., Fair, G., 1962. Design of Water-Resources Systems. Harvard University Press, Cambridge, Massachusetts.

Maier, H.R., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L.S., Cunha, M.C., Dandy, G.C., Gibbs, M.S., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D.P., Vrugt, J.A., Zecchin, A.C., Minsker, B.S., Barbour, E.J., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., Reed, P.M., 2014. Evolutionary algorithms and other metaheuristics in water resources: current status, research challenges and future directions. Environ. Modell. Software 62, 271–299.

Major, D.C., 1969. Benefit-cost ratios for projects in multiple objective investment programs. Water Resour. Res. 5 (6), 1174.

Marsh, T., Cole, G., Wilby, R., 2007. Major droughts in England and Wales, 1800–2006. Weather 62 (4).

Matrosov, E.S., Harou, J.J., Loucks, D.P., 2011. A computationally efficient open-source water resource system simulator – Application to London and the Thames Basin. Environ. Modell. Software 26 (12), 1599–1610. http://dx.doi.org/10.1016/j.envsoft.2011.07.013.

Matrosov, E.S., Padula, S., Harou, J.J., 2013a. Selecting portfolios of water supply and demand management strategies under uncertainty – contrasting economic optimisation and 'robust decision making' approaches. Water Resour. Manage 27 (4), 1123–1148.

Matrosov, E.S., Woods, A.M., Harou, J.J., 2013b. Robust Decision Making and Info-Gap Decision Theory for water resource system planning. J. Hydrol. 494, 43–58.

Mays, L.W., 2005. Water Resources Systems Management Tools. McGraw-Hill, New York.

McDonald, R.I., Green, P., Balk, D., Fekete, B.M., Revenga, C., Todd, M., Montgomery, M., 2011. Urban growth, climate change, and freshwater availability. Proc. Natl. Acad. Sci. 108 (15), 6312–6317.

Milly, P.C.D., Betancourt, J., Falkenmark, M., Hirsch, R.M., Kundzewicz, Z.W., Lettenmaier, D.P., Stouffer, R.J., 2008. Climate change – stationarity is dead: whither water management? Science 319 (5863), 573–574.

Nicklow, J., Reed, P., Savic, D., Dessalegne, T., Harrell, L., Chan-Hilton, A., Karamouz, M., Minsker, B., Ostfeld, A., Singh, A., Zechman, E., 2010. State of the art for genetic algorithms and beyond in water resources planning and management. J. Water Resour. Plan. Manage. – ASCE 136 (4), 412–432.

Padula, S., Harou, J.J., Papageorgiou, L.G., Ji, Y., Ahmad, M., Hepworth, N., 2013. Least economic cost regional water supply planning – optimizing infrastructure investments and demand management for South East England's 17.6 million people. Water Resour. Manage 27 (15), 5017–5044.

Rani, D., Moreira, M., 2009. Simulation–optimization modeling: a survey and potential application in reservoir systems operation. Water Resour. Manage.

Reed, P.M., Hadka, D., Herman, J.D., Kasprzyk, J.R., Kollat, J.B., 2013. Evolutionary multiobjective optimization in water resources: the past, present, and future. Adv. Water Resour. 51, 438–456.

Revelle, C., 1999. Optimizing Reservoir Resources: Including a New Model for Reservoir Reliability. John Wiley and Sons, New York, New York.

Sanderson, M.G., Wiltshire, A.J., Betts, R.A., 2012. Projected changes in water availability in the United Kingdom. Water Resour. Res. 48 (8), W08512.

Tang, Y., Reed, P., Wagener, T., 2006. How effective and efficient are multiobjective evolutionary algorithms at hydrologic model calibration? Hydrol. Earth Syst. Sci. 10 (2), 289–307.

Tanyimboh, T.T., Kalungi, P., 2009. Multicriteria assessment of optimal design, rehabilitation and upgrading schemes for water distribution networks. Civil Eng. Environ. Syst. 26 (2), 117–140.

Thames Water, 2010. Revised Water Resources Management Plan. Thames Water, UK.

Tsoukias, A., 2008. From decision theory to decision aiding methodology. Eur. J. Oper. Res. 187 (1), 138–161.

Veolia Water Central Limited, 2010. Water Resources Management Plan, Hatfield, UK.

Vorosmarty, C.J., Green, P., Salisbury, J., Lammers, R.B., 2000. Global water resources: vulnerability from climate change and population growth. Science 289 (5477), 284–288.

Walsh, C.L., Blenkinsop, S., Fowler, H.J., Burton, A., Dawson, R.J., Glenis, V., Manning, L.J., Kilsby, C.G., 2015. Adaptation of water resource systems to an uncertain future. Hydrol. Earth Syst. Sci. Discuss. 12 (9), 8853–8889.

Watkins, D.W., McKinney, D.C., 1997. Finding robust solutions to water resources problems. J. Water Resour. Plan. Manage. – ASCE 123 (1), 49–58.

Woodruff, M.J., Reed, P.M., Simpson, T.W., 2013. Many objective visual analytics: rethinking the design of complex engineered systems. Struct. Multidiscip. Optim. 48 (1), 201–219.

Wurbs, R.A., 1993. Reservoir-system simulation and optimization models. J. Water Resour. Plan. Manage. – ASCE 119 (4), 455–472.

WWF-UK, 2008. Thames Basin Vulnerability Report. WWF-UK, Surrey, England.

Yoon, J., Shoemaker, C., 2001. Improved real-coded GA for groundwater bioremediation. J. Comput. Civil Eng. 15 (3), 224–231.

Zeleny, M., 2005. The evolution of optimality: De novo programming. In: Coello Coello, C.A., Hernández Aguirre, A., Zitzler, E. (Eds.), Evolutionary Multi-Criterion Optimization.