MASTER

Specific growth rate control for fed-batch fermentation of bakers' yeast

Ariaans, L.J.J.M.

Award date:
1992

Link to publication

EINDHOVEN UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING
Measurement and Control Group

SPECIFIC GROWTH RATE CONTROL FOR FED-
BATCH FERMENTATION OF BAKERS' YEAST

by: L.J.J.M. Ariaans

M.Sc. Thesis
carried out from 1 January1992 to 10 December 1992
commissioned by prof. dr. ir. P. Eykhoff
under supervision of ir. M. Keulers
date: 1 December 1992

Summary

The Group Measurement and Control of the Eindhoven University of Technology and the Technology Application Unit of the Unilever Research Laboratory cooperate in a project on fed-batch fermentation of bakers' yeast. As part of this project, a controller had to be designed for laboratory scale fermentations using synthetic media. The control should maintain a set-point on the specific growth rate and obey an upper bound on the ethanol concentration. Control of the oxygen concentration should be incorporated in the design as well. Process outputs are oxygen uptake rate, carbondioxide production rate, ethanol concentration in the off-gas and dissolved oxygen tension in the broth. Inputs are glucose feed rate, air flow and stirrer speed.

The specific growth rate is observed using the algorithm of Soeterboek (1992). The specific growth rate can be estimated sufficient accurately by this observer. Theoretical analysis and simulation studies indicate, that the growth rate is not excessively sensitive to inaccuracies in the physical parameters it uses. The ethanol concentration is estimated less accurately.

The highly non-linear process is linearised as far as possible by precompensations. The linearised process is controlled by a PI-controller. μ-control is satisfactory for oxidative growth on glucose only. If ethanol production or consumption occurs, control performance is deteriorated by inaccurate ethanol estimates. Replacing these estimates by predictions during critical phases of control offers no solution, as the prediction model proved too inaccurate. If the ethanol limit is exceeded, the control brings the ethanol concentration back in the permissable range by lowering the μ set-point. Overshoot on the ethanol concentration is too large to use this scheme for ethanol control. Performance as a safeguard against excessive ethanol concentrations is satisfactory.

During the experiments, a local oxygen controller was used. Changes of the air flow and stirrer speed disturb the measurements of the process outputs. When this was recognised, the oxygen controller was further used in manual mode only. Rate limiting on air flow and stirrer speed should be implemented in the future.

Experiments were carried out for oxidative growth on glucose only. The critical growth rate of the yeast is lower than was assumed. The yield for growth on glucose increases throughout the experiment, possibly because of some adaptation mechanism within the yeast cells. The experiments suggest a rather high value for the yield in steady state. This may be caused by measuring errors or by yeast behaviour not covered by the used model.

Response to μ changes involves unexpected dynamics. A linear model for this effect was identified for small steps. Incorporating this model in a dynamic, predictive precompensation reduces overshoot on step changes and speeds up the response. It is recommended to carry out further experiments to identify a model for the dynamics involved in larger steps on μ and to incorporate this in the control structure as well. The effects during ethanol consumption or production have not been studied experimentally. These deserve further research as well.

Contents

List of symbols

Symbol	Description	Value	Unit
α	stoichiometric constant	1.05	mole/mole
a	stoichiometric constant	2.1675	mole/mole
AF	air flow	-	l/min
β	stoichiometric constant	-	mole/mole
b	stoichiometric constant	3.65	mole/mole
c	stoichiometric constant	2.35	mole/mole
C_1	experimental constant	123.0	mole/h
C_2	experimental constant	0.7	-
C_3	experimental constant	0.25	-
CPR	CO_2 production rate (concentration)	-	mole/(l·h)
CPR'	CO_2 production rate, corrected for maintenance	-	mole/h
CPR_a	CO_2 production rate (amount)	-	mole/h
d	stoichiometric constant	-	mole/mole
e	stoichiometric constant	-	mole/mole
E	ethanol concentration, general	-	mole/l
E_{max}	maximum ethanol concentration	-	mole/l
f	stoichiometric constant	-	mole/mole
g	stoichiometric constant	0.36	mole/mole
G	glucose concentration	-	mole/l
GF	glucose feed rate	1.8920	l/h
G_{in}	glucose concentration in feed	2.22	mole/l
h	stoichiometric constant	-	mole/mole
hx	hydrogen content of yeast	-	-
i	stoichiometric constant	-	mole/mole
j	stoichiometric constant	1.8740	mole/mole
k	stoichiometric constant	1.6140	mole/mole
	sample instant	-	-
K_e	saturation parameter for growth on ethanol	$2.2 \cdot 10^{-3}$	mole/l
K_i	inhibition parameter	$2.0 \cdot 10^{-4}$	mole/l
	integral gain PI-controller	-	-
K_o	saturation parameter for oxygen uptake	$3.0 \cdot 10^{-7}$	mole/l
K_p	proportional gain PI-controller	-	-
K_s	saturation parameter for glucose uptake	$1.0 \cdot 10^{-4}$	mole/l

|			
l	stoichiometric constant	1.32	mole/mole
μ	specific growth rate, general	-	h^{-1}
μ'	specific growth rate, precompensation input	-	h^{-1}
μ_d	specific growth rate, μ-controller input	-	h^{-1}
μ_{set}	specific growth rate, E-controller input	-	h^{-1}
m	stoichiometric constant	0.68	mole/mole
ms	specific rate of maintenance reaction	0.015	mole O_2/ (mole X·h)
n	stoichiometric constant	-	mole/mole
nx	nitrogen content of yeast	-	-
O_{max}	oxygen saturation concentration	$23.84 \cdot 10^{-5}$	mole/l
OUR	oxygen uptake rate (concentration)	-	mole/(l·h)
OUR'	oxygen uptake rate, corrected for maintenance	-	mole/h
OUR_a	oxygen uptake rate (amount)	-	mole/h
ox	oxygen content of yeast	-	-
p	pole location	-	-
P	oxygen concentration	-	mole/l
q_e	specific ethanol production rate	-	h^{-1}
q_s	specific glucose consumption rate	-	h^{-1}
SS	stirrer speed	-	rpm
V	volume	-	l
VX_0	initial estimate of the amount of biomass	-	mole
X	yeast concentration	-	mole/l

1. Introduction

1.1. Practical setup

A technique to cultivate bakers' yeast consists of a so called batch phase and a fed-batch phase. In the batch phase, the yeast is put in a vessel, the fermenter. Also present in this fermenter are the substances the yeast needs for its growth. Air is led through the fermenter to provide the yeast with oxygen. Stirring of the mixture further enhances dissolving the oxygen in the mixture and ensures that the medium is homogenous. No nutrients are added for several hours. This characterises the batch phase. After, in our case, about 16 hours the process enters the second phase in which feeding substances are added to the mixture. Nothing is extracted from the fermenter. This is typical for the fed-batch phase. Yeast is harvested only at the end of the process. The batch phase is needed to enter the fed-batch phase in a well-defined manner. For a further introduction to fermentation processes, see for example (Bastin and Dochain, 1990).
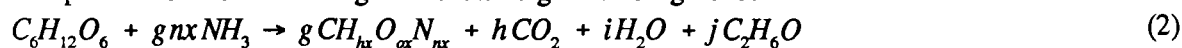
1.2. Reactions taking place

The yeast grows mainly on glucose. This is also by far the largest part of the fed-batch feed or substrate. If more glucose is added than the yeast can consume "normally", glucose is consumed at the expense of production of ethanol. If less glucose is added, this ethanol can be consumed again. If all ethanol has been consumed and the amount of glucose added is still less than the yeast can consume, the yeast simply grows at a lower rate. If the oxygen concentration is kept at a sufficiently high level, it hardly influences yeast growth. If it falls below roughly ten percent of its saturation value, there is a sharp decrease in the amount of glucose that can be consumed without producing ethanol. Apart from these reactions, a reaction seems to occur in which yeast is consumed. The energy produced in this way is used for maintenance of the yeast.

The reactions taking place are:

firstly oxidative growth on glucose

$$C_6H_{12}O_6 + aO_2 + bnxNH_3 \rightarrow bCH_{hx}O_{ox}N_{nx} + cCO_2 + dH_2O \tag{1}$$
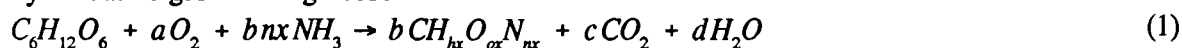
ethanol production occurs during fermentative growth on glucose

$$C_6H_{12}O_6 + gnxNH_3 \rightarrow gCH_{hx}O_{ox}N_{nx} + hCO_2 + iH_2O + jC_2H_6O \tag{2}$$

and ethanol consumption takes place during oxidative growth on ethanol

$$C_2H_6O + kO_2 + lnxNH_3 \rightarrow lCH_{hx}O_{ox}N_{nx} + mCO_2 + nH_2O \tag{3}$$
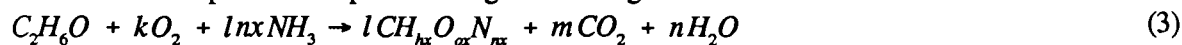
Finally, the maintenance reaction is given by

$$\alpha O_2 + CH_{hx}O_{ox}N_{nx} \rightarrow CO_2 + \beta H_2O + nxNH_3 \tag{4}$$

$C_6H_{12}O_6$ is the structure formula for glucose, C_2H_6O for ethanol and $CH_{hx}O_{ox}N_{nx}$ for yeast. The parameters hx, ox and nx represent the composition of the yeast. The parameters a to j, α and β specify the stoichiometry. The system (1), (2), (3), (4) is, from a bio-technical point of view, a highly simplified representation of the yeast activity. It is assumed however that the mentioned

system is accurate and detailed enough for the control problem at hand.

1.3. Definition of the control problem

For the study of certain bio-technological aspects of the yeast fermentation it is desired that the yeast grows according to a pre-specified pattern on the specific growth rate μ. This is what the controller should realise. Further, the ethanol concentration is not allowed to exceed a certain bound. The ethanol bound has a higher priority than the μ set-point. A sufficiently high oxygen concentration has to be maintained during the fermentation. The oxygen concentration and the specific growth rate are mutually related. In order to avoid degradation of the μ-controller performance by the oxygen controller or vice-versa it is decided to incorporate the oxygen controller in the current control problem as well.

The main costs involved in the production of yeast are the costs of glucose. Less important are the costs of the energy required for rotating the stirrer and blowing air through the broth. Minimising these costs is a secondary goal. At present, a tight μ response is considered most important.

During the fed-batch phase, the process can be influenced in three ways. The most important is the amount of glucose added per minute, the glucose feed rate (GF). The oxygen concentration can be manipulated to some extent by changing the rotation speed of the stirrer and by changing the amount of air that is blown through the broth every minute. The glucose feed rate GF, the stirrer speed SS and the air flow AF thus make up the three inputs of this process.

The air that flows out of the fermenter is analyzed by a mass spectrometer. This instrument measures the amount of oxygen extracted from the inlet air and the amount of CO_2 transferred to the air. Also, from the concentration of gaseous ethanol in the outlet air the concentration of ethanol in the broth is derived, assuming that certain exchange processes are in steady state. A probe measuring the oxygen concentration completes the setup, giving four main process outputs: oxygen uptake rate OUR, carbondioxide production rate CPR, dissolved oxygen tension (oxygen concentration normalised to its saturation value) DOT and ethanol concentration E.

A model based on Sonnleitner and Käppeli (1986) will be used as a nominal model. This model will also be used during simulations of the yeast behaviour. For further details see appendix 1.

1.4. Outline of this report

In order to control the specific growth rate, a measurement or an estimate of it should be available. An observer supplying an estimate for μ is presented in chapter 2. This observer can also be used to estimate the yeast concentration and ethanol concentration. The control structure based on simulations is described in chapter 3. The next chapter will be on some practical aspects of the implementation. Readers interested in control only can skip this chapter. Those intending to set up experiments in a similar fashion as we did, can benefit from the experience gained. In chapter 5 the experimental results and the ways in which they have affected the control structure will be discussed. Conclusions and recommendations complete the main part of this report.

2. Observer

2.1. Introduction

The observer presented in this chapter is taken from (Soeterboek, 1992). Changes to this original observer are summarised in appendix 2. It estimates on-line the biomass growth rate and the ethanol production rate from oxygen uptake rate and carbondioxide production rate measurements. By integration of these rates, the amounts of biomass and ethanol in the broth are estimated. The glucose consumption rate and the amount of glucose in the broth could be estimated along the same lines, but will be shown to be numerically ill-conditioned. Simulation studies confirm, that these glucose related quantities can not be accurately observed on line.

In the following section, the algorithm used by the observer is briefly stated. The third section will give a sensitivity analysis with respect to changes in the stoichiometric constants. Conclusions follow in the last paragraph of this chapter.

2.2. Algorithm used

This observer models the yeast activity according to (1), (2), (3), (4). Two measurements related to the reaction rates are available on line: the oxygen uptake rate OUR_a and the carbondioxide production rate CPR_a. The unknowns to be constructed are the rates for each of the reactions (1), (2), (3), (4), so it faces the problem of determining four unknowns from two data. To solve this problem, the specific rate of the maintenance reaction is taken to be a known constant. As (2) and (3) are mutually exclusive, at least one of the associated rates should be zero. Which of (2) and (3) is active can be determined from the sign of the ethanol production rate: if the production rate is positive, ethanol is produced and (2) is followed, else (3) is followed. Apart from this knowledge, the stoichiometric constants in (1), (2), (3), (4) are needed.

2.2.1. Rates estimates

If there is no ethanol production or consumption, the glucose consumption rate ($q_s \cdot VX$), biomass growth rate ($\mu \cdot VX$) and the ethanol production rate ($q_e \cdot VX$) are given by

$$
\begin{bmatrix} q_s \cdot VX \\ \mu \cdot VX \\ q_e \cdot VX \end{bmatrix} = \begin{bmatrix} \dfrac{1}{2a} & \dfrac{1}{2c} \\ \dfrac{b}{2a} & \dfrac{b}{2c} \\ 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} OUR_a - ms \cdot VX \\ CPR_a - \dfrac{ms}{\alpha} \cdot VX \end{bmatrix} + \begin{bmatrix} 0 \\ -\dfrac{ms}{\alpha} \cdot VX \\ 0 \end{bmatrix} \tag{5}
$$

For ethanol production or consumption, similar expressions are given in appendix 2. As mentioned before, the sign of $q_e \cdot VX$ is used to choose between ethanol production and ethanol consumption. Because $q_e \cdot VX$ is not available before the new rates have been calculated and the choice between ethanol production and consumption needs to be made to calculate the new rates, the sign of $q_e \cdot VX$ at the previous sample instant is used instead. If the sign of the new $q_e \cdot VX$ differs from that of the previous $q_e \cdot VX$, the assumption regarding ethanol production or

consumption was apparently wrong and the observer re-estimates the rates, using the sign of the new $q_e \cdot VX$. If this leads again to a contradiction, a warning message is displayed and the last estimates are used. If ethanol consumption is possible, i.e. $q_e \cdot VX$ is negative, but the amount of ethanol is less than 10^{-4} times the amount of biomass in moles, no ethanol consumption is assumed.

2.2.2. Amounts estimates

The estimates for the amounts VG, VX and VE are obtained from the rates estimates by integration: given the growth rate of, say, the biomass VX, the total amount of biomass at time $t = k \cdot T_s$ is obtained by

$$VX(kT_s) = VX\left[(k-1)T_s\right] + \int_{(k-1)T_s}^{kT_s} \left(\mu \cdot VX\right)(t)dt \tag{6}$$

This approach has the important drawback, that errors of any kind tend to accumulate. Periodic resetting using off-line measurements can prevent things from getting completely out of hand. This does not elimate the impact on control performance, though. Therefore the integration algorithm should be chosen carefully. A common technique to improve the accuracy of numerical integration is using high order integration algorithms. However, the sampling rate of the measuring equipment is known to be rather low. Fitting high order polynomials to the last few samples, which is what most of these techniques in fact do, is therefore not likely to give good interpolation results.

Several low order numerical integration algorithms have been investigated. Taking

$$f(x) \equiv \dot{F}(x) \tag{7}$$

the first algorithm, known as Euler backward or rectangular interpolation, is

$$\tilde{F}(k \cdot T_s) = \tilde{F}((k-1) \cdot T_s) + T_s \cdot f(k \cdot T_s) \tag{8}$$

It approximates the integrand f between two samples by a constant. The trapezoidal rule fits a straight line between the last two samples of f. This results in

$$\tilde{F}(k \cdot T_s) = \tilde{F}((k-1) \cdot T_s) + \frac{T_s}{2} \cdot \left[f(k \cdot T_s) + f((k-1) \cdot T_s)\right] \tag{9}$$

A variant of this is proposed in Soeterboek (1992). The linear interpolation is changed to a linear extrapolation to the next sample instant. Inevitable delays in the measuring equipment can thus be overcome to some extent. This yields the integration rule

$$\tilde{F}(k \cdot T_s) = \tilde{F}((k-1) \cdot T_s) + \frac{T_s}{2} \cdot \left[3 \cdot f(k \cdot T_s) - f((k-1) \cdot T_s)\right] \tag{10}$$

The performance of these rules was tested in a simulated fed-batch in which both ethanol production and consumption occurred. The following figures show the true ethanol production rate evolution and the various interpolations. For clarity, they have been corrected for delays.
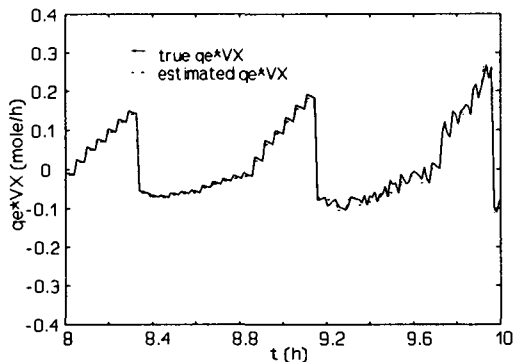
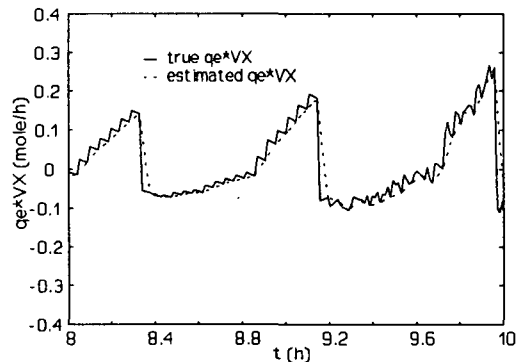figure 1: Interpolations using Euler backward

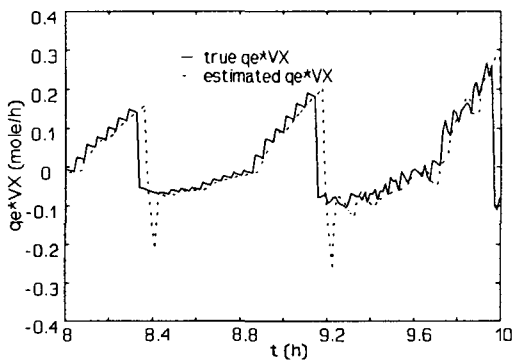figure 2: Interpolations using trapezoidal rule
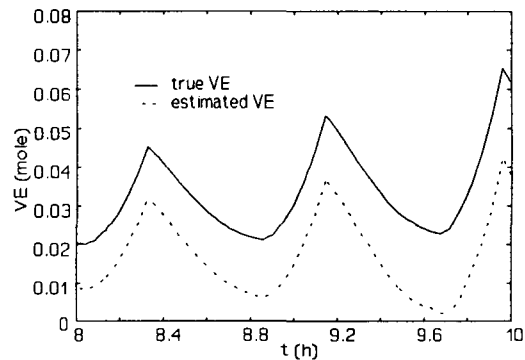

figure 3: Linear extrapolations

figure 4: Integral with Euler backward

Comparing figure 1, figure 2 and figure 3 it is clear that best results can be expected from Euler backward. Even this optimal choice can be expected to be structurally too small: the rectangular interpolation approximates the integrand fairly well, but it is always too small at the beginning of the sampling interval. This is confirmed by the actual integration results (figure 4).

The explanation for this phenomenon is as follows: at every sample instant the glucose feed rate is updated according to the newly available measurements. As yeast growth calls for exponentially increasing glucose feed rate, the overall trend in the glucose feed rate will be to increment GF at every sample instant. The reaction in the glucose concentration is nearly instantaneously. As a result, the various growth rates change instantaneously as well. During a sample interval, the yeast grows and the need for glucose increases slightly. As the glucose supply remains constant, this results in slightly decreasing growth rates. In the simulations, new measurements are taken at the end of every sample interval. Then this decrease has reached its maximum value.

For a least two reasons correcting for this decrease is tricky. Most important, in practice one can not rely on the samples being taken at the end of the sample interval. The sample rate of the measuring equipment will probably be a bit higher than that of the registration equipment, so that samples will not even be taken at a constant location within the interval. The other problem is the risk of becoming overconfident in the simulation model. The reaction to changes in the glucose feed rate may in reality not quite be as predicted by the model. Therefore it is impossible to give a reliable quantitative compensation.

2.2.3. Specific growth rate estimate

Once the biomass growth rate and the current biomass amount are available, the specific growth rate is simply obtained from

$$\hat{\mu} = \frac{\mu \cdot \hat{VX}}{\hat{VX}} \tag{11}$$

2.3. Sensitivity analysis

The stoichiometry of the observer can be changed, on line if required. To get an impression of the robustness of the observer, it is desirable to have some data regarding its sensitivity to changes in the stoichiometric constants. The reader is referred to the list of symbols for the nominal stoichiometry.

The sensitivity of a quantity y to changes in a quantity x is generally defined as

$$S_{y,x} = \frac{\partial y}{\partial x} \cdot \frac{x}{y} \tag{12}$$

Provided y(x) can be linearly approximated, the value of $S_{y,x}$ is the ratio of the relative change in y and the relative change in x: an increase in x of 5% will lead to an increase in y of $S_{y,x} \cdot 5\%$.

In the following sensitivity analysis, the dissolved oxygen tension is taken of the order 10%. For this value, the pathways being followed and the maximum reaction rates are insensitive to changes in the oxygen concentration.

2.3.1. Rates estimates

The values of OUR_a and CPR_a are not constant throughout the fed-batch phase. Apart from the maintenance influence in (5) however, the sensitivity values for the stoichiometric constants are independent of the absolute value of OUR_a and CPR_a and depend only on their ratio. Let OUR' and CPR' be defined by

$$OUR' = OUR_a - ms \cdot VX$$
$$CPR' = CPR_a - \frac{ms}{\alpha} \cdot VX \tag{13}$$

then the above argument holds exactly for OUR' and CPR'. For further details and a derivation of this claim, see appendix 2.
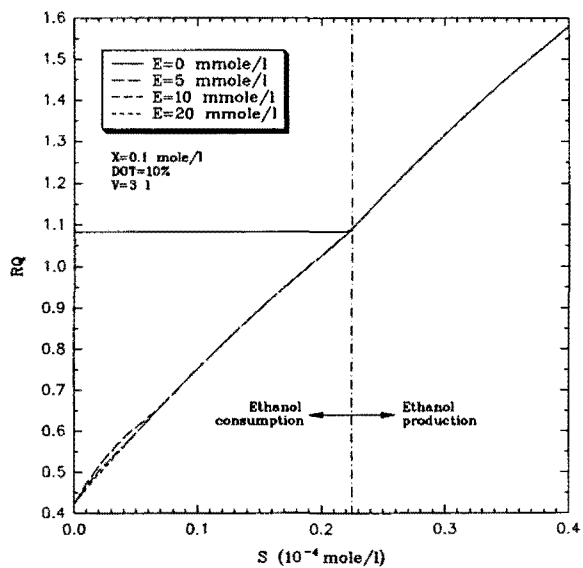
In figure 5 the value of the respiration quotient

figure 5: Respiration quotient vs. glucose concentration

RQ' is plotted. RQ' is defined as

$$RQ' \doteq \frac{CPR'}{OUR'} \tag{14}$$

From figure 5, practical extreme values for RQ' are taken. These are summarized in table A.1, appendix 2. It is recognised that especially during the batch phase, RQ' values may far exceed the upper bound given, but during the fed-batch phase, for which the controllers will be designed, this is highly unlikely to occur.

