Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Efficiently solving multiple objective optimal control problems

Peter M. Van Erdeghem * **Filip Logist** * **Ilse Y. Smets** *
Jan F. Van Impe *

* *BioTeC, Department of Chemical Engineering, Katholieke Universiteit Leuven, W. de Croylaan 46, B-3001 Leuven, Belgium. (Tel: +32-16-32.26.75, e-mail: jan.vanimpe@cit.kuleuven.be)*

optimal control; multiobjective optimisations; multiple-criterion optimisation; differential equations; process control; chemical industry

Abstract: This paper discusses an improved method for solving multiple objective optimal control (MOOC) problems, and efficiently obtaining the set of Pareto optimal solutions. A general MOOC procedure has been introduced in Logist et al. [2007], to derive optimal generic temperature profiles for a steady-state tubular plug flow reactor. This procedure is based on a weighted sum of the different costs, and combines analytical and numerical optimal control techniques. By varying the weights, the exact Pareto set has been obtained. However, it is known for the weighted sum approach that a uniform variation of the weights, not necessarily leads to an even spread on the Pareto front (and thus an accurate representation). In addition, the analytical derivations in the proposed procedure become intractable for large-scale systems. Therefore, this paper introduces two modifications: the use of (*i*) the normal constraint method (Messac and Mattson [2004]) instead of the weighted sum, and (*ii*) piecewise linear approximations instead of the analytical relations. Two examples, i.e., (*i*) a classic minimum time, minimum control effort problem, and (*ii*) a more real-life determination of optimal temperature profiles for tubular reactors, illustrate the enhanced performance and the general applicability of the procedure.

1. INTRODUCTION

In practical optimisation problems often several and conflicting objectives are simultaneously present, e.g., maximising the strength of a construction, while minimising its weight. These *multiple objective optimisation* (MOO) problems produce most often a set of optimal solutions (or the *Pareto set*) instead of one sole. While research on scalar MOO problems has attracted much attention over the years (see, e.g., Marler and Arora [2004] for an overview), much less effort is spent on *multiple objective optimal control* (MOOC) problems (i.e., when an infinite dimensional optimal control profile has to be found).

Recently, Logist et al. [2007] have proposed a procedure which allows to derive analytical optimal profiles for an exothermic tubular reactor under steady-state with conflicting conversion and energy objectives. The rationale behind this procedure is to combine (*i*) a *weighted sum* approach (which converts the multiple objective problem into a single objective problem) with (*ii*) an *analytical control parameterisation*. By varying the weights, a representation of the Pareto frontier is obtained. The procedure involves four steps. First, the set of all possible optimal arcs is derived analytically. Second, approximate piecewise constant optimal controls are computed numerically for a coarse grid of weights. From these approximate solutions the optimal arc sequences are each time identified. Based on the analytical control expressions and the optimal sequences an analytical control parameterisation is built each time. Finally, the switching positions between the different intervals are optimised over a refined weight grid.

However, several points for further improvement have been noticed. First, a uniform distribution of the weights, does not necessarily yield an equal distribution on the Pareto front, and second, the analytical derivations involved, become intractable for large-scale systems. Therefore, (*i*) the weighted sum is replaced as approach to tackle the multiple objective aspect, and (*ii*) low-order polynomials are employed to approximate the analytical relations.

These items are structured in the paper as follows. Section 2 mathematically formulates an MOOC problem, and introduces typical aspects and methods for MOO. Section 3 details the MOOC procedure by Logist et al. [2007], and proposes modifications. Section 4 introduces the two cases: (*i*) a minimum time minimum control effort problem, and (*ii*) the determination of an optimal temperature profile in a tubular reactor, while Section 5 presents the results. Section 6 finally summarises the main conclusions.

