Applied Energy 195 (2017) 114–124

Contents lists available at ScienceDirect

Applied Energy

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

Exploring multi-objective trade-offs in the design space of a waste heat recovery system

Maizura Mokhtar [a,*], Stephen Burns [b], Dave Ross [b], Ian Hunt [a]

[a] School of Engineering & the Built Environment, Edinburgh Napier University, Merchiston Campus, Edinburgh EH10 5DT, UK
[b] SRUC Research Division, Scotland's Rural College (SRUC), Roslin Institute Building, Easter Bush, Midlothian EH25 9RG, UK

HIGHLIGHTS

- Waste heat recovery system optimisation is a multi-objective optimisation problem.
- An MOEA is used to optimise a waste heat recovery system.
- Clustering discovers representative trade-offs amongst Pareto-optimal solutions.
- Combining clustering with parallel coordinates eases the analysis of trade-offs.

ARTICLE INFO

Article history:
Received 30 December 2016
Received in revised form 20 February 2017
Accepted 8 March 2017

Keywords:
Waste heat recovery
Optimisation
Multi-objective evolutionary algorithm
Mutually exclusive objective functions

ABSTRACT

A waste heat recovery system (WHRS) is used to capture waste heat released from an industrial process, and transform the heat into reusable energy. In practice, it can be difficult to identify the optimal form of a WHRS for a particular installation, since this can depend on various design objectives, which are often mutually exclusive. More so when the number of objectives is large. To address this problem, a multiobjective evolutionary algorithm (MOEA) was used to explore and characterise the trade-off surface within the design space of a particular WHRS. A combination of clustering algorithm and parallel coordinates plots was proposed for use in analysing the results. The trade-off surface is first segmented using a clustering algorithm and parallel coordinates plots are then used to both visualise and understand the resulting set of Pareto-optimal designs. As a case study, a simulation of a WHRS commonly found in the food and drinks process industries was developed, comprising of a desuperheater coupled to a hot water reservoir. The system was parameterised, considering typical objectives, and the MOEA used to build a library of alternative Pareto-optimal designs that can be used by installers. The resulting visualisation are used to better understand the sensitivity of the system's parameters and their trade-offs, providing another source of information for prospective installations.

1. Introduction

Many energy systems have behaviours that are sensitive to their parameter values, and these values need to be optimised [1,2]. Many of these systems also have multiple objectives, and these are often conflicting, meaning that in practice there is not a single design that satisfies every objective. In this situation, choosing an appropriate set of parameters first involves understanding the different trade-offs that can be made within the system's design space. One way of doing this is to identify all the solutions that are no better or worse than each other when considered across all objectives. This is known as the Pareto optimal set. Multi-objective optimisation algorithms (MOOA) are a group of optimisation techniques that are able to find good approximations of the Pareto optimal set. Because of this, they have become increasingly popular in the design of engineering systems [3–6], including the optimisation of energy systems. Examples are described in [7–17].

A Waste Heat Recovery System (WHRS) captures waste heat released from an industrial process, storing it in a form that can later be reused, for example using a hot water reservoir (HWR). A WHRS is a good example of a system with mutually conflicting objectives. Common practise in WHRS optimisation involves optimising using conventional single objective methods, typically Mixed Integer (Non-)Linear Programming (MILP or MINLP) [18–23]. In cases where there are multiple objectives, these are scalarised into a single objective. Alternatively, optimisation can be done using only a single objective, most typically minimising the cost of the installation and operations summed together, and

* Corresponding author.
E-mail address: M.Mokhtar@napier.ac.uk (M. Mokhtar).

http://dx.doi.org/10.1016/j.apenergy.2017.03.030
0306-2619/

Nomenclature

n	number of objectives	T_{ri}	input refrigerant temperature to the DSH
u	number of evolved parameters	T_{ro}	output refrigerant temperature from the DSH
k	number of clusters	T_{wi}	input water temperature to the DSH
$\dot{m}_{wd}$	mass flow rate of water into the DSH/HWR	T_{wo}	output water temperature from the DSH
$\dot{m}_{wd_{max}}$	maximum mass flow rate of water into the DSH/HWR	T_{wt}	water temperature in the HWR
m_{wt}	mass of water in the HWR	ΔT_{max}	maximum difference between T_{mx} and T_{hw}
$m_{wt_{min}}$	minimum mass of water in the HWR	DSH	desuperheater
$m_{wt_{max}}$	maximum mass of water in the HWR	EA	evolutionary algorithm
P_b	power of the backup heater	HWR	hot water reservoir
$P_{b_{max}}$	maximum power of the backup heater	MOEA	multiobjective evolutionary algorithm
T_{hw}	required/demanded hot water temperature	MOOA	multiobjective optimisation algorithm
T_{mx}	maximum water temperature in the HWR	WHRS	waste heat recovery system
T_m	mains water temperature		

the remaining objectives are implemented as constraints [18,19,21–23]. Given that single objective optimisation only provides a single solution, any information on objective trade-offs will be lost during the optimisation process, and any analysis of how the parameters affect the system efficiency will also be lost. In this work, by comparison, a Multi-Objective Evolutionary Algorithm (MOEA) was used to carry out multi-objective optimisation of a WHRS, using the results to visualise trade-offs in the design space.