It can be proven, that worst case sensitivity values for $q_s \cdot VX$ and $q_e \cdot VX$ are reached at boundary values for RQ', irrespective of the absolute values of OUR' and CPR', see appendix 2. Best case values are either equal to zero or reached at the other end of the RQ'-range. This gives the sensitivity values presented in table A.3 and table A.4, appendix 3.

In appendix 2 it is also shown, that the sensitivity values for $\mu \cdot VX$ depend both on RQ' and on OUR'/VX. According to the simulation model used, OUR'/VX is always equal to 0.2 when there is no ethanol consumption and it varies between 0 and 0.2 for ethanol consumption, always assuming that the dissolved oxygen tension is of the order 10%. A discussion of these results is postponed to paragraph 2.3.3.

The sensitivity analysis with respect to the changes in ms and α is a bit more complicated. Let $q_s \cdot VX$ be written as

$$q_s \cdot VX = A_{11} \cdot [OUR - ms \cdot VX] + A_{12} \cdot [CPR - ms \cdot VX/\alpha] \tag{15}$$

with A_{11} and A_{12} a function of the pathways being followed, then

$$\frac{\partial(q_s \cdot VX)}{\partial ms} \cdot \frac{ms}{q_s \cdot VX} = (-A_{11} - A_{12}/\alpha) \cdot \frac{ms}{q_s} \tag{16}$$

From this it follows, that the sensitivity of $q_s \cdot VX$ is at its maximum for the smallest value of q_s and at its minimum for the largest value of q_s. Without going into detail about the dynamics of the yeast process, it can be stated, that the largest value of q_s is -0.5, reached asymptotically for high glucose concentration and the minimum value is 0, reached when there is (practically) no glucose. The latter situation applies to ethanol consumption or neither ethanol consumption nor production. The value of q_s at the switching point of ethanol consumption and production depends on the oxygen concentration in the broth. For a dissolved oxygen tension in the order of 10%, this value is according to the simulation model fairly accurately equal to 0.0923 [mole glucose/(hour·mole biomass)]. Changes in α are accounted for by

$$\frac{\partial q_s \cdot VX}{\partial \alpha} \cdot \frac{\alpha}{q_s \cdot VX} = \frac{A_{12} \cdot ms}{\alpha \cdot q_s} \tag{17}$$

This yields the results of table A.5.

In order to do the same for the ethanol production rate, the boundary values for q_e need to be known. Under normal conditions these are 0 and 0.66 for ethanol production, 0 and -0.12 for ethanol consumption. This gives the results of table A.6.

For the biomass growth rate the following holds
Here A_{21} and A_{22} take the place of A_{11} and A_{12} in the glucose consumption rate analysis. When ethanol is produced, values of μ range from 0.3225 h^{-1} to 0.4693 h^{-1}. For ethanol consumption

$$\frac{\partial \mu \cdot VX}{\partial ms} \cdot \frac{ms}{\mu \cdot VX} = -\frac{A_{21} + A_{22}/\alpha + 1/\alpha}{\mu} \cdot ms$$

$$\frac{\partial \mu \cdot VX}{\partial \alpha} \cdot \frac{\alpha}{\mu \cdot VX} = \frac{(A_{22} + 1) \cdot ms}{\alpha \cdot \mu} \tag{18}$$

these values are in the range -0.01429 h^{-1} to 0.3225 h^{-1} and for production nor consumption these values are -0.01429 h^{-1} and 0.07798 h^{-1}. Results are summarized in table A.7.

2.3.2. Amounts estimates

As the only relation of the amounts estimates to the stoichiometric constants is given by their dependence on the rates estimates, the results obtained in the previous paragraph are applicable to the amounts estimates as well. Some additional remarks with respect to the conditioning and robustness of the integrations are made in this subsection.

The total amount of biomass is the least sensitive to accumulation of errors. Under normal conditions, the growth rate of the biomass is always positive, so that the relative error stays within bounds. The ethanol production rate can be both positive and negative. If this occurs, the relative error due to the above mentioned imperfections can grow without bounds. The numerical problem is ill-conditioned in this case. In the previous paragraph it was pointed out, that conditioning problems may also occur with respect to changes in some stoichiometric constants if the ethanol consumption or production is near zero. The glucose consumption rate is always positive or at least not negative. Still, the integration to obtain the total amount of glucose is always ill-conditioned, because the glucose consumption rate is compensated with a glucose feed that nearly fully compensates the glucose consumption. The net rate of change in the glucose amount, which is the difference of these two rates, is a few orders smaller than its two constituents. For this reason, the glucose amount estimate is numerically ill-conditioned and should not be used.

Sensitivity values for the amounts estimates can not be given in general terms, as they depend too much on the way in which ethanol production and consumption have alternated during the fed-batch so far. Generally speaking, a steady production of ethanol will make the ethanol estimates far less sensitive to parameter changes than alternating ethanol production and consumption, which keeps the ethanol amount comparatively low, which increases the relative deviations. Another problem is, that the rates estimates can be insensitive to changes in a parameter during ethanol production, while these changes do have influence during ethanol consumption. These considerations point out, that the effect of parameter changes on the amounts estimates depends heavily on the history of the yeast. Therefore simulations will be carried out in order to evaluate the sensitivity in a practical situation.

2.3.3. Conclusions

From table A.2 can be concluded, that the biomass growth rate can be very sensitive to all parameters involved during ethanol consumption. This will happen, when the biomass growth is almost zero. This can only be the case during ethanol consumption. Theoretically, the biomass growth rate can be close to zero for ethanol production, but this occurs only for very low oxygen

concentrations. We assume, however, that the dissolved oxygen tension is about 10%. A biomass growth rate almost equal to zero can occur in practice, after it is decided to shut off the glucose feed to consume ethanol present in the broth as quickly as possible. In all other situations, the biomass growth rate can be assumed to be clearly greater than zero.

The best case values in table A.2 are all equal to 0 for ethanol consumption. They occur when OUR' is small enough to let the term $-ms/\alpha$ completely dominate. There will be practically no biomass growth, neither on glucose, nor on ethanol, but biomass will be consumed in the maintenance reaction. Changes in the stoichiometric constants do not influence $-ms/\alpha$, so the sensitivity is 0.

According to table A.3, the ethanol production rate can be very sensitive to changes in a and c. This happens both during ethanol production and ethanol consumption, when the ethanol production rate is small, so that small absolute deviations in $q_e \cdot VX$ will lead to large relative changes.

Finally, the accuracy of the glucose consumption rate estimate is potentially endangered by inaccurate values of k and m, see table A.4. This is only of interest during ethanol consumption. As in the previous two paragraphs, sensitivity values go to infinity if the glucose consumption rate goes to 0. This can occur in practice in a phase in which glucose feed was shut off to give preference to ethanol consumption.

In all other cases the worst case behaviour of the sensitivity values is satisfactory: relative errors in rates estimates are of the order of relative errors in the stoichiometric constants used or better.

The values of $\pm\infty$ in table A.5, table A.6 and table A.7 occur, when the associated rate is equal to zero. As mentioned above, a biomass growth rate equal to zero is unlikely in practice, the glucose consumption rate can be equal to zero in some instances and the ethanol production rate being equal to zero is not at all uncommon.

As was the case with the stoichiometric constants, the reaction rates are not too sensitive to inaccuracies in ms and α, apart from the sensitivity problems that occur when rates go to zero.

The interpretation of the above mentioned values is strongly dependent on the situation to which they are applied. The worst case values of $\pm\infty$ may seem dramatic. They pose problems indeed, if the observer is part of a control structure which aims at keeping a growth at or near zero. In other cases, the values of $\pm\infty$ are less likely to give severe problems.

It is nearly impossible to derive theoretically some useful statements about the robustness of the amounts estimates. Too many variable factors play a role here. The practicality of these estimates is best assessed by means of simulations.

2.4. Simulation results

To evaluate the observer performance and in particular its sensitivity to changes in the stoichiometric constants, some simulations have been carried out. In these simulations, the input data for the process were generated by existing controllers. A SISO gpc-controller was used to keep the dissolved oxygen tension at 10% by varying the stirrer speed. Another SISO gpc-controller kept the respiration quotient at setpoint by manipulating the glucose feed rate. The RQ setpoint was switched several times between 0.8 and 1.3 to ensure that both ethanol production and ethanol consumption occurred during a simulated experiment. These controllers were used because they were readily available in the MatLab-TUE environment. The application of GPC controllers to the fermentation process is described in (Keulers and Reyman, 1991). Multiplicative noise was added to the OUR_a and CPR_a measurements: the true values were multiplied by $1+\eta$ where η was normally distributed with zero mean and $\sigma=0.01$.

2.4.1. Nominal model

The results obtained when the simulation model and the observer use the same stoichiometry are given in figure A.4 and figure A.5 in appendix 4. The specific growth rate is plotted in figure 6.
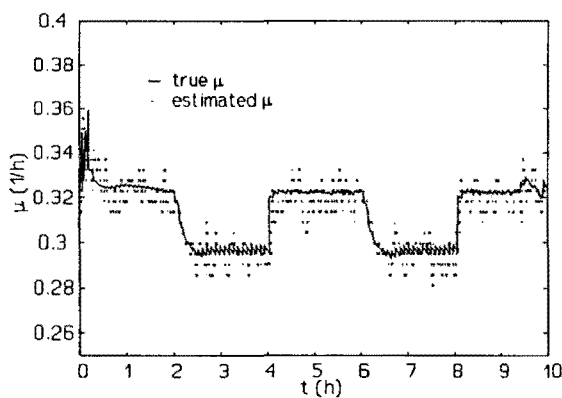

figure 6: Specific growth rate estimate

Both the biomass concentration and the ethanol concentration are estimated fairly well by the observer. Differences are somewhat larger for the ethanol concentration. Agreeing with the expectation formulated in 2.2.2, the ethanol estimate is smaller than the true value. Both estimates are rather insensitive to noise. This is to be expected, as the low pass characteristics of the integrations involved filter out most of the noise. The specific growth rate estimate is clearly more sensitive to noise. For μ no noise suppression mechanism whatsoever is present in the observer. This shows in figure 6.

2.4.2. Disturbed model

To test the sensitivity of the observer to uncertainties in the stoichiometry, the real situation would be approximated best, if the stoichiometric constants of the observer were kept at fixed values and the constants of the simulation model were changed. For two reasons preference was given to changing the stoichiometric constants of the observer:

● Changing the simulation model requires recalculation of the controller actions for each simulation. This turned out to be rather time-consuming.

● In some cases, the controllers lost control, so that the simulation results were not comparable to the other simulations.

The same set of (noisy) OUR_a and CPR_a measurements as in the previous subparagraph was taken for each experiment. All constants involved were changed by $\pm 5\%$.

The biomass growth rate and indirectly the amount of biomass appear to be the most sensitive to errors in a during ethanol production, b under all circumstances and c during ethanol consumption. This is expressed in an increasing offset between the estimated amount of biomass and the true amount. Even in these cases however, the observer keeps fairly well track of the true amount. The ethanol production rate is estimated less accurately than the biomass growth rate. Especially deviations in a and c, both during ethanol production and consumption, influence the rate estimates. This leads to a difference between the true ethanol amount and the estimated ethanol amount in the broth of approximately 50%. Changes in other constants do not have effects as pronounced as those caused by changes in a or c.

For illustration, the results obtained when a was increased by 5% are plotted. The specific growth rate estimate is shown in figure 7. The corresponding biomass and ethanol estimates are shown in figure A.6 and figure A.7, appendix 4.
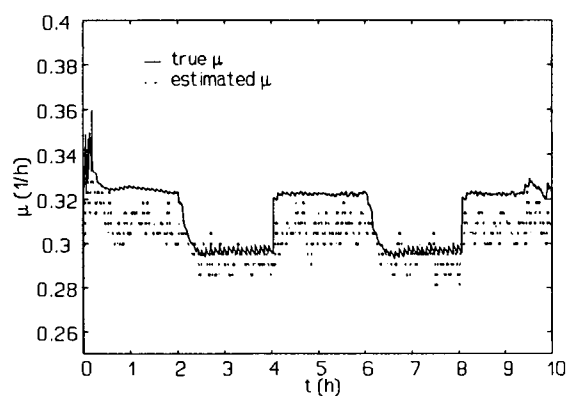

figure 7: Specific growth rate, $a=1.05 \cdot a_{true}$

2.5. Conclusions

As long as the biomass growth rate is greater than zero, sensitivity problems are not to be expected. During ethanol consumption and hardly any growth on glucose, all stoichiometric constants involved can give sensitivity problems. Simulation studies point out, that the estimated amount of biomass is influenced significantly by the value of a during ethanol production, the value of c during ethanol consumption and the value of b under all circumstances.

The ethanol production rate is estimated less accurately than the biomass growth rate, but this is caused by noise-sensitivity, not by parameter-sensitivity. Simulation results show that the ethanol production rate estimate is especially sensitive to changes in a and c, both during ethanol consumption and production, which is in close agreement with the theoretical analysis.

figure 8: Control structure based on simulation results
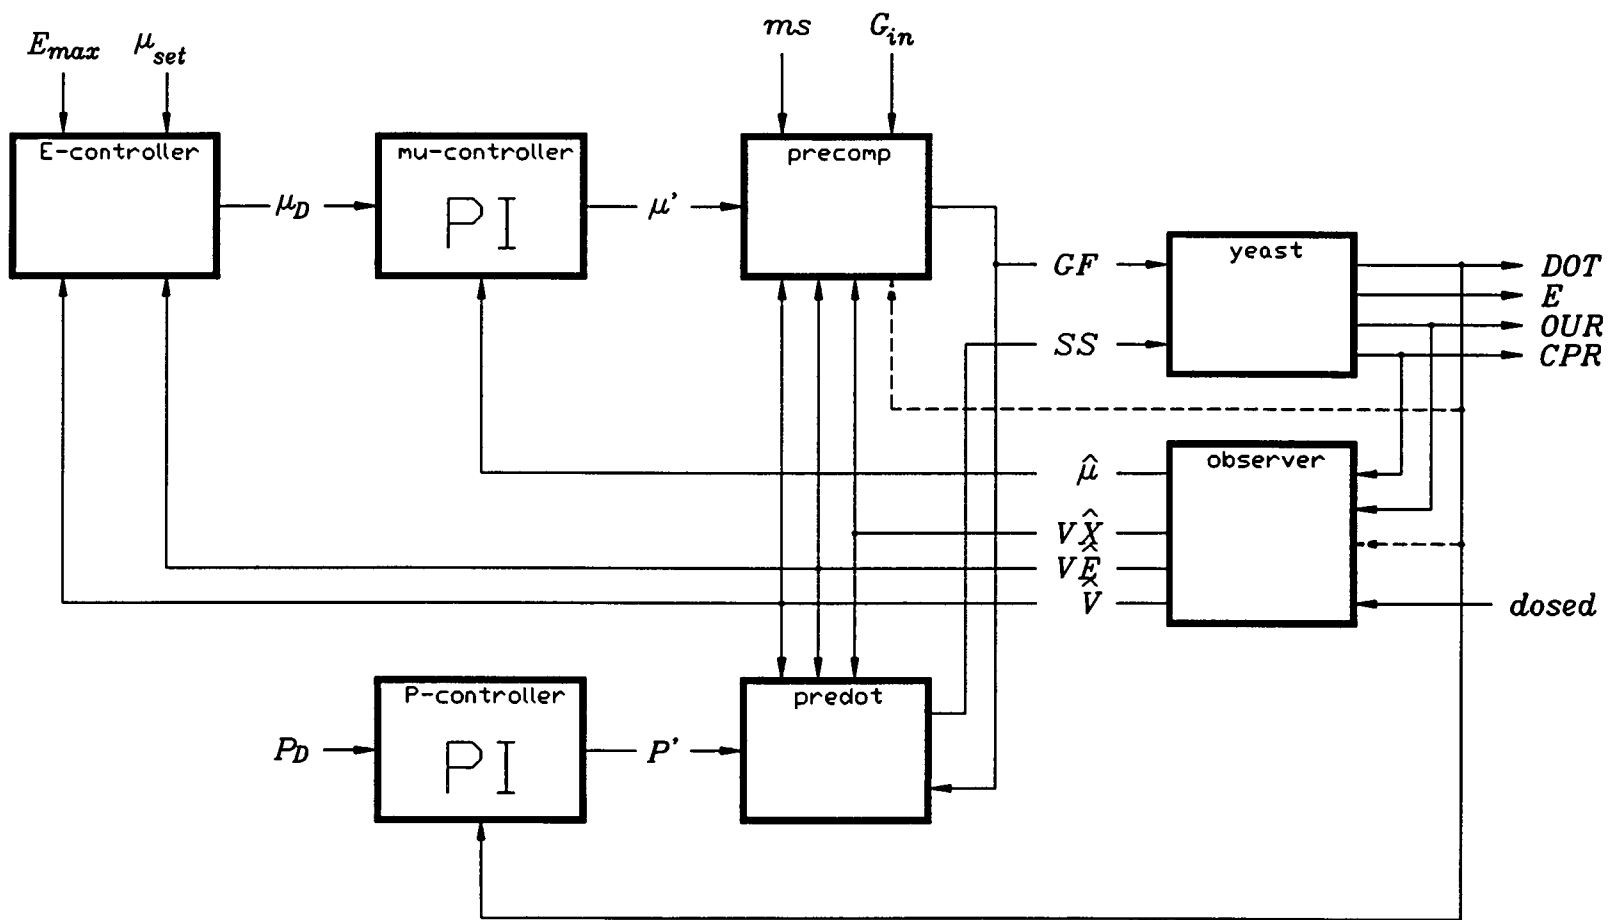
3. Control structure

3.1. General observations

It is a well known fact, that controllers for fed-batch fermentation processes need increasing gain throughout the fed-batch phase. This simply stems from the fact, that the yeast needs more glucose and more oxygen towards the end of the fed-batch phase, because the total amount of yeast has increased. Another complicating aspect of fermentation control is, that situations of ethanol production and consumption need quantitatively or, depending on the control goals, even qualitatively different control actions. The yeast process is a highly non-linear process. Therefore the behaviour in one working point will probably not correspond to the behaviour in another working point, even if the same metabolic pathways are followed. Aging of the yeast is a further inconvenience in the development of satisfactory controllers.

It is not unusual to simply accept the performance degradation resulting from these factors, (Agrawal *et al.*, 1989). A more satisfying approach is to use adaptive schemes. Many authors have used this approach, for instance (Landau *et al.*, 1990), (Pomerleau *et al.*, 1989), (Poulisse and Van Helden, 1985), (Verbruggen *et al.*, 1985), (Williams *et al.*, 1986). The inherent disadvantage of adaptive schemes is that they are always too late: they react only when things go wrong in some sense. They do not attempt to prevent things from going wrong in the first place. They also have the risk of non-converging adaptations, which may eventually result in total loss of control.

The approach adopted in the following control structure design is to linearize the process dynamics as far as possible by pre- or postcompensations. These compensations contain the available a priori knowledge about the yeast process.
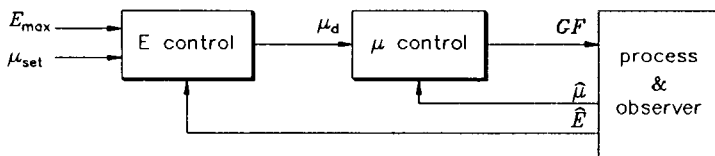
figure 9: Basic control structure idea

The basic idea of the structure is shown schematically in figure 9. As long as the ethanol limit is not exceeded, the E control simply connects μ_{set} to μ_d. If the ethanol limit is exceeded, μ_d is lowered, so that the ethanol production is stopped and ethanol consumption will occur. An advantage of this structure is, that the priority relation between ethanol limit and specific growth rate set-point is reflected automatically by this setup. Another potential advantage is that, with this structure, both μ-control and E-control can be realised. μ-control is ensured by setting the ethanol limit to an extremely large value, so that the E-control will never be activated. Setting the μ set-point to a value well above the critical value will result in violation of the ethanol limit. μ_d will be lowered until, hopefully, the ethanol limit is just maintained.

3.2. Glucose flow precompensations

Close inspection of the simulation model shows, that under practical conditions, the specific growth rate can be considered decoupled from the oxygen concentration, if the dissolved oxygen tension is sufficiently high, say at 10%. The bottleneck governing the oxygen influence on the reaction rates is proportional to (see appendix 1)

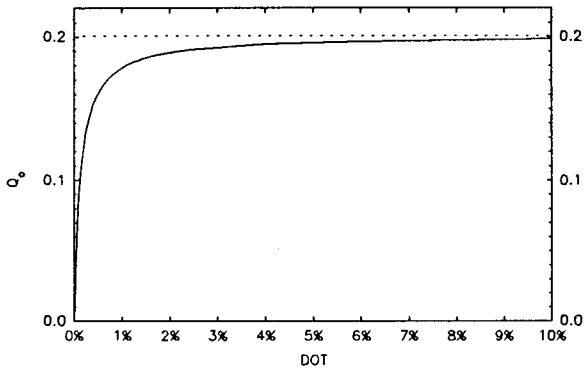$$Q_o \doteq \frac{P}{K_o + P}, \quad K_o = 3.0 \cdot 10^{-7} \text{ mole/l} \tag{19}$$

figure 10: Oxygen concentration influence on bottleneck

In order to get an impression of the influence of P on this quantity, its value is plotted in figure 10. In this figure it is demonstrated, that moderate fluctuations in P if DOT is approximately 10% have hardly any influence on the bottleneck. In practice, assuming no saturation on AF and SS, the dissolved oxygen tension will always be of the order 10% or larger, so that the oxygen concentration can be considered decoupled from the specific growth rate dynamics.

It seems likely, that the yeast's need for glucose is proportional to the total amount of yeast. This is confirmed by analyzing the simulation model. GF is therefore made proportional to $V\hat{X}$ obtained from the observer. Interpreting the maintenance effect as suggested by (1), (2), (3), (4), some additional glucose feed is needed. The yeast grows faster as a result of this extra glucose. This extra growth should exactly compensate the yeast consumption in the maintenance reaction. In steady state, the amount of glucose added per minute should be equal to the amount of glucose consumed per minute. If only oxidative growth on glucose occurs, the ideas mentioned so far result in the first precompensation:

$$GF = \left(\mu_d + \frac{ms}{\alpha} \right) \cdot \frac{V\hat{X}}{b \cdot G_{in}} \tag{20}$$

This precompensation is suggested by (O'Connor et al., 1992). It will be referred to in the following as precompensation 1. The stoichiometric constant b can be regarded as the yield of the fermentation process in mole biomass / mole glucose. G_{in} is the concentration of glucose in the feed (mole/l) and the glucose feed rate is expressed in l/h.