2. MULTIPLE-OBJECTIVE OPTIMAL CONTROL

2.1 Mathematical formulation

In general a multiple objective optimal control problem is formulated as follows:

$$\min_{u(\xi),\ \xi\in[0,\xi_f],\ \xi_f} \mathbf{J} = (J_1, \ldots, J_m)^{\mathrm{T}} \tag{1}$$

with:

$$J_i = \underbrace{h_i[\mathbf{x}(\xi_f)]}_{\text{Terminal cost}} + \underbrace{\int_0^{\xi_f} g_i[\mathbf{x}(\xi), u(\xi)]d\xi}_{\text{Integral cost}}$$

subject to the system equations:

$$\frac{d\mathbf{x}}{d\xi} = \mathbf{f}[\mathbf{x}(\xi), u(\xi), \xi] \tag{2}$$

the boundary conditions:

$$\mathbf{0} = \mathbf{I}[\mathbf{x}(0)] \tag{3}$$

$$\mathbf{0} = \mathbf{S}[\mathbf{x}(\xi_f)] \tag{4}$$

and the path constraints on the states and control:

$$\mathbf{0} \geq \mathbf{C}_1[\mathbf{x}(\xi)] \tag{5}$$

$$\mathbf{0} \geq \mathbf{C}_2[u(\xi)] \tag{6}$$

Let $\mathcal{U}$ be the feasible region, i.e., the set of all admissible controls $u(\xi)$, which induce admissible state trajectories $\mathbf{x}(\xi)$. A feasible control is said to be *inside* the feasible region, when no control or path constraint is active.

2.2 Multiple objective optimisation: concepts and methods

Contrary to single objective optimisation, typically no single global solution exists in multiple objective optimisation. Therefore, it is necessary to determine a set of points that all fit a predetermined optimality definition, which is most often the concept of *Pareto optimality.*

Definition: *A control $u^* \in \mathcal{U}$, is Pareto optimal iff there does not exist another control, $u \in \mathcal{U}$, such that $\mathbf{J}(u) \leq \mathbf{J}(u^*)$ and $J_i(u) < J_i(u^*)$ for at least one cost.*

In other words, a control is said to be Pareto optimal if there exists no other control that improves at least one objective function without worsening another.

Solution methods for MOO problems are broadly classified in two categories. A first class transforms the MOO problem into a single objective optimisation problem. Then, by varying parameters of the methods involved, often a representation of the Pareto set is obtained. This class includes the classic *weighted convex sum* of the different objectives, but also encompasses novel methods as *normal boundary intersection* (Das and Dennis [1998]), or *(normalised) normal constraint* (Messac and Mattson [2004]). These last two use a geometrically intuitive approach to obtain a uniform spread on the Pareto front when the parameters are uniformly varied. Hereby, the former introduces additional equality constraints, whereas the latter only requires additional inequality constraints. The second class of solution methods, amongst which, e.g., *genetic algorithms*, have been found to be able to generate the Pareto set directly from the multiple objective formulation (Konak et al. [2006]). For more details on MOO methods the reader is referred to a survey by Marler and Arora [2004].

3. SOLUTION STRATEGY

This section presents the original procedure by Logist et al. [2007] to compute the Pareto set, and proposes two modifications.

3.1 Original solution strategy

The original procedure is based on a weighted sum approach and involves four successive steps.

In **step 1**, analytical expressions for all possible optimal control arcs are calculated using analytical techniques from, e.g., Kirk [1970] and Srinivasan et al. [2003]. The derivation of possible optimal controls inside the feasible region is performed by eliminating the costates from the necessary conditions for optimality. This elimination can be done *explicitly*, for a control affine Hamiltonian by using Pontryagin's minimum principle (Pontryagin et al. [1962], Kirk [1970]), or *implicitly*, by a determinant procedure (Srinivasan et al. [2003]) for non-affine Hamiltonians. Control arcs which keep a constraint active, as well as the corresponding tangency conditions, are found by differentiating the active constraint with respect to the independent variable, and substituting the system equations until the control appears explicitly.

