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

On Evaluating Optimality Losses of Greenhouse Climate Controllers

G. van Straten and L.G. van Willigenburg

*Wageningen University, Systems and Control, Wageningen, The Netherlands
corresponding author's e-mail: gerrit.vanstraten@wur.nl*

Abstract: Optimal operation strategies for greenhouse crop cultivation can be computed with open loop dynamic optimization. These solutions are obtained off-line, and are valid under nominal weather conditions only. On-line, feed-back control is needed to cope with deviations from the nominal weather. One of the issues in practical control is how to link the off-line nominal solution to on-line control. One option is to use a receding horizon controller with the same goal function as used off-line, but enhanced with a term based on the co-state of the slow crop states to encapsulate the long term goals. Loss measures are introduced to evaluate this solution against various other approximate solutions proposed in the literature. To our knowledge this is the first time that various sub-optimal solutions are clearly listed and analysed. A simplified, but transparent, example is used to illustrate the various losses. One of the results is that receding horizon optimal control with an adapted goal function is superior to other more common control solutions.

Keywords: greenhouses, optimal control, protected cultivation.

1. INTRODUCTION

Control of greenhouse crop cultivation by balancing costs of resources against ultimate selling value of the crop has been the subject of several investigations. In the literature, not always a clear distinction is made between off-line dynamic optimization, which is valid for nominal (assumed) weather conditions, and on-line realisation in a control environment. Moreover, several approximate solutions have been proposed, for instance to treat the greenhouse as pseudo-static or ideally controlled on the time scale of the crop (Seginer, 1989; Ioslovich and Seginer, 1998) or to perform a two-time scale decomposition to cope with the system's stiffness (Van Henten, 1994). On the control level, hierarchical solutions are often used in practice, by trying to find optimized set-points. Focus on controller design has also been the subject of many papers, for instance Young et al. (1994), Pasgianos et al. (2003), Piñón et al. (2005), Blasco et al. (2007).

While off-line optimal control patterns can be found, using nominal weather, it is not straight-forward to see how these should be used in an on-line situation, where feed-back is needed to cope with weather variations and modelling errors, and the horizon is shorter due to the inability to obtain reliable weather forecast on a seasonal scale. Also, it is not clear what the loss is of certain simplifications, that have been chosen for believed easiness of implementation. These issues form the motivation for the work reported here.

Because of the complicated greenhouse physics and the complex dynamics of the crop, it is very difficult to fully comprehend the behaviour of optimal control solutions. In order to study this without the burden of model complications, we introduce an extremely simple but still illustrative system model, to clearly elucidate the issues raised above.

In order to fully appreciate the development below, it is important to note that the greenhouse-crop system is characterised by two major time-scales: the greenhouse responds to changes in environment on the scale of half an hour or so, whereas the crop biomass has a typical time constant in the order of weeks, or more.

2. OBJECTIVES

Our aims are:

1. To calculate the off-line optimal control trajectory for nominal weather, and compare this with an approximate solution based on time scale decomposition; such solutions are numerically attractive, and, more importantly, form the bridge between off-line calculations with nominal weather, and on-line control, as explained later.
2. To show the differences between the open loop optimal solution obtained for nominal weather and the 'dream' pattern obtained afterwards for real weather. The dream pattern serves as the reference to which other on-line solutions listed under 3-6 below should be judged. It

should be said right away that the result for real weather can be better or worse, because it may be more or less favourable compared to the nominal weather.
3. To evaluate the losses when the off-line computed fast state trajectories (available for the nominal weather only) are used as set-point for an on-line (classical) controller.
4. To point out that there may exist set-points for a classical controller that do better than under 3 above. However, the pre-requisite for finding such set-points is that the optimal solution is already available.
5. To show how two-time scale decomposition can be used to link long term seasonal goals to short term on-line control, thus realising a receding horizon controller which preserves the goal function attributes, and hence is expected to achieve the best possible on-line result.
6. To evaluate the losses of using current standard control with fixed day-night set-point pattern.

The organization of this paper is as follows. First, the simple example is presented. Next, the points above are addressed in sequence. Finally, the results will be summarized by comparing the various goal function values and losses. We conclude with a discussion, conclusions and recommendations for further research.