If the assumptions made hold indeed, this scheme is stable: if the specific growth rate is less than the set-point, more glucose is added than is consumed. Glucose is accumulated, its concentration grows and the yeast will grow faster, until the set-point for μ is reached. Simulations point out, that this steady state is reached well within a sample interval. Therefore the control structure in general and the precompensation in particular need not compensate for any dynamics in this adaptation process: changing the glucose feed rate has an instantaneous effect on the glucose concentration and the specific growth rate. This simplifies the control design significantly.

Precompensation 1 needs the following a priori knowledge and makes the following assumptions with respect to yeast dynamics
- the glucose concentration reaches its steady state within a fraction of the sampling interval
- the specific maintenance rate is known at every instant
- the yield of the oxidative growth on glucose (stoichiometric constant b) is known.

If the first assumption does not hold, the PI-control part of the μ-control faces additional dynamics. This is partly why it was incorporated in the control in the first place. If the maintenance rate is not exact, a constant offset on the realised μ will occur. This will be counteracted by the integral part of the PI-controller. The third assumption finally will give changing offsets on μ if it turns out to be invalid.

This precompensation has the drawback, that it only accounts for two of four possible pathways: oxidative growth on glucose and maintenance. Ethanol consumption and production are not considered. Taking into account these paths as well gives the second precompensation:

$$\lambda_p = \frac{\mu_d + \frac{ms}{\alpha} + (g-b) \cdot z_0}{g}, \quad \lambda_o = \frac{\mu_d + \frac{ms}{\alpha} - \frac{la}{k} \cdot z_0}{b - \frac{la}{k}}, \quad \lambda_e = \frac{\mu_d + \frac{ms}{\alpha} - 0.13 \cdot l \cdot z_1 \cdot z_2}{b} \tag{21}$$

$$GF = \max\{\lambda_p, \lambda_o, \lambda_e\} \cdot \frac{VX}{G_{in}}$$

z_0, z_1 and z_2 are defined as

$$z_0 \doteq \frac{0.2}{a} \cdot \frac{P}{K_o + P}, \quad z_1 \doteq \frac{E}{K_e + E}, \quad z_2 \doteq \frac{K_i}{K_i + G} \tag{22}$$

z_1 should be reconstructed using observed values for E. z_2 should be approximated by a constant close to 1 as normally G is much smaller than K_i. It is also the only possibility, because the value of G can not be measured or observed. Using this simplification, all data required for this precompensation are available on-line.

λ_p represents ethanol production. If λ_o is chosen, ethanol is consumed and the oxygen bottleneck is completely filled. This will be called oxygen limited ethanol consumption. If λ_e is active, ethanol is consumed and the oxygen bottleneck is not completely filled. The ethanol consumption is limited by the ethanol concentration. The situation of no ethanol production or consumption is considered a special case of ethanol limited ethanol consumption.

Theoretically the incorporation of all paths in the precompensation may seem a small step, however the amount of required a priori knowledge has increased substantially. First, the critical μ, i.e. the μ at the switching point between ethanol consumption and ethanol production, needs to be known. This μ, say μ_c, is the intersection of λ_p, λ_o and $\lambda_e(E=0)$. Secondly, the ethanol consumption rate as a function of the glucose and ethanol concentrations is used. This rate appears in (21) as

$$0.13 \cdot l \cdot z_1 \cdot z_2 \tag{23}$$

Finally, the yield of all pathways is needed. For ethanol consumption and production these are l mole biomass / mole ethanol and g mole biomass / mole glucose.

Summarizing the assumptions for the second precompensation:
- the glucose concentration reaches its steady state within a fraction of the sampling interval.
- the specific maintenance rate is known at every instant.
- the yields of all pathways are known.
- the critical μ is known.
- the ethanol consumption rate as a function of E and G is known.

This new precompensation will be called precompensation 2.

The performance of both compensations was tested by incorporating them in the control structure of figure 8. The E-controller in this figure need not be considered yet, the μ controller is a simple PI-controller. Both controllers will be discussed further on in this chapter. During the first 1.5 hours of this simulation, ethanol was consumed. For the rest of the simulated experiment, ethanol was neither consumed nor produced. The same tuning for the PI-controller was used. This gives the results of figure 11 and figure 12.

The response using precompensation 1 lies above the set-point for 1.5 hours. During this span of time, the yeast grows both on glucose and on ethanol. The PI-controller is compensating for the growth on ethanol, pulling the specific growth rate down. It does not manage to reach the set-point within 1.5 hours, though. As soon as all ethanol has been consumed completely, μ falls below its set-point. The compensation achieved so far for growth on ethanol must be undone again. This is realized in the next 1.5 hours. At three hours from the simulation start, the specific growth rate is tightly on set-point.

During the first 1.5 hours of the simulation with precompensation 2, the specific growth rate is on set-point, see figure 12. The precompensation detects growth on ethanol and compensates correctly for this. The ethanol consumption rate is limited by the oxygen concentration and therefore independent of the exact ethanol concentration. As soon as the ethanol concentration is low enough to have ethanol limited ethanol consumption, control becomes worse. This is caused by the noisy ethanol estimates. These influence the glucose feed rate in the rest of the simulation. The fluctuations in the ethanol estimates cause fluctuations in the glucose feed rate. The specific growth rate takes over these fluctuations.

To improve the performance of precompensation 2, the sensitivity to errors on the ethanol estimates during ethanol limited ethanol consumption needs to be addressed. The third precompensation, precompensation 3, uses a variant of (21). $\lambda_e(E=\hat{E})$ is used to determine which of λ_p, λ_o or λ_e should be active. In case λ_e is active, $\lambda_e(E=0)$ is then used for the calculation of GF. This will result in a peak on μ if the precompensation switches from oxygen limitation to ethanol limitation. This peak will die out because of the consumption of ethanol, which will eventually result in $E=0$ indeed. The PI-controller further aids in removing this peak. The knowledge used by precompensation 3 is the same as that for precompensation 2. It still needs the ethanol consumption rate as a function of E and G, but only to choose between ethanol limitation and oxygen limitation, not to dose the glucose feed.

The last precompensation, precompensation 4 tries to eliminate the need for accurate ethanol estimates by using predictions of the ethanol concentration instead of observer values under ethanol limitation. (Part of) the simulation model is used to generate these predictions. It is acknowledged, that this might not quite work in practical situations. The degree of correspondence between reality and simulation model is probably not high enough to give predictions that are more reliable than the observed values. It seemed nevertheless worthwhile to investigate the

possibilities of this approach. The knowledge used by precompensation 4 is again the same as that used by precompensation 2. This algorithm relies heavily on the ethanol consumption rate as a function of E and G.
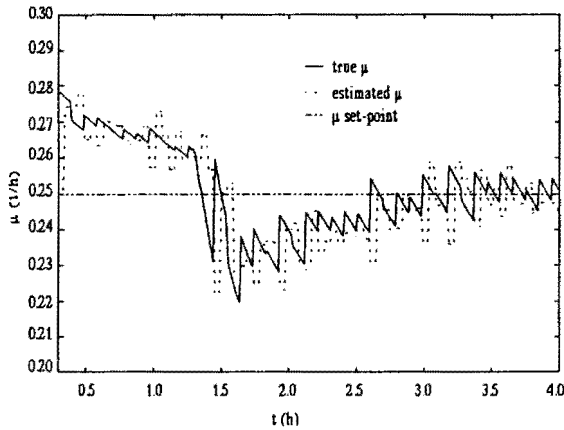

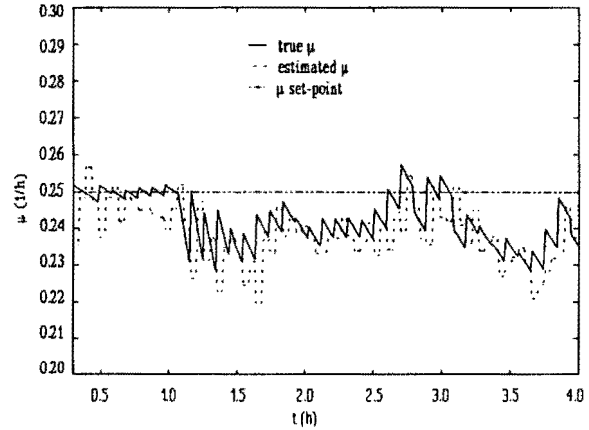figure 11: Precompensation 1, $\mu_{set}=0.25$ 1/h

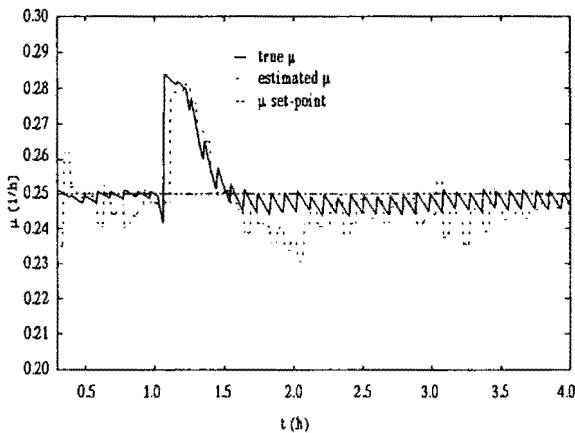figure 12: Precompensation 2, $\mu_{set}=0.25$ 1/h

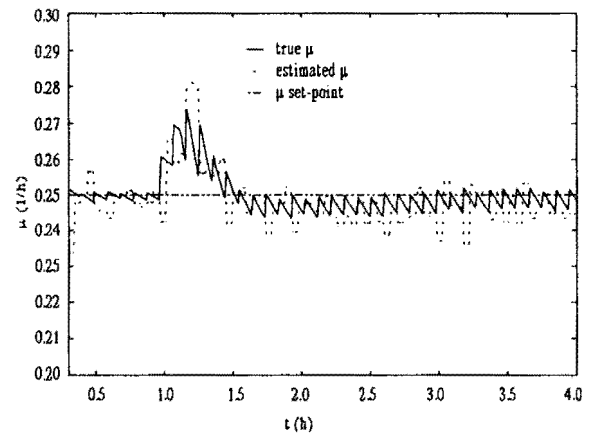figure 13: precompensation 3, $\mu_{set}=0.25$ 1/h

figure 14: precompensation 4, $\mu_{set}=0.25$ 1/h

The peak on the specific growth rate at the switch of oxygen limitation to ethanol limitation, which was expected for precompensation 3, is observed indeed. As soon as all ethanol has been consumed, the specific growth rate is steady on its setpoint. Precompensation 4 achieves little gain in performance. The peak on μ is smaller than for precompensation 3 but still rather obvious. Its dependence on the validity of the simulation model has increased significantly, however.

The μ set-point for the discussed simulations was noticeably below its critical value. Consequently the dip in the specific growth rate after all ethanol has been consumed is considerable: there is much room left in the oxygen bottleneck for growth on ethanol. The presence or absence of ethanol in the broth can make a large difference in the specific growth rate. If the oxygen bottleneck is filled more completely, the switch from oxygen limitation to ethanol limitation gives smoother response on the specific growth rate. This is demonstrated in figure A.8 and figure A.9 in appendix 4. Precompensation 1 and 2 were tested here with a μ set-point of 0.30 h^{-1} instead of 0.25 h^{-1}.

The control performance appears to be limited by the accuracy of the ethanol estimates. Improving

the accuracy of these estimates would improve control. Therefore, some extra knowledge was implemented in the observer. If the μ estimate is well below its critical value, ethanol production is disallowed. If μ is significantly above its critical value, ethanol consumption is prohibited. By keeping the margins used around the critical value large enough, the use of this knowledge may stay within justifiable bounds. Excluding ethanol production for μ below 0.29 1/h leads to a significant improvement of control performance. In the figures below, performance using such an enhanced observer is compared with performance obtained by the observer used so far. It is evident, that control has improved. The sudden fall of μ in figure 15 is caused by an incorrect recognition of the yeast state (ethanol production, oxygen limitation or ethanol limitation). This in turn is caused by the drifting ethanol concentration estimate.
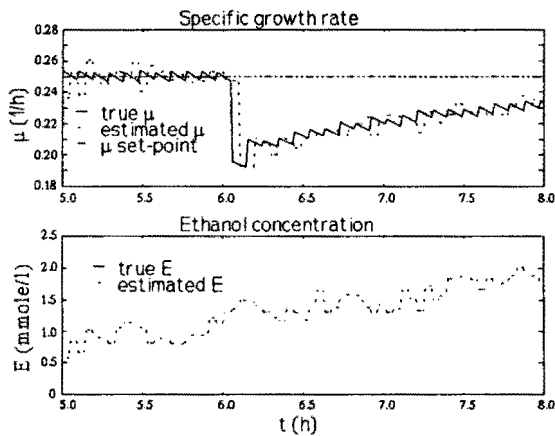

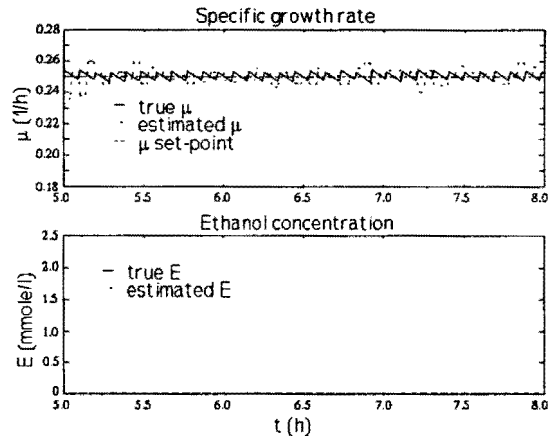figure 15: precompensation 4, original observer   figure 16: precomp. 4, enhanced observer

Another possible solution to the problems regarding the accuracy of the ethanol concentration estimates is using an ethanol measurement. Measurements of the ethanol concentration in the off-gas are available. Assuming certain exchange processes are in steady state, the ethanol concentration in the broth can be derived from these data. The noise level is about 100 ppm for these measurements. This exceeds by far the error in the observer estimates. Therefore the measurements currently available offer no solution to the sensitivity to errors in the ethanol estimates. During the experiments, first tests were carried out at Unilever Research Laboratories with an ethanol sensor. This would provide fairly accurate continuous ethanol measurements. Linking these measurements to the estimates was considered beyond the scope of this control problem.

3.3. μ-controller

3.3.1. Structure

The glucose flow precompensations are essentially feed-forward controllers of the specific growth rate. They have the usual disadvantages associated with feed-forward control:
- discrepancies between the assumptions implemented in the precompensation and the true situation are not corrected or even detected
- external disturbances are not suppressed

For this reason, the compensation should be complemented by some means of feed-back control. A simple PI-controller was chosen for this task. Two conceivable PI-structures are:
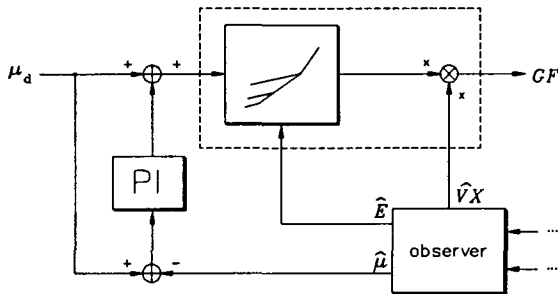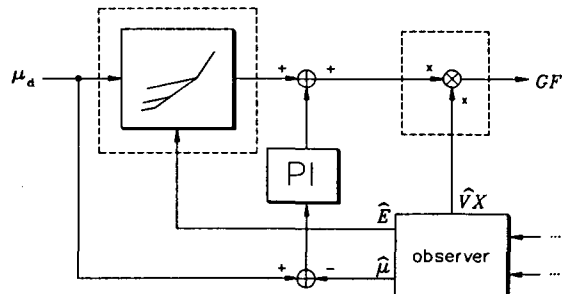

figure 17: first PI variant

figure 18: second PI variant

The parts within dotted lines are the glucose flow precompensation. The essential difference between the first and the second variant is, that the first variant influences the state recognition of the precompensation, the second variant does not. It depends on the kind of error that needs to be corrected which of these two would be appropriate. If the value of the critical growth rate is in error for example, the first structure can correct this. If the maintenance rate is wrong, the second scheme would eliminate this error. A combination of both schemes is also possible: the PI-output could be added at both places with different gains. Making either of the gains equal to zero reduces this structure to one of the variants mentioned above.

In the simulation environment, it can not be predicted, which of both situations will occur in practice. Therefore the first structure is chosen quite arbitrarily. All simulations in this chapter using a μ-PI use the variant of figure 17. As long as the precompensation does not switch between different pathways, both solutions are interchangeable.

The glucose feed rate is obviously limited to positive values. There is also an upper bound on it, caused by the maximum flow of the glucose pump. If the upper bound is reached, the integrator value in the PI-controller is not allowed to increase any further. Conversely, if the lower bound is reached, the integrator value may not decrease. This measure improves the practical robustness of the design, but does not affect normal performance.

3.3.2. Tuning

Tuning of the PI poses a theoretical problem. Apart from the delays in the measuring equipment, the nominal transfer from the input of the precompensation to the μ estimate can be modelled by a zero order model. It is nevertheless conceivable, that using a higher order model can account for the errors made by the observer in the numerical integration, for example.

To investigate this possibility, a prbs simulation experiment was designed. The input signal of the precompensation alternated between an ethanol producing specific growth rate and an ethanol consuming rate. Such rates were chosen, that the ethanol production rate of the first situation equalled as far as possible the ethanol consumption rate of the second situation. This can not be completely controlled, as the ethanol consumption rate is a function of the ethanol concentration, which is not controlled directly in this setup. Nevertheless the ethanol concentration did not drift away during the simulation. First, second and third order ARMA models were fitted using a least squares equation error criterion. Delays were corrected explicitly. This gave the pole-zero plots of

figure 19. The pole-zero cancellation for all model orders is clearly visible, so the nominal transfer faced by the PI-controller is a pure delay indeed.
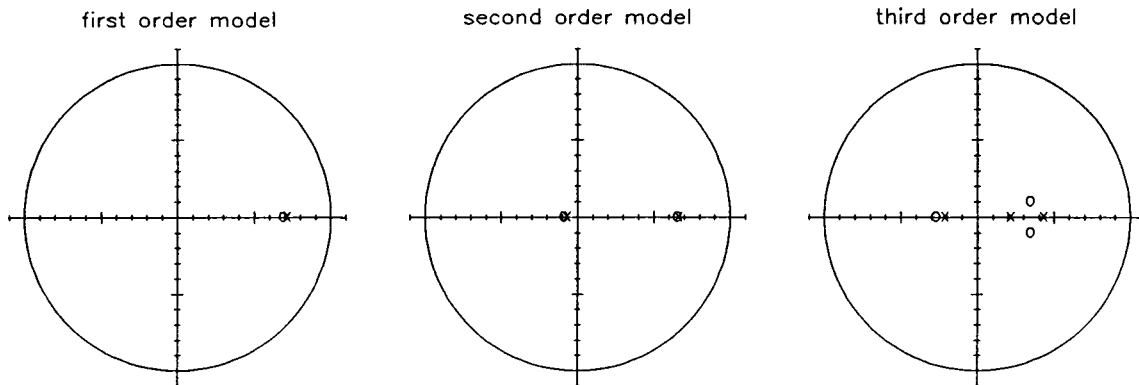
first order model   second order model   third order model

figure 19: pole-zero plots for μ_d-$\hat{\mu}$ transfer

An optimal tuning for this delay could be developed, but in the practical situation, the delay of the measuring equipment will not be constant. It will vary presumably between one and two samples. Tuning the PI based on the noise characteristics offers no solution either. Hardly anything is known about the noise on the measurements, except that the noise is probably a low-pass process and not stationary. A third possibility would be to model the structure of the modelling errors to be expected. This approach is doubtful, because one can not expect the simulation model to give a reliable indication of the structure of the dynamics it does not model correctly, because the structure of the model is based on sheer data fitting. It is not supported by physical or chemical reasoning. Even if this were allowable, it would be far from straight forward to do. This leaves us little choice than to base the tuning on trial and error. These observations imply also, that control design techniques such as pole placement and H_∞ can not be used either.

The tuning parameters found for the simulation model were a proportional gain of 0.1 and an integral gain of 0.03. On the actual yeast process, these parameters can be used as a starting point for tuning of the actual process. Then, another balance has to be found between responsiveness on the one hand and robustness and stability on the other hand.

3.3.3. Simulation results

The simulation results of the previous paragraph showed, that the current structure is able to maintain a constant set-point on the specific growth rate. Further simulations were carried out to assess its performance for a stair-case profile on the set-point and a ramp signal on the set-point. The result were satisfying, cf. figure 20. The set-point changes are rapidly followed. Between 6 and 8 hours, the true μ lies somewhat above the set-point. This is not a controller problem but an observer problem. For this value of μ, very close to its critical value, the observer has difficulties reconstructing the yeast state from the noisy measurements. Because the controller uses the observer values, these problems influence the overall control performance as well. As these effects are stored in the integrator of the PI-controller, this effect is also visible during the next two hours, in which the set-point is located significantly above the critical value. The observer has little trouble recognising the yeast state then. To illustrate the non-linearity of the input required

for this response, the glucose feed rate for this simulation is plotted in figure A.10, appendix 4.
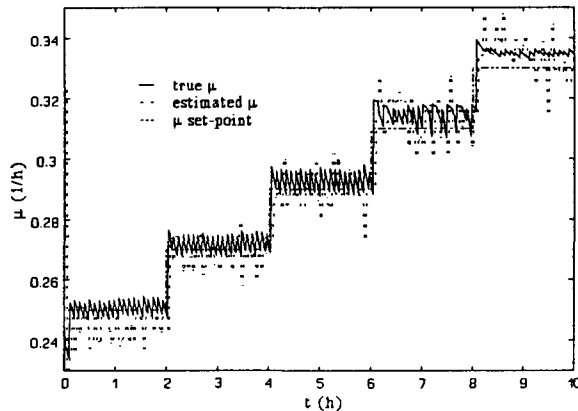

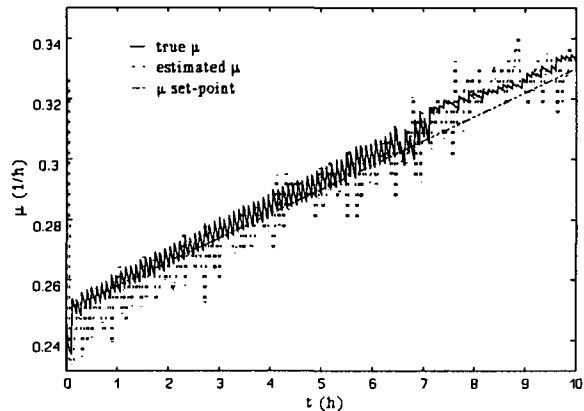figure 20: performance for staircase set-point   figure 21: response to a ramp signal

The response to a ramp signal on the μ set-point corresponds to the response to the staircase signal. For set-points near the critical growth rate the true growth rate is higher than the observed value. This effect is passed on to the two hours following this situation. The current control structure is capable of following a ramp that passes through a wide range of values: it starts off with a value far below the critical value and ends with a value significantly above it. It is not advisable to reproduce this simulation in a practical setting, because the ethanol production rate towards the end of the simulation is considerable. The ethanol concentration reaches values of the order 100 mmole/l (approximately 2000 ppm).

The saw-tooth shaped pattern superimposed on the true specific growth rate in each simulation is a result of yeast growth during a sample interval. It is the same effect as the one that causes the observer to give amounts estimates that are too small. The glucose feed rate is calculated at the start of each sample interval from the amount of yeast that is present at that moment. During the sample interval the glucose feed rate should actually increase, because the amount of yeast increases. This problem could be met by a predictor generating between two sample instants extra estimates of amounts. This is not advisable for the same practical reasons as were applicable to the observer problem. See also 3.5 for a further discussion on predictors for the yeast process.