In **step 2**, approximate optimal control profiles are determined numerically for a coarse grid of weights. Hereto, a numerical optimal control approach with a piecewise constant control parameterisation is adopted. Here, the software package MUSCOD-II (Leineweber et al. [2003]) is used, which implements a direct multiple shooting approach (Bock and Plitt [1984]).

In **step 3**, the optimal arc sequences present in the numerically obtained optimal control profiles are identified by visual inspection. This human intervention is, however, not a restriction as an automated identification procedure (Schlegel and Marquardt [2006]) can also be used.

Finally, in **step 4**, the control parameterisation is refined based on the identified optimal sequence using the analytical expressions. These analytical parameterisations with the switching positions between the different arcs as degrees of freedom, are then optimised on a refined grid, assuming the optimal arc sequence does not change. When neighbouring analytical parameterisations are different (i.e., a different optimal sequence of arcs), both sequence types have to be optimised over the refined grid, and the one yielding the lowest cost value has to be selected. Also here MUSCOD-II is employed as software environment.

3.2 Modifications

As the weighted sum exhibits some intrinsic drawbacks, it is replaced by the normal constraint method. This method is preferred over normal boundary intersection as it adds inequality instead of equality constraints. In the normal constraint method, also a series of optimisations has to be performed for a grid of parameters in order to obtain the Pareto set. Note that this modification only alters the computational steps, i.e., step 2 and 4.

As the analytical calculations become tedious for large-scale systems, this step can be omitted. The different arc types are still identified based on a piecewise constant control approximation, but a limited number of low-dimensional polynomials is now employed instead of the analytical relations. Although this representation induces approximation errors since the control is not exact any more, the deviations in cost with respect to the exact values are often small.

4. CASE STUDIES

4.1 Minimum time - minimum control effort problem

To illustrate the general applicability of the procedure, the problem of transferring a car from an initial position to a specified target in minimum time, and with a minimum control effort is considered. For instance, the following dynamic model may approximate the driving of the car,

$$\frac{dx_1}{dt} = x_2 \tag{7}$$
$$\frac{dx_2}{dt} = u \tag{8}$$

where the position x_1 [m], and the velocity x_2 [m/s] are the states, and the control u is the acceleration [m/s^2], which is in practice controlled by pushing the accelerator, or hitting the brakes. Note that in this case the time t [s] is the independent variable in the ODEs.

The aim is to drive 400 m, starting and ending at rest:

$$x_1(t = 0) = 0 \text{ m} \quad \text{and} \quad x_1(t = t_f) = 400 \text{ m} \tag{9}$$
$$x_2(t = 0) = 0 \text{ m/s} \quad \text{and} \quad x_2(t = t_f) = 0 \text{ m/s} \tag{10}$$

while minimising on the one hand, the control effort for accelerating (which can by interpreted as the fuel consumption):

$$J_1 = \int_0^{t_f} \max(0, u)dt, \tag{11}$$

and on the other hand the (scaled) travelling time:

$$J_2 = \frac{t_f}{K} = \int_0^{t_f} \frac{1}{K}dt \tag{12}$$

with $K = 20$. These are obviously conflicting objectives since a small travelling time requires a high speed, and, hence, also a large consumption of fuel for reaching this velocity. Since infinitely fast accelerating and decelerating is impossible, the control is bounded between:

$$-5 \text{ m/s}^2 \leq u(t) \leq 5 \text{ m/s}^2. \tag{13}$$

Additional constraints, are a speed, and a time limit:

$$x_2(t) \leq 40 \text{ m/s} \tag{14}$$
$$t_f \leq 100 \text{ s}. \tag{15}$$

Table 1. Process parameters.