3. THE ILLUSTRATIVE EXAMPLE

The model has two state variables, the biomass x_1, and the air temperature x_2. The rate of change is given by

$$\dot{x}_1 = p_1 d_1 x_2 , \qquad (1)$$
$$\dot{x}_2 = p_2 \left(d_2 - x_2\right) + p_3 u_1 \qquad (2)$$

where d_1, d_2 are the external inputs solar radiation and outdoor temperature, respectively, u_1 is the heat input acting as the single control variable, and $p = [p_1, p_2, p_3]$ is a vector of fixed parameters. The purpose is to balance the costs of heating against the final value of produced biomass. The goal function is defined by

$$J = \Phi\left(x(t_f)\right) + \int_{t_o}^{t_f} L(x,u,d,p)dt \qquad (3)$$

where $\Phi(x(t_f))$ are the terminal costs, and L are the running costs, and, hence, in the current example the costs (negative benefits) are given by

$$J = -p_5 x_1(t_f) + \int_{t_o}^{t_f} p_4 u_1 dt \qquad (4)$$

where the additional parameters p_4 and p_5 are the energy price and price of the sold crop at final time t_f, respectively. Note that *J* depends on weather *d*.

The nominal weather d^{nom} is given by

$$d_1^{nom} = \max\left(0, 800\sin(2\pi(t-7.8)/24)\right), \qquad (5)$$

$$d_2^{nom} = 15 + 10\sin(2\pi(t-7.8)/24). \qquad (6)$$

Realisations of the 'real' weather are given in the applications below.

In the model above it makes sense to increase the indoor greenhouse temperature during periods when there is light. In the current implementation cooling of the greenhouse is passive via the exchange term with outdoor conditions. Values for the model parameters are given in the Appendix. An eigenvalues analysis of the model reveals that the eigenvector that is dominated by the crop state has a long time constant – in fact, it is infinite because the biomass is just an integrator. The time constant of the greenhouse is $1/p_2$, being 1 h with the current parameterization.

As the goal function as well as the model are linear in the heating (control *u*), it is known from optimal control theory that the solution is bang-bang (Bryson 1999).

4. METHODS

Because the solution is bang-bang, the optimal control patterns in this case were obtained by directly optimising for the switching moments, using the NLP solver `constr` in Matlab 6.5. Due to the periodic nature of the nominal disturbances, and the integrative biomass equation, the solution repeats itself for consecutive days. Therefore, a period of two days suffices in this case as an emulation of a full season.

The receding horizon controller for on line control uses the same algorithm with a control interval of 0.2 h, and a prediction horizon of 1 h, allowing two switching instances. The PI controller has anti-windup, to limit integral wind-up during periods when the greenhouse cools down by natural exchange. The controller parameters are found by using the tuning tools in SIMULINK 2.0.

5. RESULTS

5.1 Open loop optimal control for nominal weather

First, a dynamic optimization is performed, assuming fully known, nominal weather. This is typically the situation encountered when calculating a seasonal solution in an off-line fashion. Figure 1 shows the results.

The heating is switched on a little after sunrise, as it apparently does not pay off to heat when the light is still low, and switched off quite some time before sunset, to profit from the inertia of the greenhouse. The optimal control pattern on the second day and beyond is slightly different from the first day, as the effect of the initial conditions gradually vanishes. The virtual costs (negative benefit) at final time t_f = 48h is -3.30 units, which is mainly due to the value of the crop.

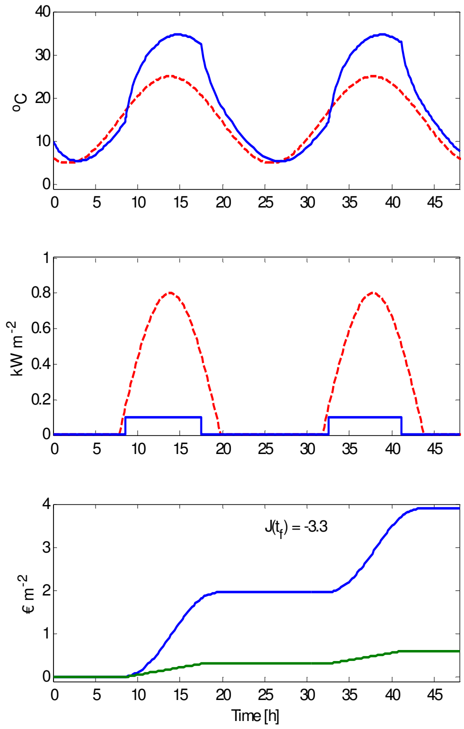

Fig. 1. Open loop optimal control with nominal weather. Top: optimal greenhouse temperature x_2 (solid) and outdoor temperature d_2 (dashed), middle: optimal heating pattern u (solid) and solar radiation d_1 (dashed), bottom: crop benefit p_5x_2 (upper line) and running heating costs (lower line).