3.4. E-controller

The ethanol controller is added to the structure mainly as a safeguard against excessive ethanol concentration. High ethanol concentrations obstruct yeast growth and can eventually destroy the yeast cells. It also has a negative influence on the quality of the yeast. If the controller manages to maintain the ethanol concentration at its maximum without significant overshoot to either side, the current control structure can be used for both ethanol control and μ control, as already mentioned at the beginning of this chapter.

3.4.1. Algorithm

The algorithm used by the controller is

$$\mu_d = \mu_{set} - K_p \cdot (\hat{E} - E_{max})_+ - K_i \cdot i(\hat{E}) \tag{24}$$

with

$$x_+ = \begin{cases} x & \text{if } x > 0 \\ 0 & \text{if } x \le 0 \end{cases} \tag{25}$$

and

$$i(\hat{E},k) = \begin{cases} i(k-1) + (\hat{E} - E_{max}) & \text{if } \hat{E} > E_{max} \\ p \cdot i(k-1) & \text{if } \hat{E} \le E_{max} \end{cases} \qquad 0 \le p < 1 \tag{26}$$

If the ethanol limit is exceeded, (24) acts as an ordinary PI-controller, feeding back both the error and the sum of errors. As soon as the ethanol limit violation has been eliminated, the intervention on μ is gradually turned off. The μ set-point converges back to its original value according to a first order step-response with a pole in p.

If a μ set-point change occurs when the ethanol limit is exceeded, the output of the ethanol controller is not allowed to increase. If passing on the set-point change would cause this, the value of i is changed so that the set-point change is absorbed completely. Downward changes of the controller output as a result of a μ set-point change are not filtered out.

The responsiveness of the controller is further improved by proper initialisation of i at the beginning of a limit violation. i is set to a value such, that the output of the controller is immediately at a value near the critical value, unless this would increase μ_d.

Just like there are physical limitations on the range of values GF can assume, μ_d is limited too. It is not allowed to take negative values. If μ_d runs into this limitation, the value of $i(k)$ is frozen, so that the integrator in the PI does not drift away. This is similar to the freezing of the integrator in the μ-PI.

3.4.2. Tuning

The asymmetric nature of the PI-part of the controller makes conventional PI-tuning criteria impractical. If K_p and K_i are chosen too small, it will take the controller too long to reach μ values for which ethanol is consumed rapidly enough. If they are chosen too big, μ values take on too small values, so that the ethanol intervention stays on too long after the violation has been eliminated. The same holds for the value of p. If it is taken too close to one, the decrease of μ_d does not die out rapidly enough, if it is too small, μ_d overshoots the critical value too far, so that instant production of ethanol occurs and a new intervention can be expected within a few sample instants.

Another problem the E controller faces, is that the ethanol production rate is proportional to X. During an experiment X increases by more than a factor 10, making control more difficult as the experiment progresses.

3.4.3. Simulation results

A simulation with a μ set-point of 0.32 1/h and an ethanol limit of 100 ppm is presented in figure 22 and figure 23. No noise was added to the measurements to make the essential issues stand out more clearly.
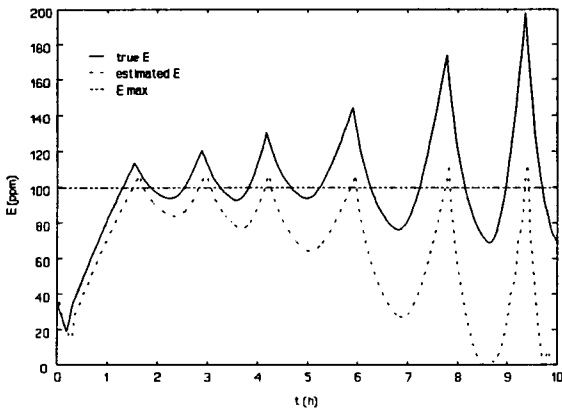

figure 22: E-control, E_{max} = 100 ppm,
μ_{set} = 0.32 h^{-1}
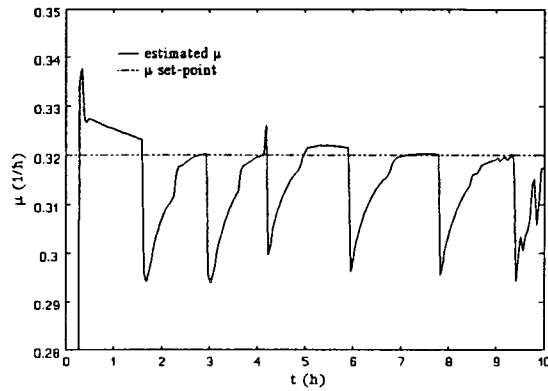
figure 23: E-control, E_{max} = 100 ppm,
μ_{set} = 0.32 h^{-1}

E control results inherently in μ set-points close to the critical value. It is for these set-points, that the observer has the most difficulties in recognising the pathways currently followed. This degrades the ethanol concentration estimates in figure 22.

The increasing overshoot in the ethanol concentration is caused by the increasing biomass concentration. The ethanol production rate is proportional to this concentration. In figure 23 is shown, that the pattern on the specific growth rate remains fairly constant.

Things get even worse if a larger μ set-point is chosen:
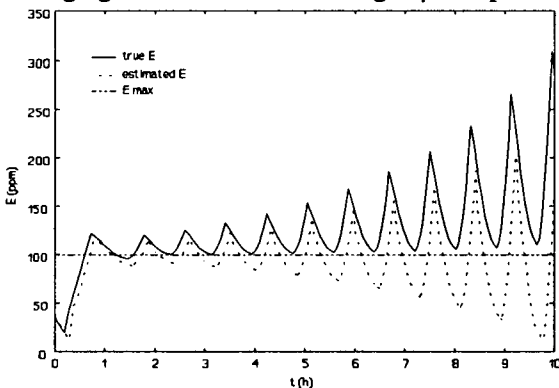
figure 24: E-control, E_{max} = 100 ppm,
μ_{set} = 0.33 h^{-1}
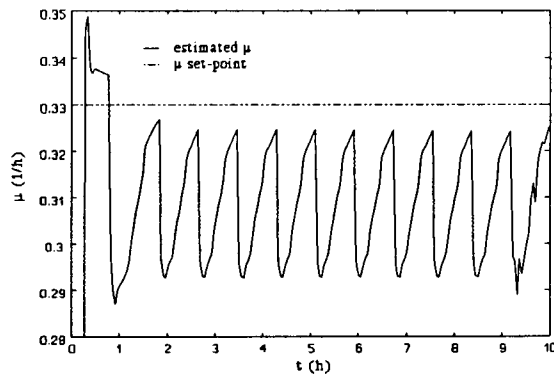
figure 25: E-control, E_{max} = 100ppm,
μ_{set} = 0.33h^{-1}

The larger μ set-point causes the μ-controller input to return faster to super-critical values in absolute sense after an ethanol controller intervention. The ethanol limit is exceeded faster now. The effect in figure 24 clearly has a higher frequency than the one in figure 22. The violation of the ethanol limit is also more severe.

3.4.4. Discussion of results

From these few simulations it is already clear, that the current structure is not suited for strict ethanol control. It could be used as a mere safeguard if the ethanol estimates could be improved, but realising true ethanol control by setting the μ set-point to large values is certainly beyond the capabilities of this controller. The specific growth rate is not an adequate input to control the ethanol concentration for the following reasons:

* The specific growth rate is not the only factor determining the ethanol consumption rate. The ethanol concentration itself is an important factor too, as is the biomass concentration.
* The performance of the μ-controller is not such, that the true specific growth rate is always instantly equal to the set-point. Running the ethanol control at a lower sample frequency than the μ control would certainly meet this problem. This could result in an concentration violation going undetected for longer than necessary. The slopes in figure 24 indicate, that this is not a useful option.
* The critical growth rate is not known exactly and may even vary somewhat over different experiments. This is certainly the most important shortcoming of this setup.
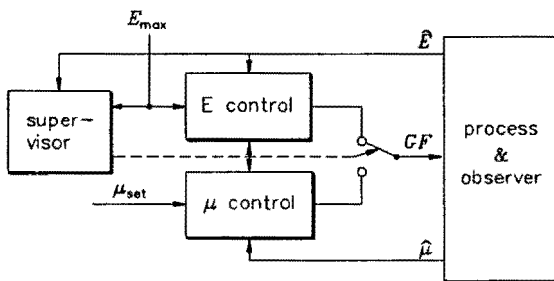

figure 26: Alternative to figure 9

An alternative to the structure of figure 9 could be the one of figure 26. In this setup, the E-controller no longer uses the specific growth rate as a process input, but acts directly upon the glucose feed rate. This approach may look promising at first, but it has some of the disadvantages from which the setup of figure 9 suffered as well. In this case, the critical glucose feed rate is not known. This gives the same problems in tuning the E-controller as were encountered in 3.4.2. Also, the problem of higher ethanol production and consumption rates when the yeast grows occurs exactly the same way in this structure.

Now no distinct solution can be offered for the ethanol control problem. If the controller is seen as just a safeguard, the controller developed here meets its targets. If true ethanol control is desired, one might look into the structure of figure 26. For the block E-control one of the ethanol controllers described in the literature could be used. See for example (Axelsson, 1989).

3.5. Predictive control

3.5.1. Advantages of predictive control

Faster controller response can be achieved by using a higher sample frequency. The true sample frequency can not be changed, unless other measuring equipment is used. It might be possible however, to identify a linear model describing the relation between process inputs and outputs over a time span of a few samples. By running the model at a higher sample frequency than the measuring equipment, the model can be used to generate intermediate samples. This is, in fact, predictive control, as the model predicts these extra samples.

This technique has three distinct advantages:

* The sample frequency is seemingly increased, which allows for a faster controller performance.
* If the interval over which the predictions are reliable can be extended to two sampling intervals of the measuring equipment, we can get rid of the two sample delay in the measurements.
* If the difference between the predicted measurements and the actual measurements exceeds a certain bound, the actual measurements can be rejected in favour of the predictions. In this way, recalibration effects and other outliers in the measurements can be removed to some extent.

3.5.2. Estimation of a linear model

It is assumed, that the oxygen tension is high enough to have no influence on the yeast behaviour: the predictions will be independent of the oxygen concentration. It is further assumed, that no ethanol production occurs. The effects of both ethanol production and consumption can impossibly be included in a single linear model. If the application of the linear model proves successful under this condition, a separate model can be estimated for ethanol production. Then a switching algorithm between the two models has to be devised as well. These assumptions leave us with one process input (*GF*) and two process outputs (*OUR* and *CPR*).

If ever a linear model is to describe this highly non-linear process in a more or less accurate way, even under the above mentioned assumptions, some pre- and postcompensations are necessary. The compensations used are basically the same as those used during the control design: inputs and outputs are scaled to the biomass estimates. Together with some provisions to compensate for the delay in the measurements and for the multi-rate situation we are in, this gives the scheme of figure 27, overleaf. Offset corrections are not drawn in this figure.

The data used for identification were collected in an experiment where *GF/VX* was excited by a PRBS signal. The minimal pulse width of the PRBS signal was 4 minutes, the over sampling rate was 2. The process in- and outputs are given in figure A.12 and figure A.11 in appendix 5. The in- and output signals were scaled to *VX*, detrended and decimated. MatLab's ARX was used to fit an ARMA model to the data. A state-space representation of this discrete time model was converted to continuous time using matrix logarithms. The continuous representation is then ready to be converted back to a discrete model using any sampling time (within reason). The impulse response for a sampling time of 4 minutes, as was the pulse width during the identification experiment, is shown in figure 28. The step response is also shown in this figure.

figure 28: Response of *OUR/VX* (detrended)

Remarkable about the step response is, that the response dies out nearly completely compared to the initial transient. This does not correspond to the physical background: if the glucose feed is increased stepwise, the oxygen uptake rate

Observer

Obs.

predictions

$\begin{cases} \text{OUR}_p \\ \text{CPR}_p \end{cases}$

measurements

$\begin{cases} \text{OUR}_m \\ \text{CPR}_m \end{cases}$

figure 27: Predictor structure

should increase stepwise as well. In steady state, the yeast will grow at a higher rate. This is not covered by the identified model.

This discrepancy between the modelled dynamics and the physical behaviour can probably be explained by the lack of a superposition theorem in the physical reality: a step can have an essentially different effect than a continuous series of pulses would have. Whatever the cause, the model seems to be valid for strictly one working point and not for a range of working points.

The optimal feedback gain $K(z)$ in figure 27 is difficult to calculate. The presence of different sample rates and delays which are not included in the model differs essentially from the usual Kalman gain determination for example. As a temporary solution, it was decided to disregard the delays after the model and to assume that measurements are available at each predictor sampling instant. This will not be the optimal gain in some sense, but it provides a working basis. The feedback gain was forced to zero at sampling instants for which no measurements are available.

3.5.3. Validation of the predictor

The predictor was validated by having it predict intermediate samples for one of the experiments carried out. An experiment was chosen in which various μ set-points were used. Thus the performance of the predictor in different working points could be established.
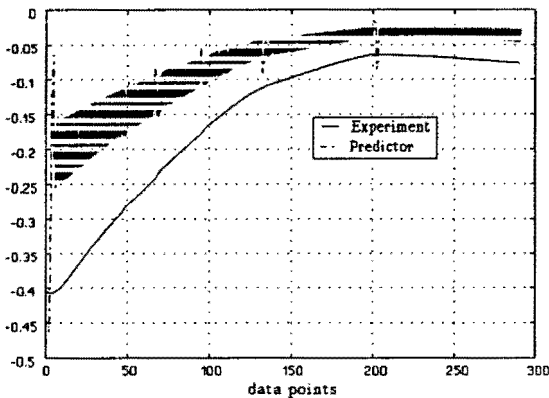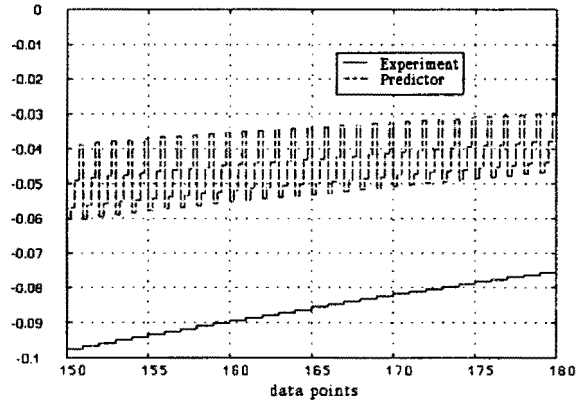
figure 29: *OUR* (experiment & predictions)

figure 30: detail of left figure

The detail in figure 30 gives a clear picture of what is going wrong: as mentioned before, the predictor is actually suited for only one working point. The working points used are apparently not equal to the one used during the identification experiment. The outputs of the predictor model converge to the outputs belonging to another working point. The offset that should remain on the outputs dies out. This happens between measuring samples. As soon as a new measuring sample becomes available, the outputs are pulled back towards the level they should be at in the first place. This gives the pattern as seen in figure 30. The current predictor structure is not suited for its purpose.

3.5.4. Alternative structure

To cope with the mentioned problems, it was decided not to use *GF/VX* as an input and *OUR/VX* as an output, but $\Delta(GF/VX)$ and $\Delta(OUR/VX)$ instead. This modifies the control structure according to figure 31.

This setup eliminates the problems related to offsets, but it also complicates tracking step-changes, as these become mere pulses on the model input. Please note, that the state associated to the integrator Δ^{-1} is also incorporated in the feedback on the states. Even though the differentiator and integrator should theoretically cancel each others effect, a new identification was done for this new structure.
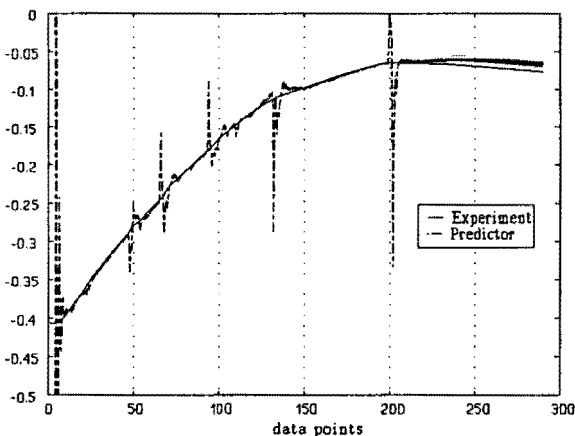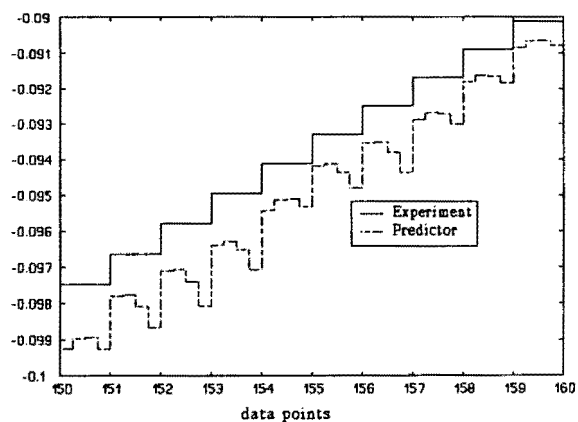

figure 32: *OUR* using alternative predictor
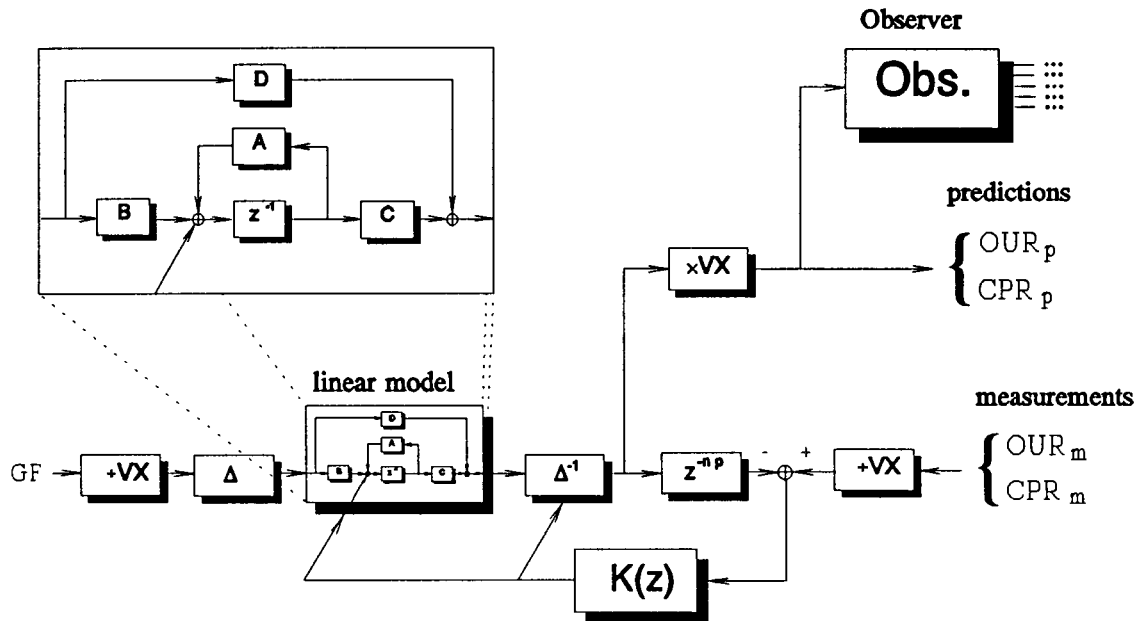

figure 33: detail of left figure

figure 31: Modified predictor structure

In figure 32 and figure 33, the performance of this predictor for the same experimental data as used before is shown. The predictor stays close to the experimental data. There are some peaks and spikes on the predicted values. These can be found to coincide with the step-changes on the μ set-point. After differentiating, these will become impulses. The response of the predictor is therefore very poor. In figure 33 it shows, that the response in between measuring samples is not good either. Again, the predictor drifts away from the "true" values between measurements and is pulled back as soon as a new measurement becomes available. The performance is worse than that of a simple zero order hold.

3.5.5. Conclusions and recommendations

The predictors tested so far were not able to give useful predictions in between measuring instants. Probably the non-linear process behaviour is the cause of the fact that a true step response is not contained in the identification data. This makes the predictor suited for only one working point, the one that was used for the identification experiment. For constant μ set-points, however, the control structure performs satisfactorily as it is, it is for set-point changes that improvement is desirable.

Identification experiments designed specifically for the identification of step responses could solve this problem. Even then it will not be guaranteed though, that the predictor is useable for a sufficiently wide range of set-points.

3.6. Stirrer speed precompensation

The principle of the stirrer speed precompensation is the same as that of the glucose flow precompensation. The yeast is considered to be in one of a limited number of states. The offset correction of the precompensation is changed according to the state that is detected. Based on the simulation model, only two states are distinguished. These are maximum oxygen consumption and sub-maximum oxygen consumption. Maximum oxygen consumption occurs during ethanol production and oxygen limited ethanol consumption. Sub-maximum oxygen consumption takes place during ethanol limited ethanol consumption.

3.6.1. Algorithm

Under maximum oxygen consumption the oxygen consumption rate in mole O_2/h is equal to

$$-\frac{dP_{1a}}{dt} = \left[0.2 \cdot \frac{P}{K_o + P} + ms \right] \cdot X \qquad (27)$$

Referring to figure 10, apart from the X-dependence this is constant.

Under sub-maximum oxygen consumption, the oxygen consumption rate is equal to

$$-\frac{dP_{1b}}{dt} = \left[0.5 \cdot a \cdot \frac{G}{K_s + G} + 0.13 \cdot z_1 \cdot z_2 + ms \right] \cdot X \qquad (28)$$

The glucose consumption rate is taken to be in steady state. As the glucose concentration in the feed exceeds by far the concentration in the broth, in good approximation it holds, that

$$\frac{GF}{VX} \cdot G_{in} = 0.5 \cdot \frac{G}{K_s + G} \qquad (29)$$

z_1 and z_2 should be reconstructed like it was done for the glucose flow precompensations. Contrary to (27), (28) is far from constant. It depends, amongst others, on the current glucose feed rate and the ethanol consumption rate.

The oxygen concentration that is taken up by the broth per unit time, is given by

$$\frac{dP_2}{dt} = C_1 \cdot SS^{C_2} \cdot AF^{C_3} \cdot \left(O_{max} - P \right) \qquad (30)$$

This should compensate (27) or (28), whichever is appropriate, and also drive the oxygen concentration in the desired direction. If the desired net time derivative of the oxygen concentration is denoted dP_s/dt, the overall precompensation algorithm is

$$SS = \left[\frac{\min\left\{ \dfrac{dP_{1a}}{dt}, \dfrac{dP_{1b}}{dt} \right\} + \dfrac{dP_s}{dt}}{C_1 \cdot AF^{C_3} \cdot (O_{max} - P)} \right]^{1/C_2} \qquad (31)$$

The state detection algorithm simply chooses the smallest from (27) and (28). The dilution effect has been neglected.

The number of assumptions and the required a priori knowledge is considerable:
- The glucose concentration is in steady state.
- Current glucose feed rate
- Current amount of biomass
- Ethanol consumption rate as a function of E and G
- Size of the oxygen bottleneck

Further, measurements or estimates are required for
- yeast concentration
- ethanol concentration
- oxygen concentration

These are available, either from the observer or from direct measurements.

The oxygen control runs at twice the sampling rate of the μ-control and the observer. At sample instants for which no new observer estimate is available, a simple zero order hold on the estimates was tested. Also linear interpolation of the estimates was tried.