C_{in}	$= 0.02$ mole $\cdot$ L^{-1}	T_{min}	$= 280$ K
E	$= 47092.5$ J $\cdot$ mole^{-1}	T_{max}	$= 400$ K
k_0	$= 10^6$ s^{-1}	$T_{w,min}$	$= 280$ K
K	$= 250000$	$T_{w,min}$	$= 400$ K
L	$= 1$ m	v	$= 0.1$ m $\cdot$ s^{-1}
R	$= 8.3145$ J $\cdot$ mole^{-1} $\cdot$ K^{-1}	β	$= 0.2$ s^{-1}
T_{in}	$= 340$ K	δ	$= 0.25$

4.2 Tubular reactor problem

The investigated reactor is a jacketed tubular plug flow reactor with a fixed length L [m], which operates under steady-state conditions. Inside the reactor an exothermic irreversible first-order reaction takes place:

$$\frac{dx_1}{dz} = \frac{\alpha}{v}(1 - x_1)e^{\frac{\gamma x_2}{1+x_2}} \tag{16}$$
$$\frac{dx_2}{dz} = \frac{\alpha\delta}{v}(1 - x_1)e^{\frac{\gamma x_2}{1+x_2}} + \frac{\beta}{v}(u - x_2) \tag{17}$$

with $x_1 = (C_{in} - C)/C_{in}$, the dimensionless concentration C [mole $\cdot$ L^{-1}] [-], $x_2 = (T - T_{in})/T_{in}$ [-], the dimensionless reactor temperature T [K], and $u = (T_w - T_{in})/T_{in}$ [-], the dimensionless jacket temperature T_w [K]. The constants α, β, γ, and δ are defined as follows:

$$\alpha = k_0 e^{\frac{-E}{RT_{in}}},\ \beta = \frac{4h}{\rho C_p d},\ \gamma = \frac{E}{RT_{in}}, \delta = -\frac{\Delta H}{\rho C_p}\frac{C_{in}}{T_{in}},$$

The exact parameter values are provided in Table 1. As the jacket temperature T_w is in practice often used to control the reactor, this variable is selected as the control $u(z)$. Note that the spatial coordinate z [m] is now the independent variable. The aim is to derive a jacket fluid temperature profile that minimises the reactant concentration at the outlet (i.e., maximises conversion):

$$J_1 = C_{in}(1 - x_1(L)) \tag{18}$$

while minimising the terminal heat loss:

$$J_2 = \frac{T_{in}^2 x_2^2(L)}{K}. \tag{19}$$

These are conflicting goals since high temperatures favour conversion. To avoid hazardous situations explicit bounds are added on the reactor and the jacket temperature.

$$x_{2,min} \leq x_2(z) \leq x_{2,max} \tag{20}$$
$$u_{min} \leq u(z) \leq u_{max} \tag{21}$$

5. RESULTS

First, the analytical results are discussed. Then, the effect of using the normal constraint method instead of the weighted sum is illustrated. Afterwards, the influence of employing low-order polynomials instead of the analytical relations is presented.

5.1 Analytical results

As mentioned, the first step consists of analytically deriving the set of optimal arcs. For brevity the symbolic calculations have been omitted, and only the results are provided (see Table 2 for an overview).

Table 2. Analytical relations for the different case studies.

Control		Car problem	Tubular reactor
At control bounds	$u_{min}(z)$	-5 m/s^2	$(T_{w,min} - T_{in})/T_{in}$
	$u_{max}(z)$	5 m/s^2	$(T_{w,max} - T_{in})/T_{in}$
At state bounds	$u_{path}(z)$	0 m/s^2	$x_2 - \frac{\alpha\delta}{\beta}(1-x_1)e^{\frac{\gamma x_2}{1+x_2}}$
		@ $x_{2,max} = 40$ m/s	@ $x_{2,min} = (T_{min} - T_{in})/T_{in}$ or @ $x_{2,max} = (T_{max} - T_{in})/T_{in}$
Inside feasible region	$u_{inside}(z)$	0 m/s^2	impossible

For the car problem, the set contains the following four elements: u_{min}, u_{max}, $u_{path} = 0$ m/s^2, and $u_{inside} = 0$ m/s^2. Here, the first two are the upper and lower limit on the acceleration. The third one, i.e., nor breaking, nor accelerating, is required to maintain a constant velocity in order to stay on the upper speed limit. By chance this is also the relation for a control inside the feasible region.