5.2 Open loop optimal control for 'real' weather

In a real situation, the weather is, of course, not known over the full control horizon. This is one major reason why some form of feed-back is needed (the other reason is modelling errors). Nevertheless, when the weather has been realised, it is possible to calculate, afterwards, the optimal solution that could have been obtained when all information would have been available right from the start (cf. e.g. Van Henten et al., 1997). Although this "dream pattern" cannot be reached in practice, it is a good benchmark to which other solutions can be rated. All other solutions will be sub-optimal, and it is interesting to compare the sub-optimality losses of various propositions.

In order to keep matters simple, here a very basic change in the weather relative to the nominal weather is used, by introducing a 80% solar radiation drop for two hours on the second day. The pattern and optimal solution are presented in Figure 2.

It is seen that the optimal value drops from 3.30 to 2.97, which must be attributed to the availability of less light. This is also seen from the lower biomass production (compare final values in bottom panels of Figures 1 and 2). There is also a clear effect on the optimal control pattern itself, as the heating stays off at the second day around sun-rise, as the optimisation 'knows' that there will be a light drop later on.

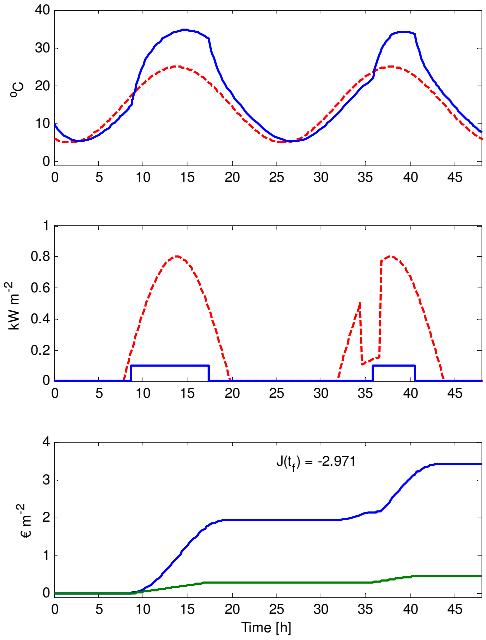

Fig. 2. Open loop optimal control with "real" weather. Meaning of all lines as in Figure 1.

5.3 Closed loop control using optimal state trajectories as set-point

Next, we turn our attention to closed loop control. One idea that comes into mind is just to use the nominal control pattern to control the greenhouse. It is easy to see that this form of open loop control is not a good idea, as this does not provide the necessary feed-back.

On the other hand, it is frequently felt that using the temperature patterns obtained for smooth nominal weather are good candidates as set-points for a local controller. Offering the nominal optimal temperature trajectories as set-points to a local PI controller under actual weather conditions, i.e. with the light dip on the second day results in J = -2.73. The controller performance is shown in Figure 3. Despite excellent tracking, it is clear that there is a loss compared to the best achievable control (-2.97). One reason for the loss is that there is heating on the second day when it is not required, as can be seen by comparison with Figure 2. Although this is partly an artefact of the current example, as the light drop affects the crop only, it is illustrative of what happens when a set-point derived for nominal conditions – which is the only possibility in practice - is chased under changed conditions. The other reason may be the dynamics introduced by the feed-back controller.

To more clearly see the effect of the introduced controller dynamics the PI controller is used with the nominal weather.

The pattern is very similar to Figure 3. We find that J = -3.09, which is a loss of 6% with respect to the optimal value of -3.30. This loss is solely due to the extra dynamics introduced by the feed-back loop, as all other information is exactly known.