3.6.2. Simulation results

Simulations point out, that, again, the ethanol concentration is the bottleneck in the control performance. This is needed to reconstruct z_j. Either the ethanol estimate from the observer is used, or the ethanol concentration predicted by precompensation 4. Together with the choice between interpolating linearly between two samples or using no interpolation at all, this gives four options for the ethanol concentration estimates / predictions.

Any benefit to be expected from linear interpolation was more than compensated by the phase shift introduced by this technique. As the ethanol predictions were, (on the simulation model!) more accurate than the ethanol estimates, best results were obtained when using these predictions without interpolation. They will be presented in the next section.

3.7. Oxygen controller

3.7.1. Algorithm

To compensate for errors in the stirrer speed precompensation and to suppress external disturbances, a PI-controller is put in series with the precompensation. This PI-controller compares the measured oxygen concentration with the set-point. As the precompensation requires the time-derivative of the oxygen concentration instead of the oxygen concentration itself, the PI-controller is followed by a numerical differentiator using Euler backward:

$$y(k) = \frac{u(k) - u(k-1)}{T_s} \tag{32}$$

This is open for improvement.

If the oxygen concentration is above its set-point, more energy is spent on stirring the mixture and blowing air through it than is needed. If it is below its set-point, ethanol production is likely. This

means glucose is used less efficiently, as more glucose is needed to grow yeast on ethanol formed during earlier fermentative glucose consumption than to grow yeast on glucose directly. It also disturbs the specific growth rate controller, as this assumes, that the oxygen concentration does not influence the specific growth rate. This is only the case if oxygen concentrations are not too low. The PI controller takes these observations into account by an asymmetric weighing of deviations from the set-point: deviations towards zero are weighed more heavily than deviations away from zero.

3.7.2. Tuning

The tuning of the PI control is based on simulation results. As it was the case for the μ- and E-controllers, insufficient data are available to tune the PI theoretically. Further manual tuning of the controller on the real process is unavoidable.

3.7.3. Simulations

A typical controller response was obtained in a simulation were all paths were followed. The achieved oxygen concentration and the associated (true) specific growth rate are plotted in figure 34 and figure 35.
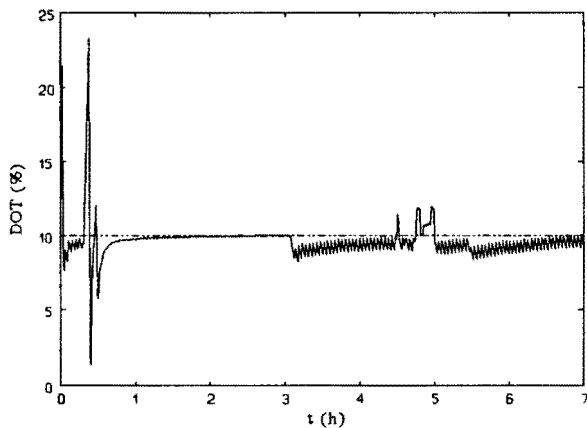
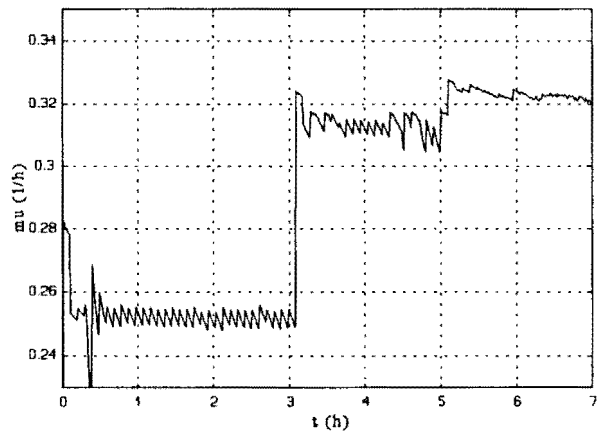

figure 34: Oxygen controller performance   figure 35: μ during P-controller test

The peak on the oxygen concentration after 0.4 hours is a result from the transition from oxygen limited ethanol consumption to ethanol limited ethanol consumption. This is only detected by the controller after it occurs. The time constants associated to the oxygen concentration are so small, that if such a switch occurs during a sample interval, it will be noticeable at the end of the interval. The mentioned peak is over-compensated, resulting in a dip during the next sample instant. This oscillatory behaviour dies out within a few samples. Then the oxygen concentration is restored to its set-point.

After about three hours μ is set to a value for which limited ethanol production occurs. Apparently the need for oxygen is more dependent on the amount of biomass in this case as it was during ethanol consumption, as the saw-tooth pattern observed on the specific growth rate in figure 35 now appears on the oxygen concentration as well. Apart from this effect, the oxygen controller has no difficulties with ethanol production. This also applies to the situation in which

significant ethanol production occurs. This is demonstrated in the simulation between five and seven hours since the simulation start.

4. Implementation

4.1. Fermentation setup

The experiments are carried out on laboratory scale. The fermenter is designed specifically to provide a controlled environment. A schematic view of a fermenter is given in figure 36. Heating and cooling elements are stuck through the bottom plate. During yeast growth, which is an exothermic reaction, only cooling elements are needed. Feeding substances for the fed-batch phase and air enter via a metal pipe. The stirrer is located in the centre of the bottom plate. It is driven by an external motor over a magnetic coupling. A temperature sensor and a drain complete the bottom plate. pH and oxygen probes enter via the top plate. Base, acid and anti-foam solution can also be added to the mixture from the top plate. To avoid infections, all media added to the yeast are sterilized. Any air entering the fermenter is filtered.
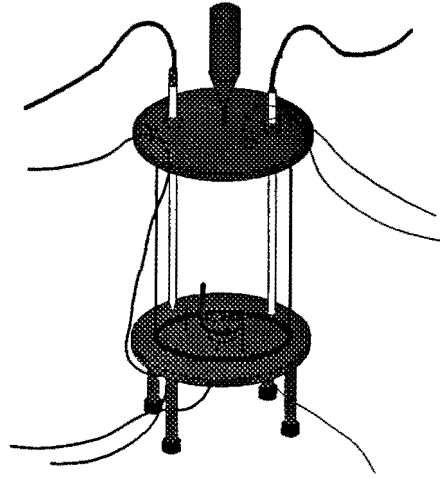

figure 36: Fermenter

pH and temperature are controlled by local controllers and are assumed to be constant at values of 5.0 and 30°C, respectively. The total volume of the fermenter is about 13 litres. The volume of the broth is 4.4 litres at the start of the fed-batch phase and can grow to 6 to 7 litres. The yeast involved is baker's yeast (*Saccharomyces cerevisiae*) and feeding takes place using a synthetic medium. In industrial practice, other yeast strains are used as well and the feeding substance are molasses, which contain, apart from glucose, a number of other substances, which have various side-effects on the fermentation process. Also, the scale of the fermenter is a lot bigger. In the end, control of this process is the objective. For the project described in this report, these aspects should be born in mind but need not be incorporated in the implementation, nor will any experiments be carried out involving these points.

4.2. Overall implementation structure

The control structure was implemented using a combination of existing hard- and software, and software written especially for this purpose. The overall structure is as follows:
The oxygen controller and the stirrer speed precompensation of chapter 3 were not used. A local controller was used instead. This controller runs at a significantly higher sample rate than the μ-controller. It is a PID-controller with dead-zone. Temperature and pH are controlled by on/off controllers. All local controllers are implemented in the Applicon units. They are governed by CDAS. This program runs on a microVAX. Data are read from and written to an applicon unit, which contains the necessary D/A and A/D converters for reading out the sensors and controlling the actuators. CDAS has several other functions as well. It is a data acquisition system too. Measurements read from the Applicon unit are logged for later inspection but can also be viewed on-line. Analysis of the off-gas is performed by a mass-spectrometer. The mass-spectrometer and CDAS offer these services to twelve fermenters at the same time. If all twelve fermenters require

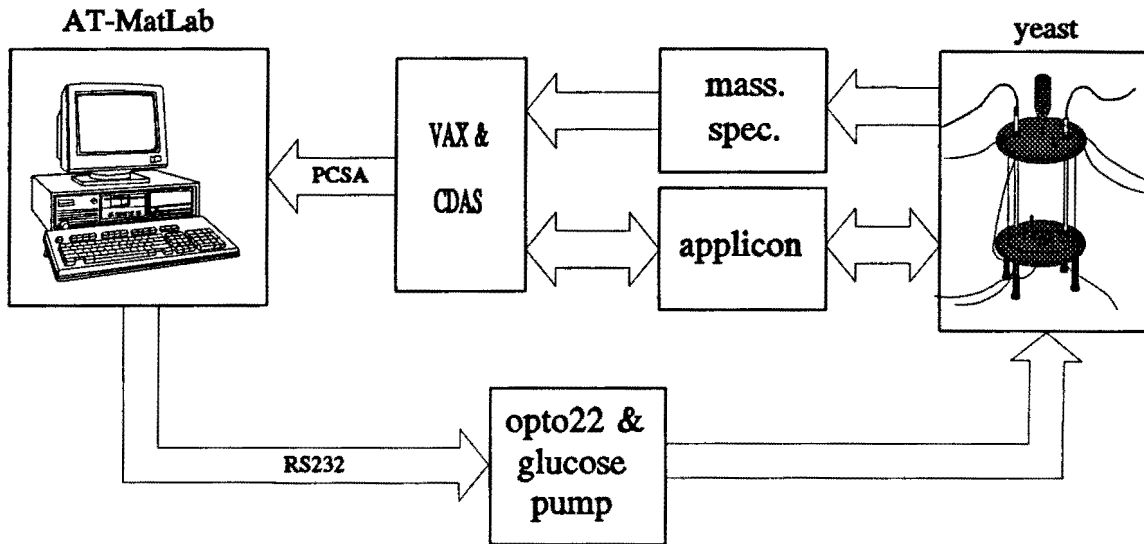AT-MatLab   yeast

figure 37: Overall structure of the implementation

the mass-spectrometer's services, its maximum sampling rate is lowered noticeably. For our experiments, the measurements are written to a shared file. This file is read into a PC via the PCSA network. Any data available to CDAS can be passed on to the PC in this way. The actual controllers are implemented on the PC.

As the stirrer speed and air flow are controlled locally, only the glucose feed rate is to be controlled by the PC. This is done using an opto22 board. This board contains several A/D and D/A converters. It communicates with the computer using the serial port. Only one D/A converter is used. A simple parallel resistor was used to match its output to the analogue input of the glucose pump.

4.3. DOS and MatLab considerations

The μ- and E-controllers are programmed in a combination of the MatLab programming language and ordinary C. The C files are linked to MatLab using Mex-files. All I/O routines are written in C. To avoid the problems associated with protected mode DOS-extenders, PC-MatLab was used at first. This soon ran out of memory, despite combining several Mex-files to one Mex-file, avoiding some of the memory overhead associated to Mex-files. Other memory saving measures proved insufficient as well. Therefore the switch was made to AT-MatLab. After making small changes to the opto22 drivers and resolving a Tandon-DOS compatibility problem, all routines functioned properly.

To give an impression of the structure of the controller program, its flow chart is drawn in figure 38. For the interested reader, the complete source text is given in a separate report. It will suffice to give here some general remarks on the implementation.

As MS-DOS was never intended to run real-time applications, it is entirely the responsibility of the controller program that things happen at the right time. This is achieved by constantly moni-
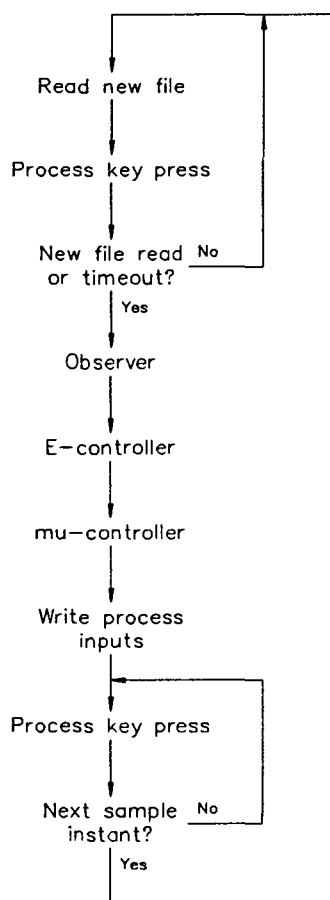
figure 38: flow chart of control program

toring the system clock. In the flow chart this is shown in the lower part of the main loop.

Ensuring the integrity of the file containing the mass spectrometer measurements appears to conflict with sharing it with CDAS or, more precisely, with the behaviour of MS-DOS if a sharing violation occurs. By opening the file both in CDAS and in the control program with the right sharing options, file integrity can easily be ensured. If DOS tries to open the file while this is not allowed by CDAS, a critical error is generated. DOS reports the problem and gives the well-known 'Abort, Retry, Fail' message. Apart from corrupting the screen display, this message will hold the program until the operator presses 'F' or 'R'. Pressing 'A' immediately terminates AT-MatLab, so that the control program is stopped as well. Even opening the file at both sides without denying any right to other processes caused this problem in some cases. This problem was solved by calling AT-MatLab from another program that intercepts the critical errors. If a critical error is caused by a sharing violation, this program takes care of choosing the 'Fail' option without any error message or required operator action.

Using AT-MatLab with Mex-files handling screen output and keyboard input is a combination that is not recommended as an ideal solution to any programming problem. Calling low-level BIOS routines requires frequently shifting the 80286 processor in and out of protected mode. This slows I/O down significantly.

Some data is stored in C static variables. These variables are cleared whenever a DOS program is called from MatLab. It is therefore not allowed to use any !-command from within the program.

4.4. Other experiences

Although it may sound obvious, it is worth mentioning that the glucose pump should be powerful enough. During one of our first experiments, the pump turned out to be unable to rotate at the low rate specified. In the points where the friction with the tube was at its maximum, the pump just stopped rotating.

A remark related to this is, that controlling the pump by a feed-back controller should seriously be considered. Not only would this have prevented the error situation mentioned above, it can also trigger an alarm when the glucose tube has broken. From our experience this is a real possibility. It also eliminates the need for periodic recalibration of the pump. A feed-back signal is already available for this purpose, as the weight of glucose remaining in the supply is measured by a balance.

Two sizes are available for the cooling elements. The shorter ones are used most commonly. For high specific growth rates and high stirrer speeds, the cooling capacity of these elements is

insufficient to keep the temperature at the set-point of 30 °C. This can be solved by placing the heating element in series with the cooling element, yielding two cooling elements, so twice the cooling capacity. A more elegant solution is to use a long element in the first place. Related to this problem is the following: the local oxygen controller keeps the dissolved oxygen tension on set-point by manipulating the stirrer speed. Only if the stirrer speed runs into its lower or upper bound, the air flow is changed. Because this approach can result in unnecessarily high stirrer speeds, it is sometimes advisable to raise the air flow manually. This gives lower energy costs and requires less cooling capacity.

The end of the batch phase was detected manually. A sharp increase in the oxygen concentration signals that all ethanol has been consumed. This could be implemented in the controllers, so that the switch is made automatically. As several operator actions had to be taken at the beginning of the fed-batch phase, this was not considered worthwhile yet.

5. Experimental results

5.1. Introduction

In chapter 3 a control structure was designed and tuned based on simulations. Control goals were first to maintain a constant set-point on μ and then to track changing set-points. Once these goals were realised, provisions were added to obey a limit on the ethanol concentration and to keep the dissolved oxygen tension on set-point. As pointed out in the previous chapter, the control of the oxygen concentration was realised during the experiments by existing controllers. These will not be tested explicitly in the following experiments, as their performance has already been established in other experiments. Apart from the omission of oxygen control, it was intended to follow a path in the real experiments similar to the one followed in the simulations: first some experiments without ethanol production and consumption with constant μ set-point were planned. Then the performance for changing set-points and for ethanol production and consumption was to be tested. Finally, the ethanol control would be tested.

Unfortunately, too many unexpected phenomena were encountered in the experiments without ethanol. It was not possible to do experiments for ethanol production or consumption. In the experiments that *were* carried out, several discrepancies between the simulated yeast behaviour and the actual behaviour were identified. They will be discussed in this chapter. Some required only retuning of a few control parameters, others called for more structural changes.

5.2. Batch results

The batch phase as such is not part of the control problem. All batch results are similar. The results below are taken from one of the first experiments.
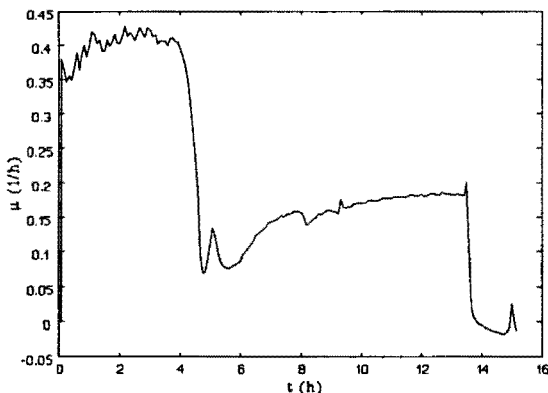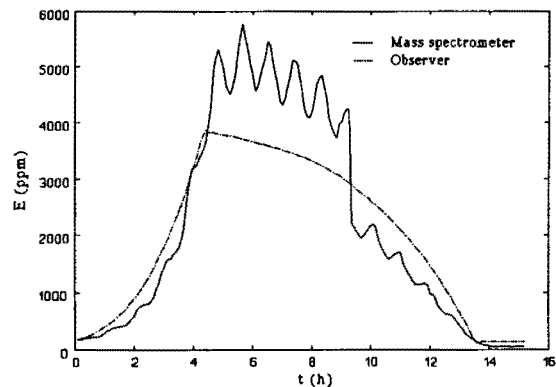

figure 39: μ during batch phase

figure 40: E during batch phase

During the first four hours of the experiment there is excess glucose in the broth. Heavy ethanol production occurs. When the glucose has been consumed, the ethanol produced earlier is consumed again. If this ethanol has been consumed as well, which occurs after about 14 hours from the experiment start, growth falls to near zero values. Due to the maintenance reaction, growth is even slightly negative. The mass spectrometer recalibrates every hour and every 12 hours a more thorough recalibration is performed. The latter caused the fall in the ethanol concentration after approximately 9 hours.

The batch phase ends in a well defined state. This means, that the fed-batch phase is entered in a known and reproducible way. The initial conditions for the fed-batch phase are summarised in the following table:

table 1: Fed-batch initial conditions

Quantity	value	unit
Biomass concentration	10.5	g/l
Volume	4.45	l
Ethanol concentration	0	ppm
Glucose concentration	0	g/l
Dissolved oxygen tension	>90	%
pH	5.0	-
Temperature	30	°C

The rest of this chapter is concerned with the fed-batch phase only.

5.3. Critical growth rate

The critical growth rate during the simulations was approximately 0.32 h^{-1}. It seemed safe to choose a μ set-point of 0.25 h^{-1} when trying to avoid ethanol production. Experimental results proved us wrong. Referring to figure 41 and figure 42 the critical specific growth rate was identified to be approximately 0.23 h^{-1}. This error in the critical μ negatively influences both the precompensation performance and the observer performance. The precompensation assumes an oxygen bottleneck that is considerably wider than corresponds to the real situation. Therefore the precompensation reacts more aggressively to ethanol fluctuations than necessary. The observer inhibited ethanol production for growth rates below 0.28 h^{-1}. The ethanol estimates consequently do not track the ethanol production that occurs during the first part of the experiment.
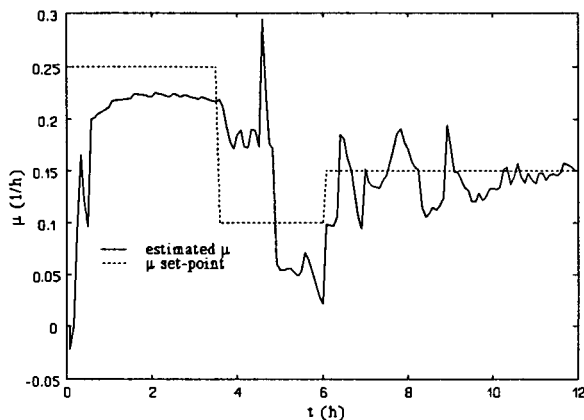

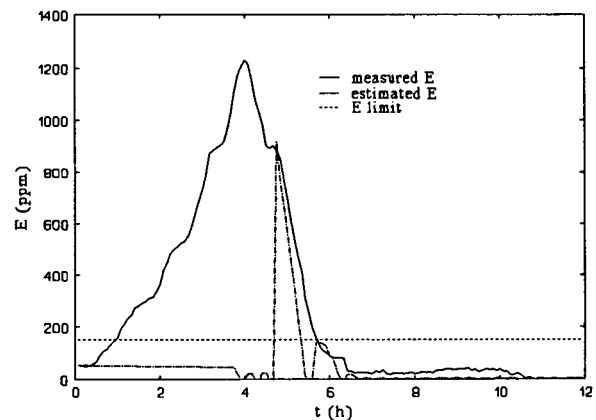
figure 41: μ during first experiment

figure 42: E during first experiment

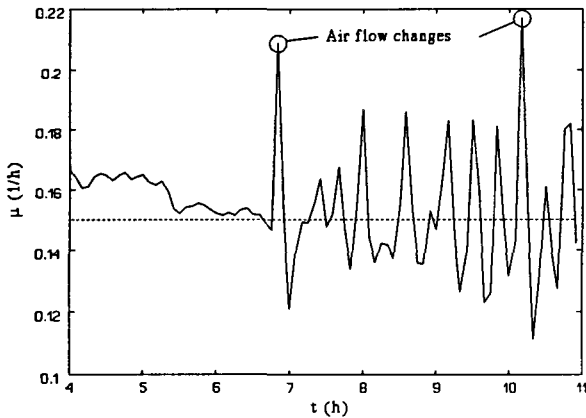## 5.4. Air flow and stirrer speed influence

figure 43: Air flow influence on μ estimates

For high air flow rates, the noise level on *OUR* and *CPR* and consequently on μ are higher. The changes in the air flow rate itself cause a large peak on the μ estimate. Both air flow changes in figure 43 were air flow *raises*. All other operating parameters were left constant. Nevertheless the system starts oscillating after the first air flow raise. The behaviour of the μ estimates coincides with the stirrer speed behaviour. It is therefore assumed, that air flow and stirrer speed changes disturb the *OUR* and *CPR* measurements.

After seven hours, the specific growth rate starts oscillating in figure 43. The only parameter changed was the air flow. It is therefore assumed, that the specific growth rate is not really oscillating, but that it is a problem with the measuring equipment. The higher air flow apparently amplifies the disturbances introduced by the stirrer speed influence. It is not likely that these disturbances represent actual yeast behaviour. To verify this hypothesis, the oxygen controller was switched off during the next experiments. Stirrer speed and air flow were set manually to values high enough to keep the dissolved oxygen tension above 20%. This confirmed the hypothesis, as no sudden increase in the noise level of μ was observed after raising the air flow.

Slowly changing the air flow and stirrer speed reduced the effect on the μ estimates to acceptable levels. Some rate limiting should therefore be implemented in the oxygen controller. Because the oxygen controller was implemented locally, this was not possible for us. The oxygen control was therefore kept in manual mode for the rest of the experiments. In the future, such a rate limiting should be implemented, though.