The possible arcs for the tubular reactor have been derived in Logist et al. [2008], and involve: (i) the upper and lower limit of the jacket fluid temperature (u_{min} and u_{max}), (ii) and a nonlinear relation u_{path} to keep the reactor at its lower or upper temperature limit ($x_{2,min}$ or $x_{2,max}$). For this case, an arc inside the feasible region is impossible.

5.2 Normal constraint vs. weighted sum

For both (bi-objective) problems, the convex weighted sum results in a single scalar cost $J = (1 - A)J_1 + AJ_2$ which involves a trade-off parameter A. By varying A from zero to one, the Pareto set can be obtained. The normal constraint method, however, introduces parameters to move in the *utopia hyperplane* (i.e., the plane that connects all individual minima in the cost space). In the current situation only one parameter $W \in [0, 1]$ is required to move along the utopia line. Thus, varying W between zero and one leads here to a representation of the Pareto front.

In both approaches a uniform grid of 11 points is employed for the parameter A or W in order to construct an approximation of the Pareto set.

The resulting Pareto frontiers for the car problem are depicted in Fig. 1. Clearly, the normal constraint method yields a more uniform spread, and, hence, a more accurate representation. The corresponding optimal control profiles are displayed in Fig. 2, while the state trajectories based on the normal constraint method can be found in Fig. 3. Note that for visibility reasons the time axis has been scaled by the final time t_f. As can be seen, both methods yield an optimal control sequence of u_{max}-u_{path}-u_{min} for high A- or W-values. This situation corresponds to accelerating as fast as possible until the speed limit $x_{2,max}$ is reached. Then, this speed has to be maintained for a while, whereas towards the end, the brake is hit as hard as possible in order to arrive at rest at the destination. For higher values, however, an arc inside the feasible region u_{inside} replaces the path constrained arc u_{path}. In these cases, an intermediate velocity $x_{2,intermediate}$ is maintained after an initial acceleration, in order to limit the fuel consumption. It is also clear from Fig. 2 that the control profiles obtained from the normal constraint method yield a more gradual evolution between the two extremes.

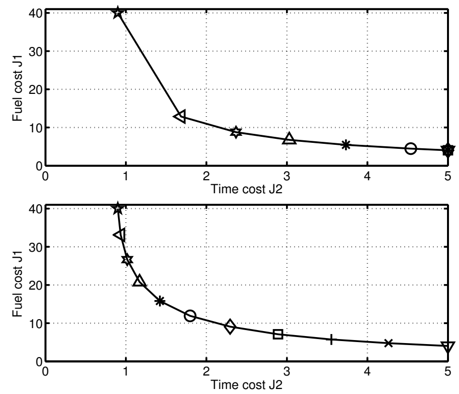

Fig. 1. Car problem: Pareto set generated by weighted sum (top) and normal constraint (bottom).

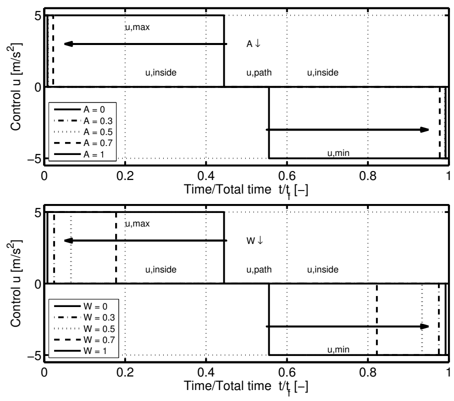

Fig. 2. Car problem: control profiles obtained with: weighted sum (top) and normal constraint (bottom).