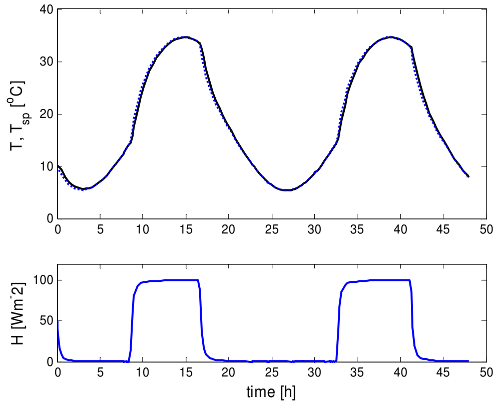

Fig. 3. PI control using optimal (nominal) temperature (T_{sp}, solid line) as set-point, with “real” weather. The state x_2 (dashed) is indicated by T , the control u by H.

5.4 Sensitivity to near-optimal set-points

When the optimal set-point trajectory derived from Figure 2 is approximated by a piece-wise linear approximate pattern as shown in Figure 4 – as an example of how implementation may look like in practice – we obtain J = -2.90 for nominal weather and J = -2.56 for actual weather. In both cases this is a loss of 6% compared to the correct pattern. This suggests that seemingly small differences in set-point can already lead to appreciable losses.

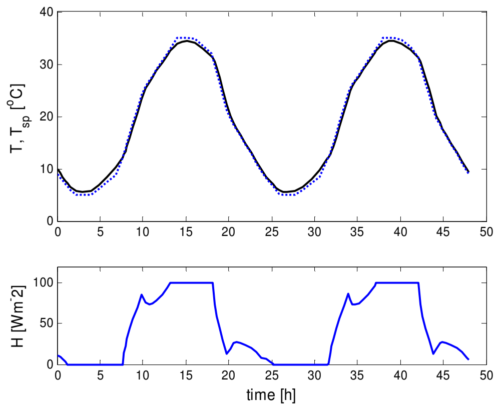

Fig. 4. PI control using an approximation of the optimal (*nominal*) temperature (T_{sp}, solid line) as set-point, with “real” weather. The state x_2 (dashed) is indicated by T , the control u by H.

5.5 Receding Horizon Control with co-state information

In optimal control problems the so-called co-state equation plays an important role. There are as many co-states or ad-joint variables (λ) as states. They can be interpreted as the sensitivity of the goal function to the states (Stengel, 1994), and can therefore be seen as shadow prices for a marginal increase in state value. The general co-state equation reads

$$-\dot{\lambda} = \left(\frac{\partial f}{\partial x}\right)^T \lambda + \left(\frac{\partial L}{\partial x}\right)^T, \quad \lambda(t_f) = \left(\left.\frac{\partial \Phi}{\partial x}\right|_{t_f}\right)^T. \qquad (7)$$

Working this out for the slow state (biomass) in the simple example gives

$$-\dot{\lambda}_1 = 0, \quad \lambda_1(t_f) = p_5 \,. \qquad (8)$$

This shows that in this simple example the marginal value of a unit of biomass is constant. This is because, in the simple example, the biomass increases linearly when light and temperature are constant. It should be noted that this is in agreement with the linear phase associated to full foliage coverage in classical expo-linear growth.

Two-time scale decomposition now suggests (Van Henten, 1994, van Straten et al., 2002; van Henten and Bontsema, 2007) that an appropriate goal function to be used on the short time scale is given by the adjusted goal function

$$J_{rhoc}(t, t_H) = \int_t^{t+t_H} \left(p_4 u_1 + \lambda_1 \dot{x}_1\right) dt \,, \qquad (9)$$

The integration is taken over the moving window from the current time t to $t+t_H$, with t_H the prediction horizon.
This approach is based on the observation that the co-state associated to the slow variable does not depend much on the actual realization of the weather. It can therefore be evaluated from an ensemble of possible weather patterns.

The first real application of this principle was for tomato (Tap, 2000). Note that the goal function (9) preserves the direct economic costs of the resources, and adheres to it a value that should be given to the state increment of the slow state, despite the fact that these revenues will only be obtained at the final time.