5.5. PI tuning

The proportional and integral feed-back gain of the PI controller in the μ-controller were initially tuned based on simulations. Some early experiments were dedicated to tuning these parameters based on experiments. If feed-back gains are chosen too low, responsiveness is needlessly sacrificed. It they are chosen too high, disturbances will cause badly damped oscillations or even undamped oscillations and instabilities. An example of the latter effect is shown in figure A.13, appendix 5. The oscillations have a period of four samples. They are mainly caused by the two sample delay of the measuring equipment.

In figure A.14 the final PI tuning performance is shown. The dips in μ marked 'ethanol influence' are caused by faulty ethanol concentration estimates. This problem will be addressed in the next paragraph. Nevertheless, this figure demonstrates, that the PI is now tuned such, that disturbances on μ cause damped oscillations. Later in the experiment control becomes worse. This can be attributed to the stirrer speed and air flow interference. The oxygen controller was not switched to manual mode in this experiment yet.

To aid the PI-controller in making set-point changes, a two sample delay will filter the μ set-point. When the set-point remains constant, the delay has no effect. This implies for instance, that the oscillations mentioned above are not removed by this measure.

5.6. Precompensation structure

No ethanol production occurs in the current series of experiments. The true ethanol concentration will be zero. If the precompensation tries to compensate for ethanol consumption because it sees a non-zero ethanol concentration, extra noise is introduced in the glucose feed rate. For an example of this, see the dips marked 'ethanol influence' in figure A.14. It was therefore decided to use precompensation 3, which does not compensate for growth on ethanol in the glucose feed rate, unless there is excess amount of ethanol. It was also decided to have the PI act upon the precompensation according to figure 17 instead of figure 18. This avoids that the precompensation assumes ethanol production if values for μ become too high.

5.7. Precompensation tuning

During the experiments it was observed, that set-point changes were always accompanied by a considerable overshoot. To find out whether this was caused by the yeast or by the control, the glucose feed rate was studied. To remove the effects related to growth, it was normalised to the estimated amount of yeast first. This showed, that the amount of glucose needed per mole biomass seemingly decreased during an experiment. This was observed in several experiments. The most pronounced example is shown in figure 44.
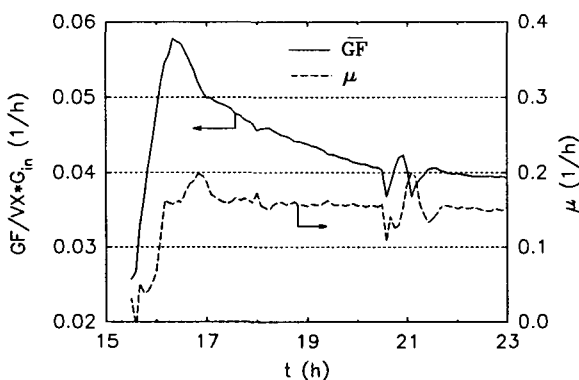

figure 44: Apparent increase in yield

Several factors can theoretically cause this effect. If the initial amount of biomass is estimated too low, more glucose is *seemingly* needed per mole yeast to achieve a certain growth. As yeast grows, the error on the initial estimate becomes of less relative importance. The normalised glucose feed rate should then converge to a constant value. Attempts to give μ and GF a similar pattern show, that only unrealistically large errors on VX_0 could give the results of figure 44. Off-line measurements of VX_0 using dry weight data contradict this possibility.

An error in the specific maintenance rate can also deform the graph of $\overline{GF}$ to the pattern shown in figure 44. The maintenance rate should have a value over three time larger than its current value to match μ and $\overline{GF}$. This results in significantly lower estimates of the amount of yeast, as more yeast is consumed. This is contradicted by the dry weight data as well, which indicates that the observer gives reliable estimates of the amount of biomass.

Another possibility is a non-linearity in the pump. If $\overline{GF}$ remains constant over a period of time, GF actually increases because of growth. If the yield of the pump for large values of the glucose

flow is higher than expected, this explains the observed anomaly. There should be a severe non-linearity, however, to give the result shown. From t=17h to t=21h, the amount of yeast has not even doubled. Calibration experiments for the pump operated the pump from 10% up to 90% of its full capacity. Their results do not support the assumed non-linearity.

Apart from measuring errors a further explanation is, that the yeast adapts to the requested growth rate. Finding a physical mechanism for such an adaptation is a challenge for biochemists and biotechnologists. The adaptation is certainly not covered by (1), (2), (3), (4). From a control point of view, the time constant of the adaptation is large enough to be absorbed by the PI-controller. The current control structure can keep μ constant, albeit a little above its setpoint. As long as the observed phenomenon is not modelled adequately, it seems best to accept this solution. Before the effect can be incorporated in the control, it should be clear, whether the effect is retriggered after a new set-point change, whether time-constants are the same for different specific growth rates and so on.

Now it is clear, that one of the assumptions for all precompensations does not hold. The yields of all pathways are not known constants. It was decided to leave any dynamical effects out of the precompensation and tune it based on experimental steady state values. Steady states can only be expected if neither ethanol production nor consumption occurs. Therefore the experimental tuning can only cover oxidative growth on glucose and maintenance. Several experiments were carried out. The μ set-point was set to various values and was kept constant until a steady state seemed to be reached. A straight line was fitted to the experimental data using a least squares criterion. The slope of this line was considered most important, as any errors in the offset will be corrected by the integral part of the PI-controller. This will then influence control performance no further. The tuning found was

$$\frac{GF}{VX} \cdot G_{in} = 0.2005 \cdot \mu_d + 0.0093 \tag{33}$$

In figure A.15, appendix 5, the experimental data and the resulting tuning are plotted. In figure A.16 the tuning is validated against the results of another experiment. The offset is clearly not optimal for this experiment. As the error in the offset will be corrected by the PI-controller, a precompensation transfer with a corrected offset is plotted as well. The gain of (33) then clearly corresponds to these experimental data.

According to (1), (2), (3), (4), the value of 0.2005 in (33) is equal to $1/b$, so that b should be equal to 4.99. The value used by the observer is 3.65. These values differ 37%. Observer performance seems too accurate to support the value of 4.99. A possible explanation is, that the actual yeast behaviour is not in the model set (1), (2), (3), (4). This means no physical interpretation for the value of 0.2005 is available. Before bothering biochemists for providing a realistic interpretation of (33), measuring errors should be ruled out. Several factors can cause such errors.

- the yield of the pump may be higher than assumed. This includes the transfer from input current of the pump to yield of the pump and the conversion from opto22 driver input to D/A converter input.
- the true concentration of glucose in the fed-batch feed may be higher than expected.
- the observed value for VX is possibly too large.

The decrease in weight of the glucose supply is logged. The logged data seem to confirm, that the yield of the pump is too high. The total volume of the broth at the end of an experiment is lower

than the integration of the assumed glucose feed rate predicts. This can be attributed to either a lower yield of the pump than expected or evaporation effects. The latter are quite likely, as blowing air through the broth enhances evaporation. The hypothesis of higher pump yield is contradicted only by the calibration experiments. These take half an hour to an hour, depending on the number of data points used for calibration and the time taken per data point. The pump is driven in a feed-forward way only. Errors occurring if the pump is used over a longer period of time are therefore not detected. For this reason, the calibration experiments may not be totally reliable. It is nevertheless felt, that further backup for this hypothesis is needed before accepting it. This could be done by means of a calibration experiment covering several hours.

It is highly unlikely, that the concentration of glucose in the fed-batch feed is higher than expected. The error in the amount of glucose in the fed-batch feed is less than 1 ‰. The confidence in the observed values is also such, that an error of 37% in VX seems impossible.

5.8. Filtering set-point changes

From other experiments at Unilever Research Laboratories it is known, that the yeast produces ethanol if the glucose concentration increa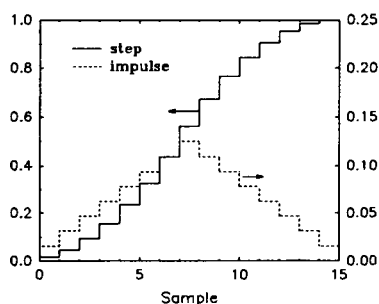ses suddenly. As a rule of thumb, an increase of 0.1 g/l in the glucose concentration is sufficient to trigger this effect. This may obstruct set-point changes. In order to avoid this effect, μ set-point changes are filtered by a finite impulse response filter. The step response of this filter should start and finish smoothly. This suggests the triangular impulse response shown in figure 45. A typical value for the length of the impulse response is taken in this figure. At every set-point change it is verified whether the current length is sufficient to keep the maximum change in glucose feed rate below a critical value. This critical value is 0.1 g/l glucose divided by the sampling time, yielding an upper bound on the change in the glucose feed rate. As the transfer from $\Delta\mu_{set}$ to ΔGF is proportional to the amount of yeast, the length of the impulse response will increase towards the end of the experiment. Large set-point changes will also induce an increase in the impulse response length. The dc-gain of the filter is always equal to 1: set-point changes should be smoothened, not changed in magnitude.

figure 45: fir-filter response

Unfortunately, this approach proved unsuccessful in removing the overshoot from set-point changes. The overshoot remained even roughly the same. The only effect of this filter turned out to be a pointless delay on the set-point change.

5.9. Dynamic precompensation

As overshoot remained on the response to step-changes after retuning the precompensation and rate limiting μ_{set}, the suspicion arose that the overshoot was caused by unmodelled yeast dynamics. To investigate this, set-point changes were requested more or less in open loop: the PI-controller was switched off during these trials. As GF still depends on VX the process is not completely operated in open loop. μ_{set} was changed from 0.048 h^{-1} to 0.055 h^{-1}. Larger changes could not be investigated, because they caused the specific growth rate to drift away to ethanol producing values. The suspicion was confirmed: the change was accompanied with overshoot.

The dynamic behaviour of the yeast will be accounted for in the control structure by means of a block "dynamic precompensation." Together with the FIR filter of the previous paragraph, the final structure will be as depicted in figure 50, page 54. The dynamic precompensation should compensate the yeast dynamics just discovered. It operates on a feed-forward basis. Therefore a model was fitted to the experimental data using a least squares output error criterion. A second order model proved sufficient to describe the overshoot, see figure 46. The model poles

figure 46: Step response

and zeros are given in table 2. Removing the zero in -57.7310 deteriorated the fit, so it was decided to leave it in the model. The transfer function of the dynamic precompensation is the inverse of the model. The unstable pole in -57.7310 is relocated to 1/-57.7310.
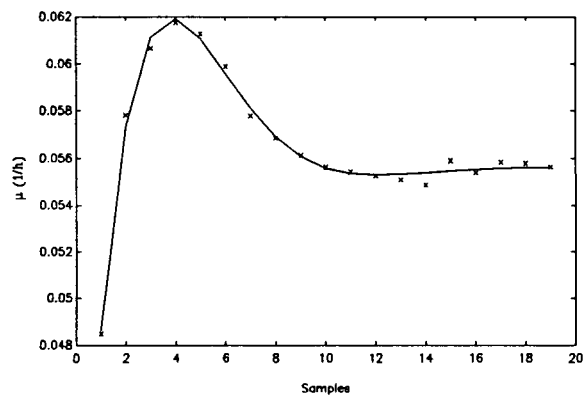
table 2: Model poles and zeros and resulting controller

model		controller	
zeros	poles	zeros	poles
0.6430+0.2581i	0.8422	0.8422	0.6430+0.2581i
0.6430-0.2581i	-57.7310	1/-57.7310	0.6430-0.2581i

As the precompensation operates on a feed-forward basis, its sample rate need not be limited by the sample rate of the measuring equipment. To improve control performance, the precompensation is run at twice the sample rate of the measuring equipment. The transfer function is converted to continuous time and then back to discrete time using MatLab's d2c and c2d. Theoretically, even higher sample rates can be used. The higher the sample rate, however, the more sensitive the precompensation will be to errors in the linear model. This is undesirable, because the model is based on a few samples only and because the true process is probably not in the model set used. The precompensation response and the precompensation & process response to a step change are depicted in figure 47 and figure 48.
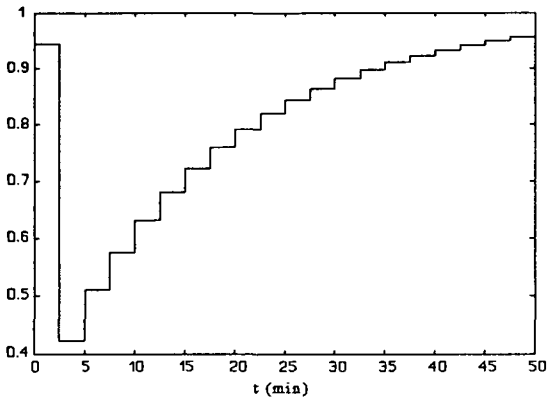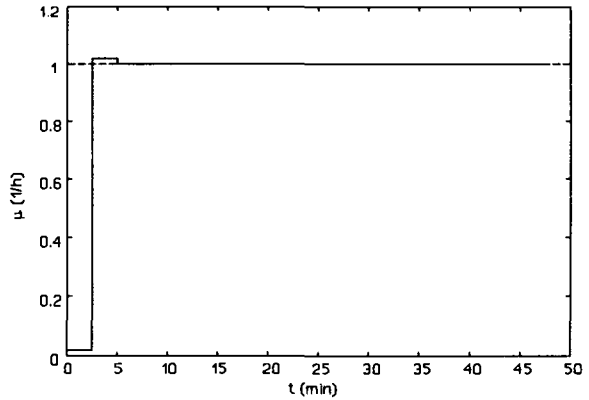
figure 47: dynamic precomp. step-response

figure 48: step-response with dynamic
precomp.

The response of the process is improved by the addition of a dynamic precompensation. In figure 49 results are shown for steps of the order 0.01 h^{-1}. This is comparable to the size of the step on which the prediction model is based. Within roughly one or two hours a set-point change is completed. In previous experiments, set-point changes took 3 hours and longer. Overshoot is still present but small compared to experiments without dynamic precompensation.
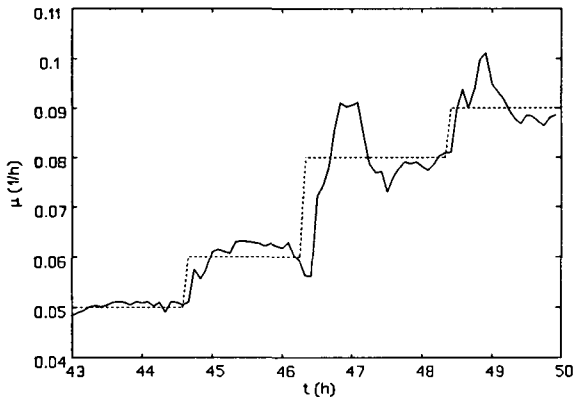

figure 49: Response to small steps

For large steps (not shown in figure 49), large overshoot is still present on the response. Yeast behaviour for large steps is apparently not a linear extrapolation from the behaviour for small steps. This is not surprising. Remarkable about the response to large steps with dynamic precompensation is, that no ethanol production occurs, even though the specific growth rate reaches temporarely values for which ethanol production was observed in other experiments.

5.10. Future work

Little experience has been gained with respect to the dynamic precompensation. First results are promising enough to look deeper into the possibilities of this approach. As the prediction model in the dynamic precompensation is based on a single step-response, further open-loop experiments are desirable to provide a more solid basis for this model. A next step could be to identify a model for larger steps. The dynamic precompensation should have a mechanism to choose the appropriate model based on the size of the set-point change.

Once step changes are mastered, it may be worthwhile to investigate the possibilities for μ profiles that do not necessarily contain constant parts. Control during ethanol production and consumption has not been intentionally tested yet. Ethanol estimates were not accurate enough for this. In the past months, a continuous ethanol sensor has become available. Incorporating this in the current control structure is likely to improve performance, especially during ethanol production and consumption. For ethanol production and consumption, the dependence of μ on GF/VX and E probably has to be identified in experiments designed explicitly for this purpose.

figure 50: Final control structure
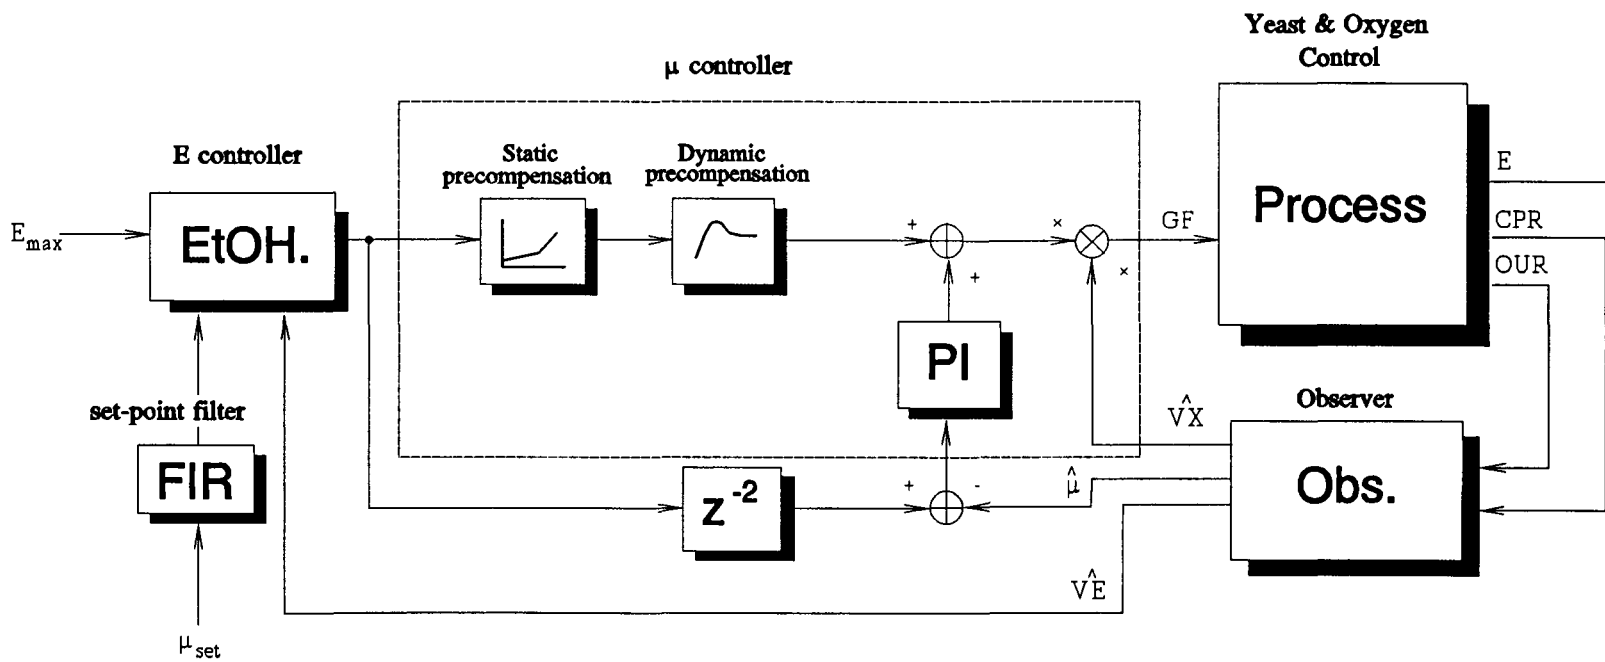
6. Conclusions and recommendations

Theoretical analysis as well as simulation studies show, that the specific growth rate estimate is rather insensitive to errors in the observer parameters. The ethanol concentration estimate is more sensitive to errors in these data.

The control structure is able to maintain a constant set-point on μ while obeying a limit on the ethanol concentration. Performance deteriorates if ethanol production or consumption occurs. In the simulations this could be counteracted by incorporating extra a priori knowledge in the control. If the limit on the ethanol concentration is exceeded, the ethanol concentration is successfully brought back within the allowed range. Overshoot is too large to consider the current control structure a candidate for ethanol control.

In experiments, occasional ethanol production or consumption indicates that control is poor under these conditions. Accurate information on the the ethanol concentration is badly needed then. This may be realised by using an ethanol sensor, which became available in the course of this research. New yeast properties were discovered: the yield for oxidative growth on glucose increases throughout an experiment, but converges to a constant value. The time constant associated to this behaviour is large enough to be handled by the control structure.

Changes in air flow and stirrer speed disturb the oxygen uptake rate and carbondioxide production rate measurements, and consequently the specific growth rate estimates. To enable μ control the oxygen controller should be switched to manual mode, until a rate limitation has been implemented on *AF* and *SS*.

Open-loop experiments uncovered dynamic effects during μ changes. Compensating for these dynamics by a predictive feed-forward controller improves performance. A prediction model is currently only available for small steps. Models for larger steps should be identified in further experiments.

Only part of the yeast behaviour is covered by the experiments carried out. Ethanol production and consumption have not been investigated properly. The control structure has provisions for these situations but these are tuned to the simulation model. Identification experiments should be carried out first to tune the structure to the real process.

References

[1] **Agrawal, P., Koshy, G., Ramseier, M.,** (1989), *An algorithm for operating a fed-batch fermentor at optimum specific growth rate*, Biotechnology and bioengineering, Vol. 33, pp. 115-125

[2] **Axelsson, J.P.,** (1989), *Modelling and control of fermentation processes*, PhD thesis, Lund institute of technology, Sweden.

[3] **Bastin, G. and Dochain, D.** (1990), *On-line estimation and adaptive control of bioreactors*, Elsevier, Amsterdam

[4] **O'Connor, G.M., Sanchez-Riera, F., Cooney, C.L.,** (1992), *Design and evaluation of control strategies for high cell density fermentations*, submitted to Biotechnology and bioengineering

[5] **Keulers, M., Reyman, G.,** (1991), *The application of the GPC algorithm to a fed-batch fermentation process. A simulation study.* Proceedings of the European Simulation Conference, Copenhagen.

[6] **Landau, I.D., Samaan, M., M'Saad, M.,** (1990), *Robust performance oriented adaptive control for biotechnological processes, a tutorial*, Proceedings of the American Control Conference 1990, pp. 2684-2687, San Diego

[7] **Pomerleau, Y., Perrier, M., Dochain, D.,** (1989), *Adaptive nonlinear control of the bakers' yeast fed-batch fermentation*, Proceedings of the American Control Conference 1989, pp. 2424-2429, Pitsburg

[8] **Poulisse, H.N.J., Helden, C. van,** (1985), *Adaptive LQ control of fermentation processes*, IFAC Modelling and Control of Biotechnological Processes, pp. 43-47, Noordwijkerhout, Holland

[9] **Queinnec, L., Dahou, B., M'Saad, M.,** (1991), *On adaptive control of fed-batch fermentation processes*, Proceedings of the European Control Conference 1991, Grenoble

[10] **Soeterboek, A.R.M.** *et al.*, (1992), *Estimation of biomass concentration and specific growth rate of fermentation processes*, Internal Unilever communication

[11] **Sonnleitner, B., Käppeli, O.,** (1986), *Growth of saccharomyces cerevisiae is controlled by its limited respiratory capacity: formulation and verification of a hypothesis*, Biotechnology and bioengineering, Vol. 28

[12] **Verbruggen, H.D., Eeldering, G.H.B., Broecke, P.M. van,** (1985), *Multiloop controlled fed-batch fermentation processes using a selftuning controller*, IFAC Modelling and Control of Biotechnological Processes, pp. 121-126, Holland

[13] **Williams, D., Yousefpour, P., Wellington, E.M.H.,** (1986), *On-line adaptive control of a fed-batch fermentation of Saccharomyces cerevisiae*, Biotechnology and bioengineering, Vol. 28, pp. 631-645

Appendix 1: Dynamic model

The model of Sonnleitner and Käppeli (1986) assumes, that the rates for the reactions (1), (2), (3), (4) depend only on the current concentrations of glucose, biomass, ethanol and oxygen. These four quantities together with the total volume make up the states of this model.