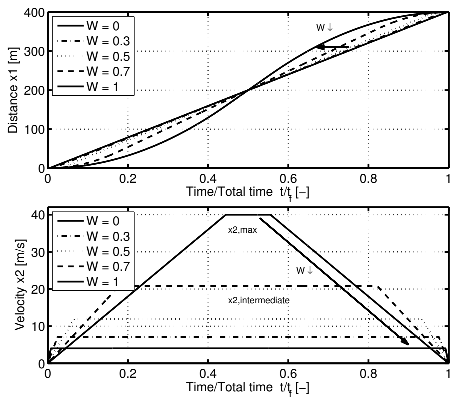

Fig. 3. Car problem: state profiles obtained with normal constraint: distance (top) and velocity (bottom).

Similar conclusions apply to the plug flow reactor case (see Fig. 4). A uniform grid of weights A results in a highly skew distribution on the Pareto front, whereas the normal constraint yields an accurate approximation. Figs. 5 and 6 display the optimal control profiles for both methods, and the state profiles corresponding to the normal constraint results, respectively. Clearly, the optimal jacket fluid temperature profiles consist of a sequence u_{max}-u_{min}-u_{path}-u_{min}, where the last *min* interval shrinks (and eventually vanishes) for decreasing A- or W-values. The explanation is that heating the reactor as fast as possible, and maintaining the upper reactor temperature limit, stimulates conversion, whereas cooling the reactor towards the outlet reduces the terminal heat loss. Also here, the control profiles generated by the normal constraint method, exhibit a more gradual evolution between the two extreme cases.

In summary, the normal constraint method outperforms the classic weighted sum approach. Although, the presented cases only involve two objectives, the procedures are extendable to more objectives. However, this will lead to more parameters to be varied and, hence, larger computation times.

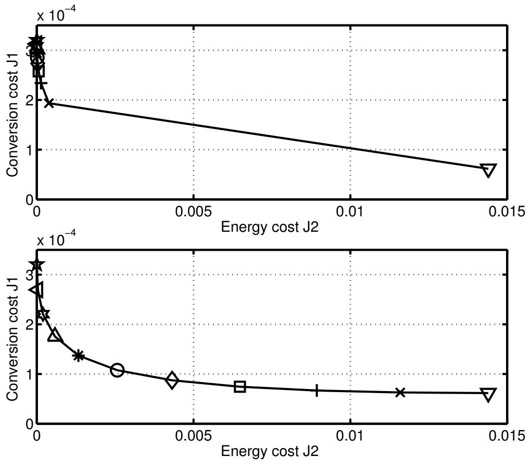

Fig. 4. Plug flow reactor: Pareto set generated by weighted sum (top) and normal constraint (bottom).

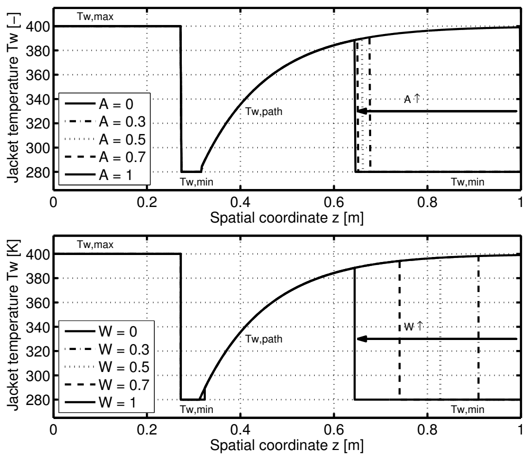

Fig. 5. Plug flow reactor: control profiles obtained with: weighted sum (top) and normal constraint (bottom).

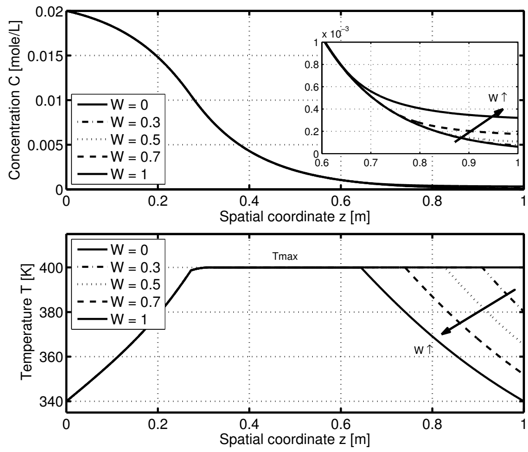

Fig. 6. Plug flow reactor: state profiles obtained with normal constraint: concentration (top) and reactor temperature (bottom).