Based on this goal function, it is now possible to build a receding horizon optimal controller (RHOC). At any sampling instant the control trajectory is computed over a finite horizon t_H needed to locally optimize the goal function (9), based on an estimate of the current state, the current weather, and an expectation of the weather over the prediction horizon. An example is given in Figure 5. Only the first calculated control action is applied during the remainder of the sampling interval. Hence, in the situation of Figure 5 the heating will be turned on. When the next sampling instant is reached (in this case after 0.2 h) the optimisation is repeated, using a new estimate of the state derived from actual observations, and a new observation and expectation of the weather. In this way, both feed-forward as well as feed-back is provided. Note that the RHOC is not a standard Model Predictive Controller (MPC) as the goal function is different and has an economic basis, rather than a control performance basis. A forecast of the weather is needed, but as the horizon is short, this is much less a problem than when computing seasonal optimisations. In the calculations reported below, a so called lazy man prediction is used, meaning that *d* over the prediction horizon will be the same as the most recent measured value.

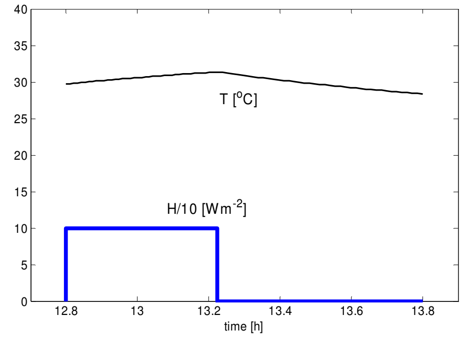

Fig. 5. Example of *predicted* optimal temperature trajectory and control pattern over the prediction horizon of 1 h, at t = 12.8 h

Figure 6 shows the ensemble of all predicted temperature trajectories over two days and also presents the optimal control actions trajectories as overlapping blocks. It is interesting to note that although the RHOC predicts switching off the heating somewhere half way the one hour horizon (Figure 5), the actual on-time is longer because at later times the original prediction is overruled by the newly calculated control at the next control interval.
Using the actually realized control, the realized cost function can be computed. The result is -2.82, which in this case is 5% less than the achievable value.

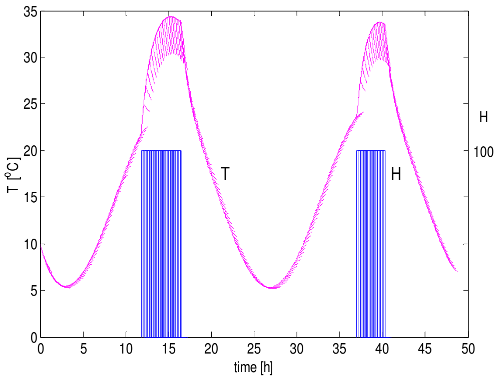

Fig. 6 RHOC temperature predictions and optimal control trajectories, with a control interval of 0.2 h and horizon of 1 h, real weather

5.6 PI control with day-night pattern

The classical greenhouse settings are a standard day-night pattern. Unlike Dai et al. (2006), who try to find these set-points by optimisation, using the greenhouse and a crop model, here, day-night settings 28 °C from 9:00 to 19:00, and of 5 °C during the night, with a one hour gradual change, were used. These values were inspired by the optimal result and were already forced in a direction to obtain more or less bang-bang heating. The result of -2.63 is slightly better than with the set-points obtained from optimal control, but still 7% below the RHOC controller, and 11% below the theoretical ideal.

6. SUMMARY OF THE RESULTS

Table 1. Summary of the results

Nominal Weather	*-J* profit	% loss
1 OL; Optimal control (5.1)	3.30	0
2 CL; PI, set-points from nominal dynamic optimisation (5.3)	3.09	6
3 CL; PI, set-points from approximated (stylised) nominal dynamic optimisation (5.4)	2.90	12
Real Weather	*-J* profit	% loss
4 OL; "Dream" pattern optimal	2.97	0
5 CL; PI, set-points from nominal dynamic optimisation (5.3)	2.73	8
6 CL; PI, set-points from approximated (stylised) nominal dynamic optimisation (5.4)	2.56	14
7 RHOC (5.5)	2.82	5
8 PI, fixed day-night (5.6)	2.63	11

For convenience, profit is reported here, which is –*J*.

7. DISCUSSION AND CONCLUSIONS

The use of any kind of on-line feed-back control always leads to losses with respect to the *a posteriori* optimal "dream" pattern. On the basis of a simple example we have shown how to evaluate the sub-optimality losses of various controller options.