The reaction rates are basically given by Monod kinetics. Some rules have been added to combine the four reactions of (1), (2), (3), (4). These are

- The maintenance reaction takes place at a rate that is proportional to the total amount of yeast. The effect of the other state variables, including the oxygen concentration, is neglected. The maintenance reaction is an extension to (Sonnleitner and Käppeli, 1986).
- Ethanol consumption and ethanol production are mutually exclusive.
- Oxidative growth on glucose has the highest priority as both glucose consumption and oxygen consumption is concerned. If more glucose is present than can be consumed with the available oxygen concentration, fermentative growth on glucose occurs. If more oxygen is present than is needed for the consumption of glucose, oxidative growth on ethanol occurs. Growth on ethanol occurs only in the presence of ethanol, of course.

These rules can be viewed as a bottleneck principle:

    ▨ Oxidative growth on glucose
    ◰ Oxidative growth on ethonol
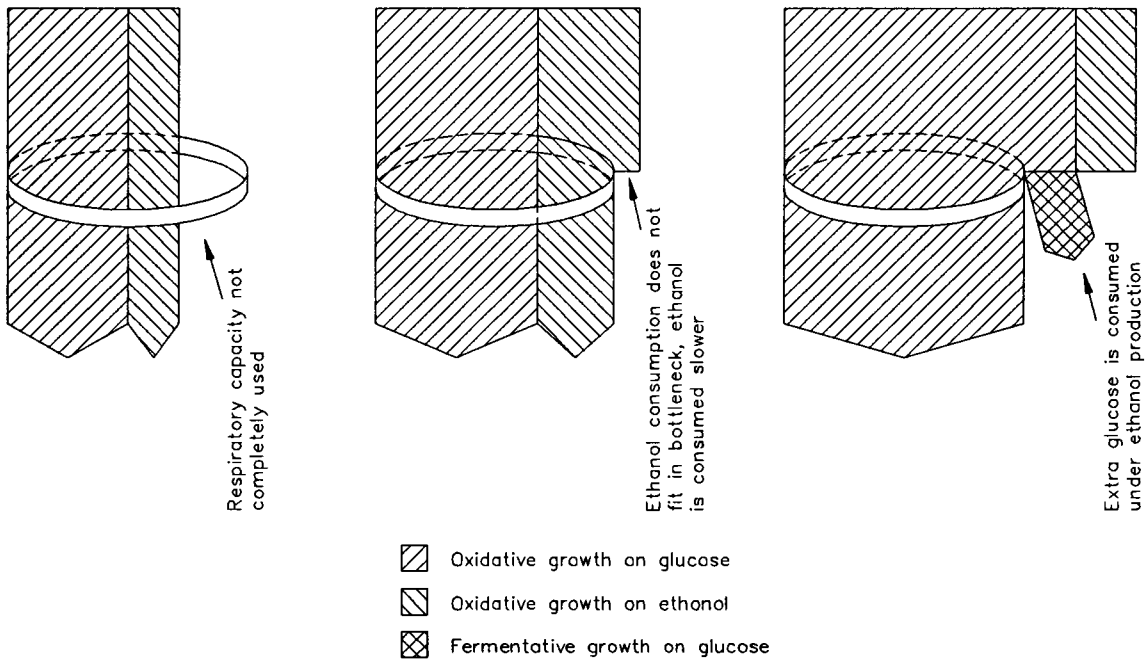   ▩ Fermentative growth on glucose

figure A.1: Bottleneck principle

The width of the arrows in figure A.1 represents the rates of the three pathways. The maintenance reaction is not considered in this bottleneck principle. The widths above the bottleneck represent the reaction rates that would occur if none of the other reaction took place. The widths under the bottleneck are the reaction rates after being fitted in the bottleneck. In the left-most picture is shown, the situation in which the respiratory capacity is not fully utilised. The respiratory capacity is sufficient to consume both glucose and ethanol without extra restrictions. In the second figure, the rates for ethanol consumption and glucose consumption together do not fit in the bottleneck. The rate of the ethanol consuming path is then reduced, until they do. In the right-most figure the glucose consumption rate on its own does not fit into the bottleneck. Part of the glucose is then

consumed under ethanol production.

The width of the bottleneck is a function of the oxygen concentration. More precisely, it is proportional to

$$\frac{P}{K_o + P}, \quad K_o = 3.0 \cdot 10^{-7} \tag{A.1}$$

These assumptions lead to the following model:

If

$$0.5 \cdot \frac{G}{K_s + G} \geq \frac{0.2}{a} \cdot \frac{P}{K_o + P} \tag{A.2}$$

it is decided that the yeast is overfed; else if

$$0.13 \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} < \frac{1}{k} \cdot \left[0.2 \cdot \frac{P}{K_o + P} - 0.5 \cdot a \cdot \frac{G}{K_s + G} \right] \tag{A.3}$$

the yeast is underfed under ethanol limitation; otherwise the yeast is underfed under oxygen limitation.

<div align="center">Yeast is overfed</div>

$$\dot{G} = -0.5 \cdot \frac{G}{K_s + G} \cdot X - \frac{GF}{V} \cdot (G - G_{in})$$

$$\dot{X} = \left[0.2 \cdot \frac{b}{a} \cdot \frac{P}{K_o + P} + g \cdot \left[0.5 \cdot \frac{G}{K_s + G} - \frac{0.2}{a} \cdot \frac{P}{K_o + P} \right] - \frac{ms}{\alpha} \right] \cdot X - \frac{GF}{V} \cdot X$$

$$\dot{E} = j \cdot \left[0.5 \cdot \frac{G}{K_s + G} - \frac{0.2}{a} \cdot \frac{P}{K_o + P} \right] \cdot X - \frac{GF}{V} \cdot E$$

$$\dot{P} = \left[-0.2 \cdot \frac{P}{K_o + P} - ms \right] \cdot X + C_1 \cdot SS^{C_2} \cdot AF^{C_3} \cdot \left(O_{max} - P \right) - \frac{GF}{V} \cdot P \tag{A.4}$$

$$\dot{V} = GF$$

$$OUR = \left[0.2 \cdot \frac{P}{K_o + P} + ms \right] \cdot X$$

$$CPR = \left[0.2 \cdot \frac{c}{a} \cdot \frac{P}{K_o + P} + h \cdot \left[0.5 \cdot \frac{G}{K_s + G} - \frac{0.2}{a} \cdot \frac{P}{K_o + P} \right] + \frac{ms}{\alpha} \right] \cdot X$$

Yeast is underfed, ethanol limitation

$$\dot{G} = -0.5 \cdot \frac{G}{K_s + G} \cdot X - \frac{GF}{V} \cdot (G - G_{in})$$

$$\dot{X} = \left[0.5 \cdot b \cdot \frac{G}{K_s + G} + 0.13 \cdot l \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} - \frac{ms}{\alpha} \right] \cdot X - \frac{GF}{V} \cdot X$$

$$\dot{E} = -0.13 \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} \cdot X - \frac{GF}{V} \cdot E$$

$$\dot{P} = \left[-0.5 \cdot a \cdot \frac{G}{K_s + G} - 0.13 \cdot k \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} - ms \right] \cdot X +$$

$$C_1 \cdot SS^{C_2} \cdot AF^{C_3} \cdot (O_{max} - P) - \frac{GF}{V} \cdot P$$

$$\dot{V} = GF$$

(A.5)

$$OUR = \left[0.5 \cdot a \cdot \frac{G}{K_s + G} + 0.13 \cdot k \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} + ms \right] \cdot X$$

$$CPR = \left[0.5 \cdot c \cdot \frac{G}{K_s + G} + 0.13 \cdot m \cdot \frac{E}{K_e + E} \cdot \frac{K_i}{K_i + G} + \frac{ms}{\alpha} \right] \cdot X$$

Yeast is underfed, oxygen limitation

$$\dot{G} = -0.5 \cdot \frac{G}{K_s + G} \cdot X - \frac{GF}{V} \cdot (G - G_{in})$$

$$\dot{X} = \left[0.5 \cdot b \cdot \frac{G}{K_s + G} + \frac{l}{k} \cdot \left[0.2 \cdot \frac{P}{K_o + P} - 0.5 \cdot a \cdot \frac{G}{K_s + G} \right] - \frac{ms}{\alpha} \right] \cdot X - \frac{GF}{V} \cdot X$$

$$\dot{E} = -\frac{1}{k} \cdot \left[0.2 \cdot \frac{P}{K_o + P} - 0.5 \cdot a \cdot \frac{G}{K_s + G} \right] \cdot X - \frac{GF}{V} \cdot E$$

$$\dot{P} = \left[-0.2 \cdot \frac{P}{K_o + P} - ms \right] \cdot X + C_1 \cdot SS^{C_2} \cdot AF^{C_3} \cdot (O_{max} - P) - \frac{GF}{V} \cdot P$$

$$\dot{V} = GF$$

(A.6)

$$OUR = \left[0.5 \cdot a \cdot \frac{G}{K_s + G} + \left[0.2 \cdot \frac{P}{K_o + P} - 0.5 \cdot a \cdot \frac{G}{K_s + G} \right] + ms \right] \cdot X$$

$$CPR = \left[0.5 \cdot c \cdot \frac{G}{K_s + G} + \frac{m}{k} \cdot \left[0.2 \cdot \frac{P}{K_o + P} - 0.5 \cdot a \cdot \frac{G}{K_s + G} \right] + \frac{ms}{\alpha} \right] \cdot X$$

Some properties of this model are analyzed further by means of the figures on the next pages. In figure A.2 the behaviour of the specific growth rate is analyzed. The specific ethanol production rate is investigated by figure A.3. The critical glucose concentration as shown in figure A.3 is the glucose concentration at the switching point of ethanol limited ethanol consumption and oxygen limited ethanol consumption. It is *not* the concentration at the switching point between ethanol production and consumption as used in other parts of this report. For the μ-controller, figure A.2

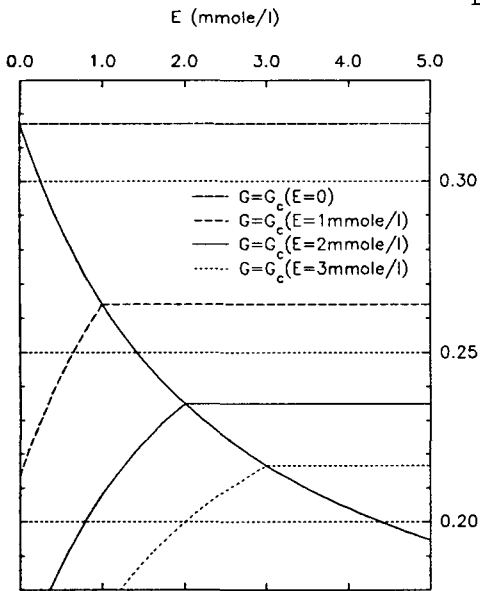is of interest, figure A.3 could be used for an ethanol controller design.
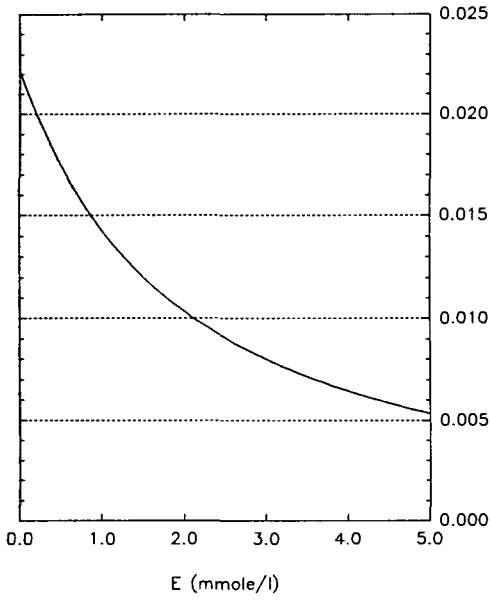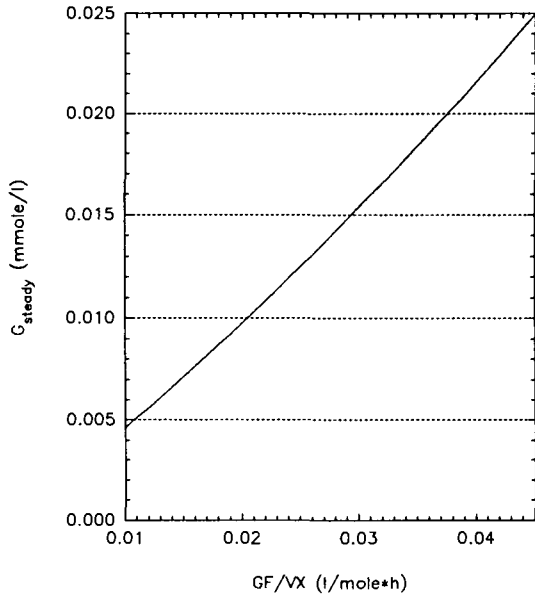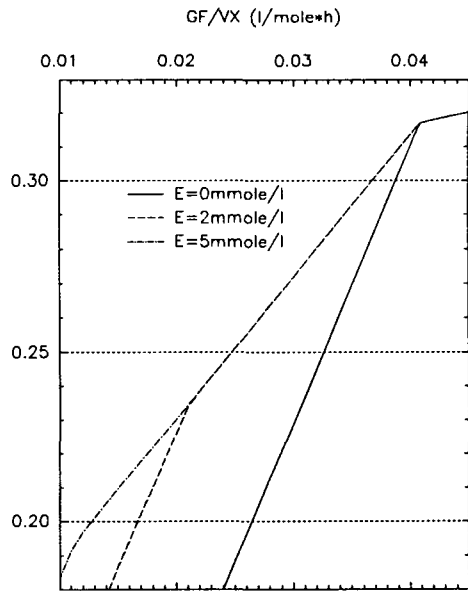
a: Critical glucose concentration   b: Steady state glucose concentration
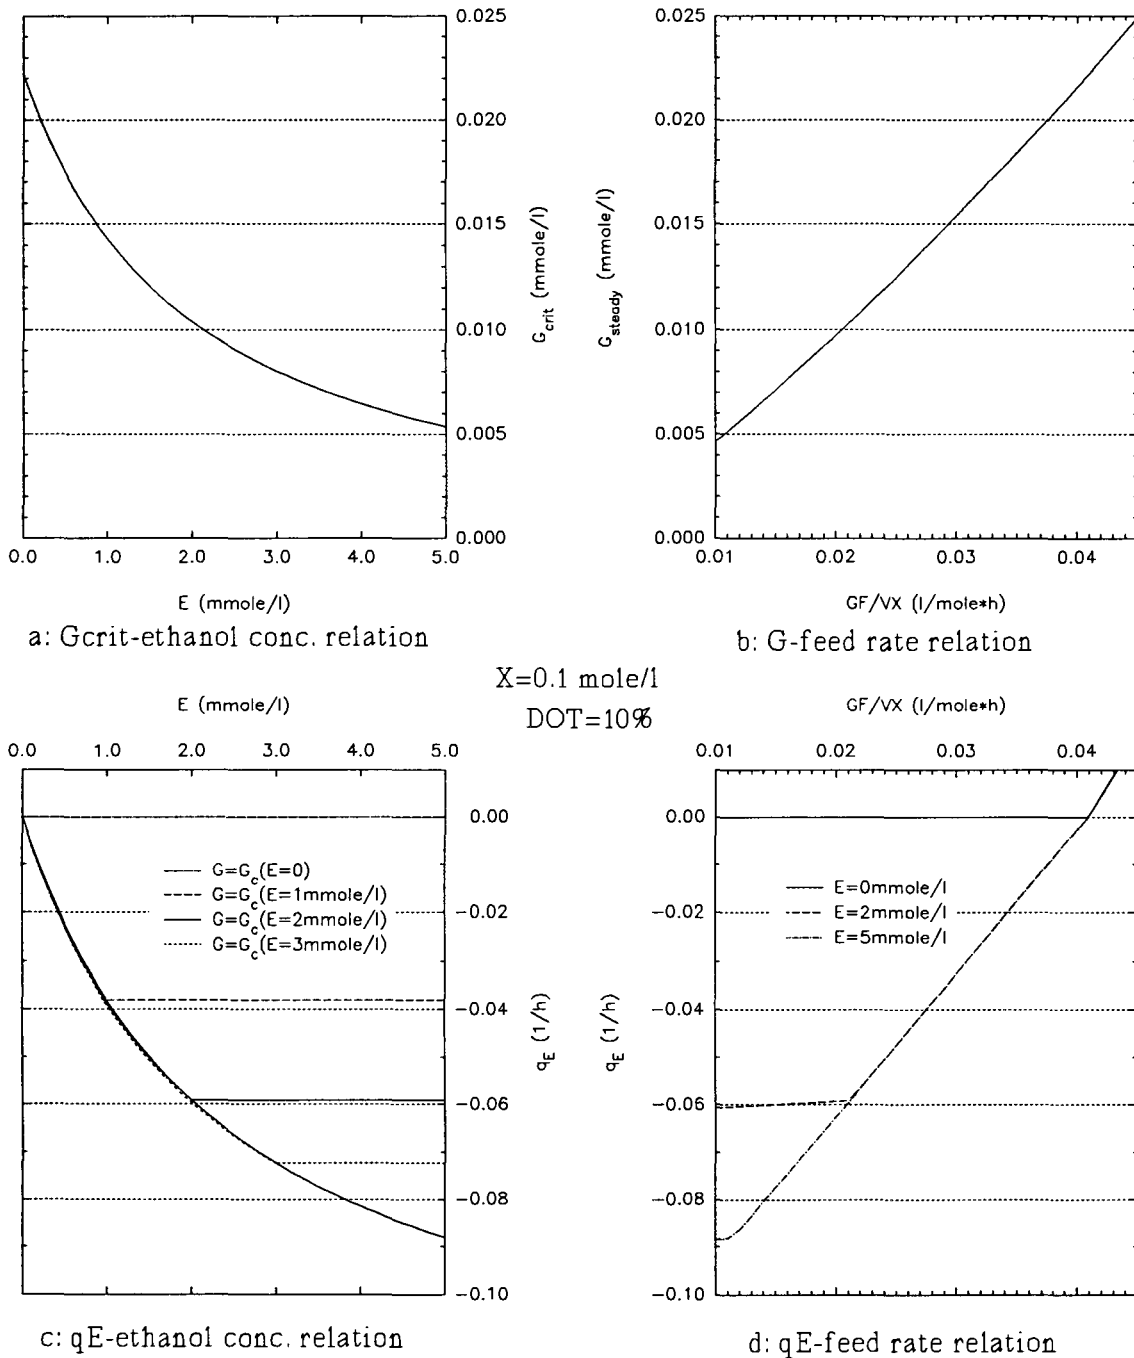
X=0.1 mole/l

DOT=10%

c: μ-ethanol conc. relation   d: μ-feed rate relation

figure A.2: μ behaviour

a: Gcrit-ethanol conc. relation

b: G-feed rate relation

X=0.1 mole/l
DOT=10%

c: qE-ethanol conc. relation

d: qE-feed rate relation

figure A.3: q_e behaviour

Appendix 2: Additional observer information

2.1. Calculations for the sensitivity analysis

The observer algorithm can be written as

$$
\begin{bmatrix} q_s \cdot VX \\ \mu \cdot VX \\ q_e \cdot VX \end{bmatrix} = A_i \cdot \begin{bmatrix} OUR_a - ms \cdot VX \\ CPR_a - \dfrac{ms}{\alpha} \cdot VX \end{bmatrix} + \begin{bmatrix} 0 \\ -\dfrac{ms}{\alpha} \cdot VX \\ 0 \end{bmatrix} \quad i \in \{p,\, c,\, 0\} \tag{A.7}
$$

The matrix used during ethanol production shall be denoted A_p, the one used during ethanol consumption is A_c and the remaining matrix, used if there is no consumption or production of alcohol, is A_0. The following holds:

$$
A_p = \begin{bmatrix} \dfrac{h-c}{ah} & \dfrac{1}{h} \\[2mm] \dfrac{bh-gc}{ah} & \dfrac{g}{h} \\[2mm] -\dfrac{jc}{ah} & \dfrac{j}{h} \end{bmatrix} \quad
A_c = \begin{bmatrix} \dfrac{m}{am-kc} & -\dfrac{k}{am-kc} \\[2mm] \dfrac{bm-lc}{am-kc} & -\dfrac{bk-la}{am-kc} \\[2mm] \dfrac{c}{am-kc} & -\dfrac{a}{am-kc} \end{bmatrix} \quad
A_0 = \begin{bmatrix} \dfrac{1}{2a} & \dfrac{1}{2c} \\[2mm] \dfrac{b}{2a} & \dfrac{b}{2c} \\[2mm] 0 & 0 \end{bmatrix} \tag{A.8}
$$

For the sensitivity analysis, derivatives of A_i, $i \in \{p,\, c,\, 0\}$ are needed. These are given below.

$$
\frac{\partial A_p}{\partial a} = \begin{bmatrix} -\dfrac{h-c}{a^2 h} & 0 \\[2mm] -\dfrac{bh-gc}{a^2 h} & 0 \\[2mm] \dfrac{jc}{a^2 h} & 0 \end{bmatrix} \quad
\frac{\partial A_c}{\partial a} = \begin{bmatrix} \dfrac{-m^2}{(am-kc)^2} & \dfrac{mk}{(am-kc)^2} \\[2mm] -m\dfrac{bm-lc}{(am-kc)^2} & k\dfrac{bm-lc}{(am-kc)^2} \\[2mm] \dfrac{-mc}{(am-kc)^2} & \dfrac{kc}{(am-kc)^2} \end{bmatrix} \quad
\frac{\partial A_0}{\partial a} = \begin{bmatrix} \dfrac{-1}{2a^2} & 0 \\[2mm] \dfrac{-b}{2a^2} & 0 \\[2mm] 0 & 0 \end{bmatrix} \tag{A.9}
$$

$$
\frac{\partial A_p}{\partial b} = \begin{bmatrix} 0 & 0 \\ \dfrac{1}{a} & 0 \\ 0 & 0 \end{bmatrix} \quad
\frac{\partial A_c}{\partial b} = \begin{bmatrix} 0 & 0 \\ \dfrac{m}{am-kc} & \dfrac{-k}{am-kc} \\ 0 & 0 \end{bmatrix} \quad
\frac{\partial A_0}{\partial b} = \begin{bmatrix} 0 & 0 \\ \dfrac{1}{2a} & \dfrac{1}{2c} \\ 0 & 0 \end{bmatrix} \tag{A.10}
$$

$$
\frac{\partial A_p}{\partial g} = \begin{bmatrix} 0 & 0 \\ \dfrac{-c}{ah} & \dfrac{1}{h} \\ 0 & 0 \end{bmatrix} \quad
\frac{\partial A_c}{\partial g} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \quad
\frac{\partial A_0}{\partial g} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \tag{A.11}
$$

$$\frac{\partial A_p}{\partial h} = \begin{bmatrix} \dfrac{c}{ah^2} & \dfrac{-1}{h^2} \\[2ex] \dfrac{gc}{ah^2} & \dfrac{-g}{h^2} \\[2ex] \dfrac{jc}{ah^2} & \dfrac{-j}{h^2} \end{bmatrix} \qquad \frac{\partial A_c}{\partial g} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_0}{\partial g} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad (A.12)$$

$$\frac{\partial A_p}{\partial j} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \dfrac{-c}{ah} & \dfrac{1}{h} \end{bmatrix} \qquad \frac{\partial A_c}{\partial j} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_0}{\partial j} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad (A.13)$$

$$\frac{\partial A_p}{\partial k} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_c}{\partial k} = \begin{bmatrix} \dfrac{mc}{(am-kc)^2} & \dfrac{-am}{(am-kc)^2} \\[2ex] c\dfrac{bm-lc}{(am-kc)^2} & -a\dfrac{bm-lc}{(am-kc)^2} \\[2ex] \dfrac{c^2}{(am-kc)^2} & \dfrac{-ac}{(am-kc)^2} \end{bmatrix} \qquad \frac{\partial A_0}{\partial k} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad (A.14)$$

$$\frac{\partial A_p}{\partial l} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_c}{\partial l} = \begin{bmatrix} 0 & 0 \\ \dfrac{-c}{am-kc} & \dfrac{a}{am-kc} \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_0}{\partial l} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad (A.15)$$

$$\frac{\partial A_p}{\partial m} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad \frac{\partial A_c}{\partial m} = \begin{bmatrix} \dfrac{-kc}{(am-kc)^2} & \dfrac{ka}{(am-kc)^2} \\[2ex] -c\dfrac{bk-al}{(am-kc)^2} & a\dfrac{bk-al}{(am-kc)^2} \\[2ex] \dfrac{-ac}{(am-kc)^2} & \dfrac{a^2}{(am-kc)^2} \end{bmatrix} \qquad \frac{\partial A_0}{\partial m} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{bmatrix} \qquad (A.16)$$

The sensitivity of the estimated specific glucose consumption rate is given by

$$S_{q,x} = \frac{\partial\left[(A_i)_{1,1}\cdot OUR' + (A_i)_{1,2}\cdot RQ'\cdot OUR'\right]}{\partial x} \cdot \frac{x}{(A_i)_{1,1}\cdot OUR' + (A_i)_{1,2}\cdot RQ'\cdot OUR'}$$

$$= \frac{\left[\frac{\partial A_i}{\partial x}\right]_{1,1} + RQ'\cdot\left[\frac{\partial A_i}{\partial x}\right]_{1,2}}{(A_i)_{1,1} + RQ'\cdot(A_i)_{1,2}}\cdot x \qquad i\in\{p,c,0,\}, \quad x\in\{a,b,c,g,h,j,k,l,m\}$$

(A.17)

Apparently the sensitivity is only dependent on RQ' and independent of the absolute value of OUR'. This property was already mentioned in the text. Best and worst case values are found at minimum and maximum values for $S_{q,x}$. These are found, either at the end of the permitted range for RQ' for a given i, or at a point in which

$$\frac{\partial S_{q,x}}{\partial RQ'} = \frac{\left[\frac{\partial A_i}{\partial x}\right]_{1,2}\cdot(A_i)_{1,1} - \left[\frac{\partial A_i}{\partial x}\right]_{1,1}\cdot(A_i)_{1,2}}{\left[(A_i)_{1,1} + RQ'\cdot(A_i)_{1,2}\right]^2}\cdot x$$

(A.18)

changes sign. Another possible "best case" situation occurs, when the quantity defined by (A.17) passes through zero. Formula (A.18) does not change sign. This holds not only for the $q_c\cdot VX$ estimate, but also for the $q_s\cdot VX$ estimate.