As a case study, a simulation of a type of WHRS commonly found in the food and drinks process industry was developed, which involves a desuperheater connected to a HWR, providing hot water for a production plant's intermittent internal cleaning process. After identifying a number of common objectives for this type of system, the MOEA was used to find a set of trade-off designs within this objective space.

One disadvantage of MOOA and MOEA approaches is the visualisation of results, especially when n number of objectives and u number of parameters are large. A common method of visualising the results is to use a scatter plot, or decision maps, as shown in [14,10,13], but in general these can only show three objectives or parameters at once, which is limiting. Furthermore, the correlation between parameter and objective values are not explicitly shown, and in practice, analysis of the correlation is only applied to a selected few solutions, as depicted in [7–9]. Typically, they are the solutions at the extrema of the Pareto-optimal set, giving limited insight to the attributes of the other Pareto-optimal solutions found. Consequently many studies limit themselves to $n \leqslant 3$ objectives, as shown in [9,13,11,17,8].

To aid in the analysis of the results from high-dimensional multi-objective optimisation, and to identify a reduced set of representative designs, this paper proposed an alternative method of visualising the Pareto-optimal solutions. The solutions are first clustered into k-number of clusters, either in the design space or the parameter space, to identify the degree of commonality between the solutions. For each cluster identified, parallel coordinates [24,25] are used to visualise the high-dimensional solution space and objective space as a pair of two-dimensional plots; one for each of the spaces.

Parallel coordinate plots are used to visualise the Pareto-optimal solutions, each for the u-dimensional solution space and n-dimensional objective space. The correlation between a solution and its objective values in a specific cluster are identified by the common colour used in both plots. This method of visualisation can therefore reduce the number of figures (and tables) to depict the results significantly, down to $2k$ figures – one for each of the two spaces. The significant reduction in the number of figures used eases in the analysis of the trade-offs between the Pareto-optimal solutions.

The paper is organised as follows: Section 2 provides a brief introduction to the MOOA used in this work. Section 3 introduces WHRS and gives an overview of the case study. Section 4 presents results and analysis using various multidimensional visualisation methods, including decision maps and parallel coordinates plots. Section 5 concludes the paper.

2. Multi-Objective Evolutionary Algorithm (MOEA)

Whilst many forms of optimisation can be generalised to the multi-objective case, in practice the most widely used forms of MOOA are based around evolutionary algorithms (EA). EAs are a class of population-based metaheuristic optimisation algorithms. As described in [26]: *the initial population is a random sample of search points, a selection mechanism then discards search points with poor objective values, and variation operators derive a new population of search points from those that remain. This new population then replaces the previous population, and the process of selection and variation are repeated until an optimal solution is found, or some other termination criterion is met.* The search points, in our case, are vectors of parameter values. The variation operators are crossover, which recombines two search points by swapping vector elements, and mutation, which randomly replaces one or more vector elements to create a new search point. Population-based metaheuristics, such as EAs, carry out a relatively broad search of an optimisation space, and consequently are often able to find better solutions than local search metaheuristics, such as hill climbing or simulated annealing.

A multi-objective evolutionary algorithm (MOEA) is a specialised form of EA, and in this work, a popular MOEA called NSGA-II [27] was used. The main difference between NSGA-II and a single-objective EA lies in how selection takes place. Rather than only propagating the best search points from one generation to the next, NSGA-II first carries out a ranking of the search points in the population. Search points which are no worse than any other when considered across all the objectives are known as dominating search points, and are given a rank of 1. Those which are only dominated by rank 1 search points are assigned rank 2, etc. After ranking, the first half of the ordered population is then copied directly to the next generation, and the remainder of the population is filled by applying the variation operators.

NSGA-II also uses a diversity preservation method, known as crowding, to encourage search in regions of objective space which have not been previously explored. For the search points that are of the same rank, those that are more dissimilar to the others are preferred for selection for the next generation population. This not only results in an approximation of the Pareto optimal set which