5.3 Polynomial approximations vs. analytical relations

In this section the approximation error is studied when polynomials are employed in step 4 of the original procedure instead of the exact analytical relations. Since only in the plug flow reactor exhibits a nonlinear arc, i.e., u_{path}, this case is concentrated on.

Several approximations are proposed and their parameters are determined by minimising the total weighted sum for $A = 0.5$. PWL1 is a single linear function, while PWL2 consists of two continuous piecewise linear parts. The original analytical parameterisation (AP) has three degrees of freedom (i.e., the switching positions), whereas PWL1 and PWL2 add, respectively, one and two slopes as parameters to be optimised. For reasons of comparison, also PWC5U and PWC5ZU are selected. The former consists of a uniform piecewise constant approximation with five control interval of fixed length (i.e., five degrees of freedom). The latter allows also an optimisation of the interval lengths (i.e., nine degrees of freedom).

Fig. 7 displays the analytical solution and the resulting optimised profiles for the different approximations. Fig. 8 depicts the corresponding concentration and reactor temperature profiles. Here, it is observed that PWL2 approximates AP fairly well. PWL1 on the other, experiences more problems since it possesses one degree of freedom less, and apparently squeezes the little *min* control part. Although both PWC5U and PWL2 have an equal number of degrees of freedom, PWC5U does not succeed in providing an accurate approximation. Allowing free interval lengths (PWC5ZU), however, introduces more flexibility and results in a more accurate representation. The same trend is observed, when the total cost values are displayed for a range of trade-off values (Fig. 9). The AP yields the lowest cost, but the increase for PWL2 and PWC5ZU is limited. PWL1 and PWC5U on the other hand, cause significant cost increases. Thus, employing a few piecewise linear parts instead of analytical relations often only induces small increases in the cost. Consequently, the analytical derivations of step 1 can be omitted, which enables the applicability of the proposed procedure to larger systems.

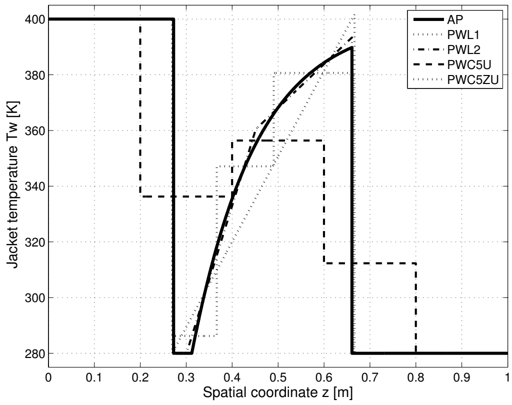

Fig. 7. Plug flow reactor: optimised control approximations.

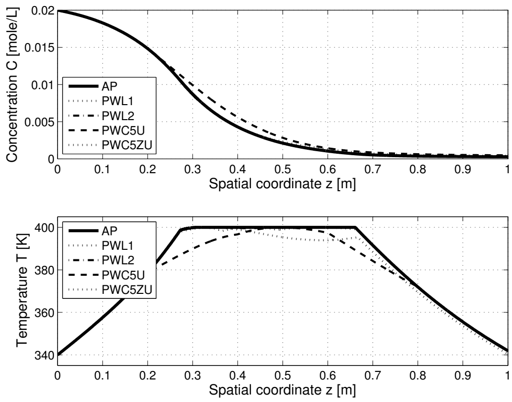

Fig. 8. Plug flow reactor: concentration (profiles) and temperature profiles (bottom).