Clearly, the example is a gross oversimplification of any real greenhouse-crop system, but it was deliberately chosen as such to clearly mark the features we wished to elucidate. We are quite confident that the trends observed are similar to those in the real world, and it is very challenging to repeat the loss evaluation for realistic crop-greenhouse models.

The results so far suggest that chasing fixed day-night set-points may lead to larger losses than if the set-points are obtained from optimisation under nominal disturbances, but even then the losses can be quite large. This can be attributed to the fact that the crop does not really need the target temperature when the light regime is different than expected. Sloppy approximation of the optimal controller set-points will also lead to losses. With deviating weather, performance of the PI controllers deteriorates further.

It should be noted that in all these cases the deviation from the nominal weather was much less than what can be expected in a real case, where much larger fluctuations can occur.

Receding horizon control on the basis of an adjusted goal function that uses the same economic parameters, derived from the optimal solution via two-time scale decomposition outperforms all other feed-back controllers, and is expected to yield the best economic result achievable in an on-line situation.

8. REFERENCES

Blasco, X., M. Martiñez, J. M. Herrero, C. Ramos, and J. Sanchis. 2007. Model-based predictive control of greenhouse climate for reducing energy and water consumption. Computers and Electronics in Agriculture 55 (1):49-70.

Bryson, A.E. 1999. Dynamic Optimization: Addison-Wesley.

Dai, J. , W. Luo, X. Qiao, and C Wang. 2006. Model-based decision support system for greenhouse heating temperature set point optimization. Nongye Gongcheng Xuebao/Transactions of the Chinese Society of Agricultural Engineering 22 (11):187-191

Ioslovich, I., and I. Seginer. 1998. Approximate seasonal optimization of the greenhouse environment for a multi-state-variable tomato model. Transactions of the American Society of Agricultural Engineers 41 (4):1139-1149.

Pasgianos, G. D., K. G. Arvanitis, P. Polycarpou, and N. Sigrimis. 2003. A nonlinear feedback technique for greenhouse environmental control. Computers and Electronics in Agriculture 40 (1-3):153-177.

Piñón, S., E. F. Camacho, B. Kuchen, and M. Peña. 2005. Constrained predictive control of a greenhouse. Computers and Electronics in Agriculture 49 (3):317-329.

Seginer, I. 1989. Optimal greenhouse production under economic constraints. Agricultural Systems 29 (1):67-80.

Stengel, R.F. 1994. Optimal Control and Estimation. Dover Edition ed: Dover Publ, Mineola, NY, USA.

Tap, R.F. 2000. Economics-based Optimal Control of Greenhouse Tomato Crop Production, Systems and Control, Wageningen Agricultural University, Wageningen, The Netherlands.

Van Henten, E. J., J. Bontsema, and G. Van Straten. 1997. Improving the efficiency of greenhouse climate control: An optimal control approach. Netherlands Journal of Agricultural Science 45 (1):109-125.

Van Henten, E.J. 1994. Greenhouse Climate Management: An Optimal Control Approach, Agricultural University Wageningen, Wageningen.

van Henten, E.J., and J. Bontsema. 2007. Time-scale decomposition of an optimal control problem in greenhouse climate management. Control Engineering Practice forthcoming.

van Straten, G., L.G. van Willigenburg, and R.F. Tap. 2002. The significance of crop co-states for receding horizon optimal control of greenhouse climate. Control Engineering Practice 10 (6):625-632.

Young, P. C., M. J. Lees, A. Chotai, W. Tych, and Z. S. Chalabi. 1994. Modelling and PIP control of a glasshouse micro-climate. Control Engineering Practice 2 (4):591-604.

APPENDIX

Model parameters for the illustrative example:

model:
p_1 = 7.5e-8 kg(dw) m^{-2} (Wm^{-2})$^{-1}$ ^{o}C^{-1} h^{-1} ; p_2 = 1 h^{-1} ; p_3 = 0.10 ^{o}C (Wm^{-2})$^{-1}$ h^{-1}.
goal function:
p_4 = 4.55e-4 € (Wh)$^{-1}$; p_5 = 136.4 € kg(dw)$^{-1}$.
initial conditions:
$x_1(0)$=0 kg(dw)m^{-2}; $x_2(0)$=10 ^{o}C.
PI controller:
$K_p = 40$; $K_I = 30$.