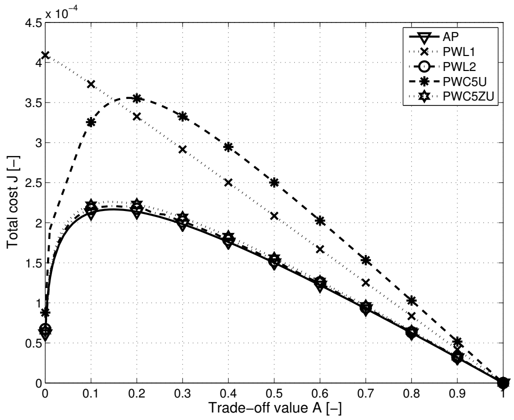

Fig. 9. Plug flow reactor: total cost for different optimised approximations.

6. CONCLUSIONS

This paper has proposed two improvements to the general strategy for solving multiple objective optimal control (MOOC) problems, introduced by Logist et al. [2007].

First, the weighted sum approach has been replaced by the normal constraint method, which enables an even spread along the Pareto set by uniformly varying the method's parameters. This modification allows a more efficient computation. Second, the symbolic derivations required to obtain analytical expressions for the set of optimal arcs may be omitted, since low-dimensional polynomials may provide an accurate approximation of the analytical control, and induce only a small increase in cost when optimised. Hence, large systems for which the symbolic calculations are intractable, can now be treated.

The two adaptations have been illustrated for a classic minimum time minimum control effort problem and a more real-life tubular reactor case study. Consequently, the strategy has been shown to be flexible, general, and not restricted to a particular MOOC problem.

ACKNOWLEDGEMENTS

Work supported in part by Projects OT/03/30 and EF/05/006 (Center-of-Excellence Optimization in Engineering) of the Research Council of the Katholieke Universiteit Leuven, and by the Belgian Program on Interuniversity Poles of Attraction, initiated by the Belgian Federal Science Policy Office. The scientific responsibility is assumed by its authors.

REFERENCES

H.G. Bock and K.J. Plitt. A multiple shooting algorithm for direct solution of optimal control problems. In *Procs of the 9th IFAC world congress*, 1984.

I. Das and J.E. Dennis. Normal-Boundary Intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM Journal on Optimization*, 8:631–657, 1998.

D.E. Kirk. *Optimal Control Theory: an Introduction*. Prentice Hall, Englewood Cliffs, New Jersey, 1970.

A. Konak, D.W. Coit, and A.E. Smith. Multi-objective optimization using genetic algorithms: a tutorial. *Reliab. Eng. Syst. Saf.*, 91:992–1007, 2006.

D.B. Leineweber, A. Schäfer, H.G. Bock, and J.P. Schlöder. An efficient multiple shooting based reduced SQP strategy for large-scale dynamic process optimization. Part II: Software aspects and applications. *Computers and Chemical Engineering*, 27:167–174, 2003.

F. Logist, P. Van Erdeghem, I.Y. Smets, and J.F. Van Impe. Multiple-objective optimisation of a jacketed tubular reactor. In *Procs of the European Control Conference*, 963–970, 2007.

F. Logist, I.Y. Smets, and J.F. Van Impe. Derivation of generic optimal reference temperature profiles for steady-state exothermic jacketed tubular reactors. *Journal of Process Control*, 18:92–104, 2008.

R.T. Marler and J.S. Arora. Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26:369–395, 2004.

A. Messac and C.A. Mattson. Normal constraint method with guarantee of even representation of complete Pareto frontier. *AIAA Journal*, 42:2101–2111, 2004.

L. Pontryagin, V. Boltyanskiy, R. Gamkrelidze, and Y. Mishchenko. *The mathematical theory of optimal processes*. Wiley - Interscience, New York, 1962.

M. Schlegel and W. Marquardt. Detection and exploitation of the control switching structure in the solution of dynamic optimization problems. *Journal of Process Control*, 16:275–290, 2006.

B. Srinivasan, S. Palanki, and D. Bonvin. Dynamic optimization of batch processes I. Characterization of the nominal solution. *Computers and Chemical Engineering*, 27:1–26, 2003.