The situation for $\mu\cdot VX$ is a bit more complicated because of the term $-ms/\alpha$.

$$S_{\mu,x} = \frac{\partial\left[(A_i)_{2,1}\cdot OUR' + (A_i)_{2,2}\cdot RQ'\cdot OUR' - ms\cdot VX/\alpha\right]}{\partial x}\cdot$$

$$\cdot\frac{x}{(A_i)_{2,1}\cdot OUR' + (A_i)_{2,2}\cdot RQ'\cdot OUR' - ms\cdot VX/\alpha}$$

$$= \frac{\left[\frac{\partial A_i}{\partial x}\right]_{2,1} + RQ'\cdot\left[\frac{\partial A_i}{\partial x}\right]_{2,2}}{(A_i)_{2,1}\cdot OUR' + (A_i)_{2,2}\cdot RQ'\cdot OUR' - ms\cdot VX/\alpha}\cdot OUR'\cdot x$$

(A.19)

$$= \frac{\left[\frac{\partial A_i}{\partial x}\right]_{2,1} + RQ'\cdot\left[\frac{\partial A_i}{\partial x}\right]_{2,2}}{(A_i)_{2,1} + (A_i)_{2,2}\cdot RQ' - \dfrac{ms/\alpha}{OUR'/VX}}\cdot x \qquad i\in\{p,c,0,\}, \quad x\in\{a,b,c,g,h,j,k,l,m\}$$

A candidate for the location of an extreme value for $S_{\mu,x}$ is a point in which

$$\frac{\partial S_{\mu,x}}{\partial RQ'} = \frac{\left[\frac{\partial A_i}{\partial x}\right]_{2,2}\cdot(A_i)_{2,1} - \left[\frac{\partial A_i}{\partial x}\right]_{2,1}\cdot(A_i)_{2,2} - \left[\frac{\partial A_i}{\partial x}\right]_{2,2}\cdot\dfrac{ms/\alpha}{OUR'/VX}}{\left[(A_i)_{2,1} + RQ'\cdot(A_i)_{2,2} - \dfrac{ms/\alpha}{OUR'/VX}\right]^2}\cdot x$$

(A.20)

changes sign. When there is no ethanol consumption nor production, both RQ' and OUR'/VX are constant, so that there is only one sensitivity value. During ethanol production, RQ' varies and

ethanol consumption, both RQ' and OUR' vary. Now the denominators of (A.20) and (A.19) pass through zero, giving worst case sensitivity values of ∞.

2.2. Changes made to the original observer

During the simulations, two changes were made to the original observer. First, the integration algorithm used is currently based on constant interpolation between two sample instants instead of on linear extrapolation. This improved the short term accuracy of the observer in the simulations. It can not be verified, whether this applies to the real behaviour as well. The off-line measurements needed to verify this simply do not provide enough information: the rate with which they can be taken is too low and the uncertainty in their results is too high.

The second change was to prohibit ethanol consumption for high specific growth rates and ethanol production for low specific growth rates. Determining which growth rates are to be considered high and which are to be considered low requires the value of the critical growth rate. The discrepancy between the true critical growth rate and the expected critical growth rate was so large, that is was no longer considered safe to use this knowledge in practical situations. This a priori knowledge was consequently removed from the observer.

Appendix 3: Observer sensitivity values

table A.1: *RQ'* boundary values

	lowest value	highest value
Ethanol production	1.08	1.60
Ethanol consumption	0.42	1.08
Otherwise	1.08	1.08

table A.2: Sensitivity of biomass growth rate

	Ethanol production		Ethanol consumption		Otherwise
	worst case	best case	worst case	best case	
a	-0.92	-0.86	$\pm\infty$	0.00	-0.50
b	1.04	0.98	$\pm\infty$	0.00	1.00
c	-0.13	-0.12	$\pm\infty$	0.00	-0.50
g	0.06	0.00	0	0	0
h	-0.06	0.00	0	0	0
j	0	0	0	0	0
k	0	0	$\pm\infty$	0.00	0
l	0	0	$\pm\infty$	0.00	0
m	0	0	$\pm\infty$	0.00	0

table A.3: Sensitivity of ethanol production rate

	Ethanol production		Ethanol consumption	
	worst case	best case	worst case	best case
a	∞	2.08	∞	0.00
b	0	0	0	0
c	∞	-2.08	∞	0.00
g	0	0	0	0
h	-1.00	-1.00	0	0
j	1.00	1.00	0	0
k	0	0	-1.64	-1.64
l	0	0	0	0
m	0	0	0.64	0.64

table A.4: Sensitivity of glucose consumption rate

	Ethanol production		Ethanol consumption		Otherwise
	worst case	best case	worst case	best case	
a	0.24	0.15	0.64	0.00	-0.50
b	0	0	0	0	0
c	-1.24	-0.78	-1.64	0.00	-0.50
g	0	0	0	0	0
h	-0.37	0.00	0	0	0
j	0	0	0	0	0
k	0	0	$+\infty$	0.00	0
l	0	0	0	0	0
m	0	0	$-\infty$	0.00	0

table A.5: Sensitivity of the glucose consumption rate to changes in ms and α

	Ethanol production		Ethanol consumption		Otherwise	
	worst case	best case	worst case	best case	worst case	best case
ms	0.06	0.01	∞	0.06	∞	0.06
α	-0.08	-0.01	$-\infty$	-0.11	$-\infty$	-0.03

table A.6: Sensitivity of the ethanol production rate to changes in ms and α

	Ethanol production		Ethanol consumption	
	worst case	best case	worst case	best case
ms	∞	0.00	$-\infty$	-0.02
α	∞	0.02	$-\infty$	-0.11

table A.7: Sensitivity of the biomass growth rate to changes in ms and α

	Ethanol production		Ethanol consumption		Otherwise	
	worst case	best case	worst case	best case	worst case	best case
ms	-0.12	-0.08	$\pm\infty$	-0.11	$\pm\infty$	-0.49
α	-0.05	-0.04	$\pm\infty$	-0.10	$\pm\infty$	-0.32

Appendix 4: Simulation results
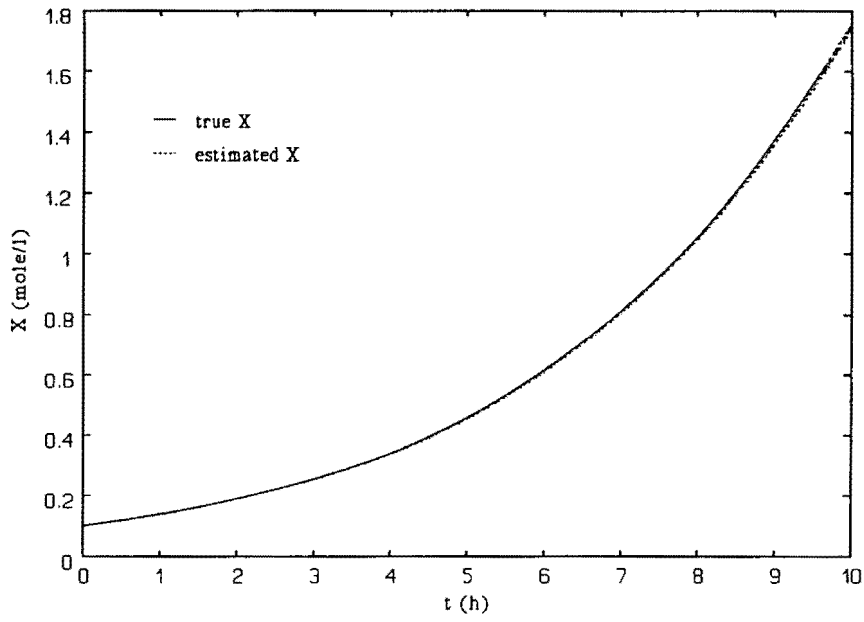
figure A.4: Biomass concentration estimate, nominal model
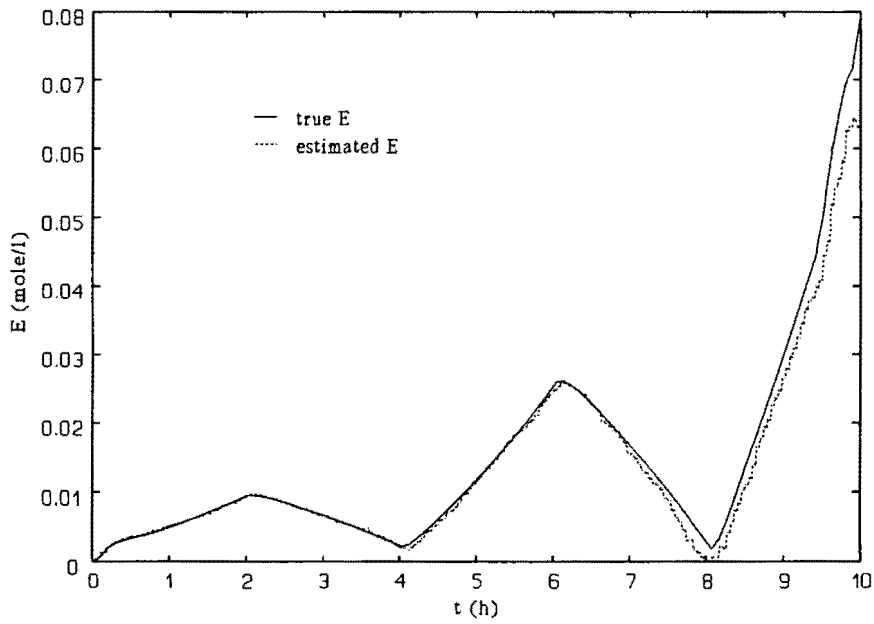
figure A.5: Ethanol concentration estimate, nominal model
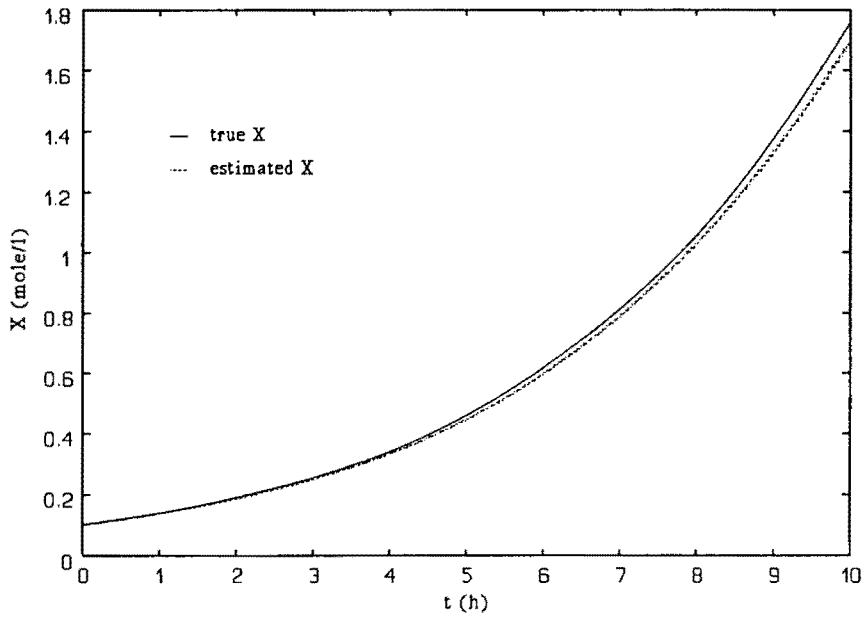
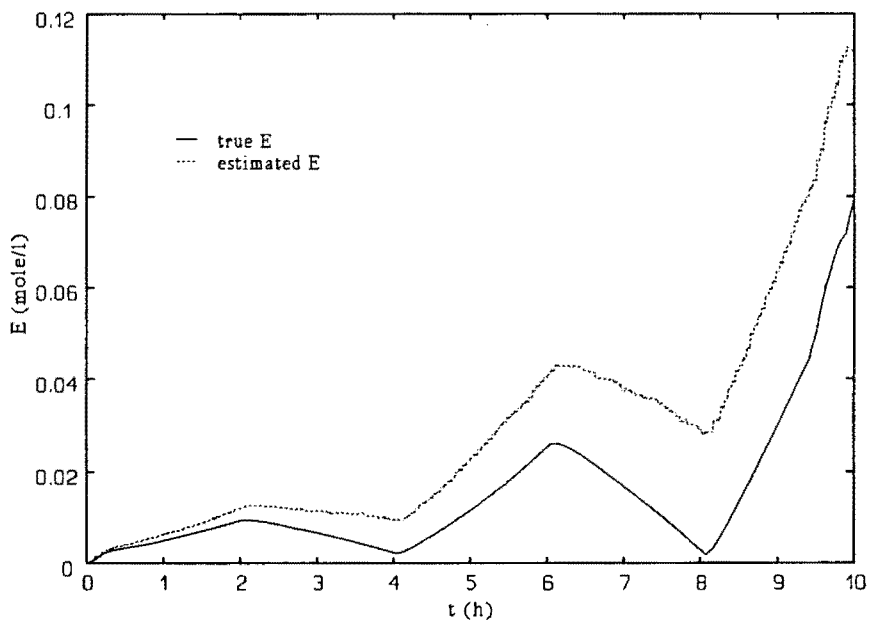figure A.6: Biomass concentration estimate, $a = 1.05 \cdot a_{true}$

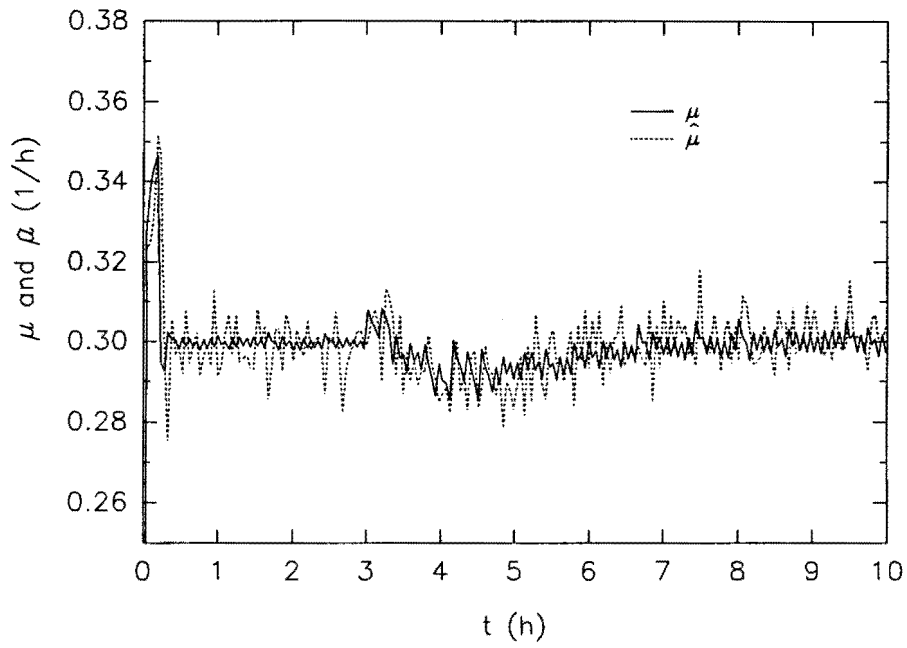figure A.7: ethanol concentration, $a = 1.05 \cdot a_{true}$

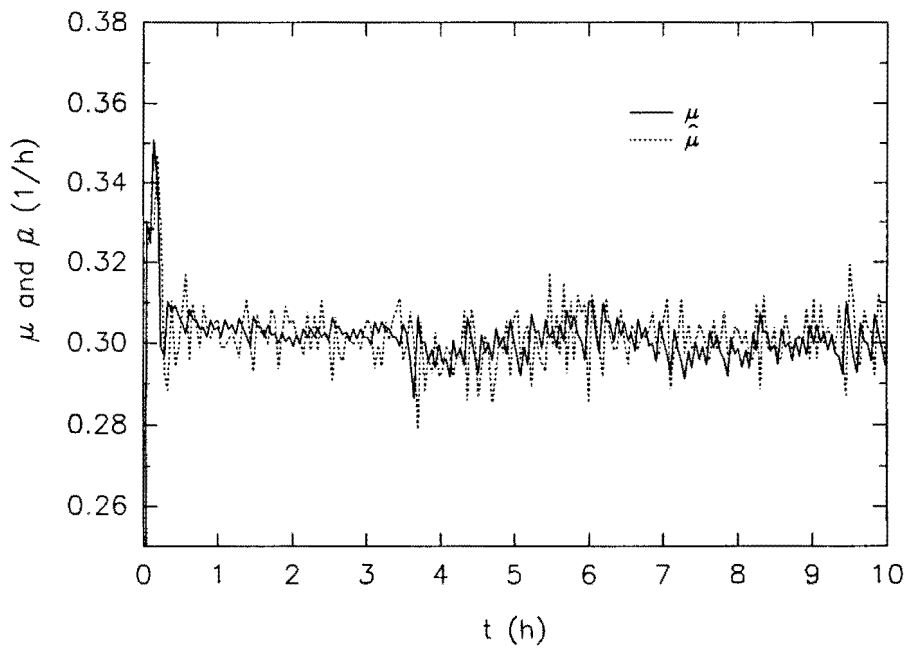figure A.8: Specific growth rate, precompensation 1, $\mu_{set} = 0.30$ h^{-1}

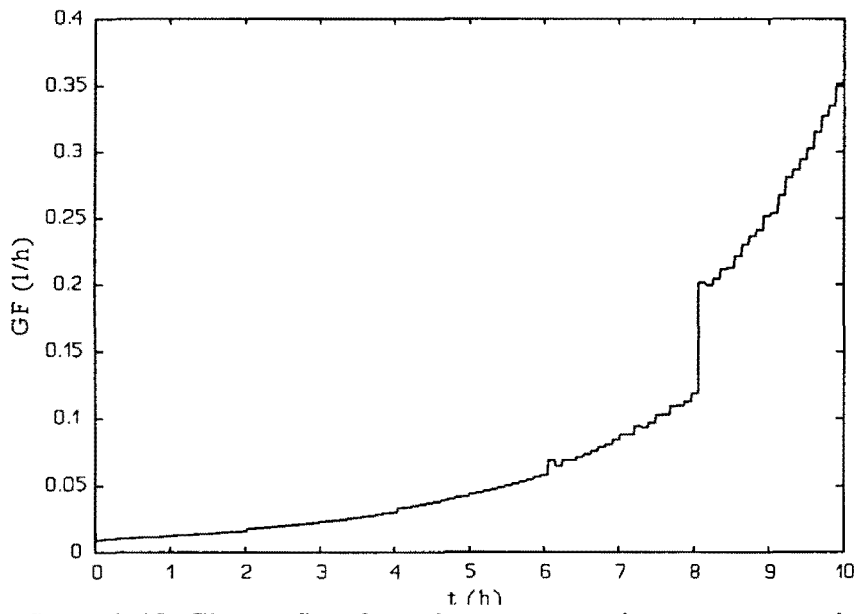figure A.9: Specific growth rate, precompensation 2, $\mu_{set} = 0.30$ h^{-1}

figure A.10: Glucose flow for stair-case μ set-point, precompensation 4

Appendix 5: Experimental results
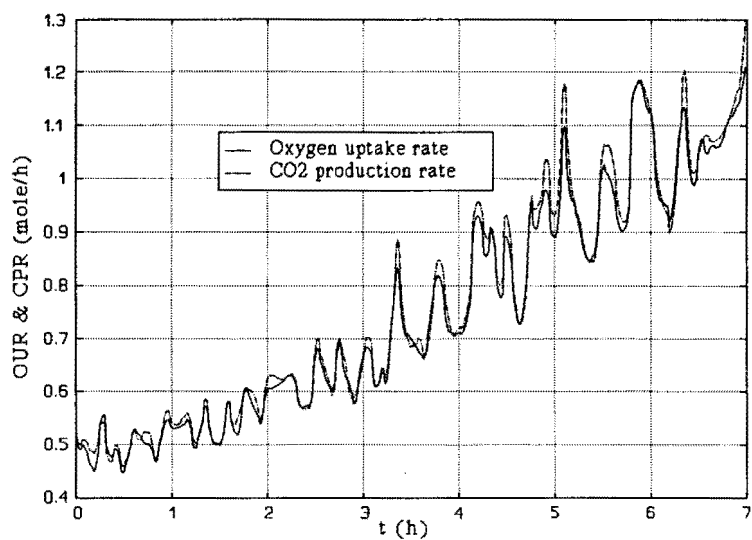
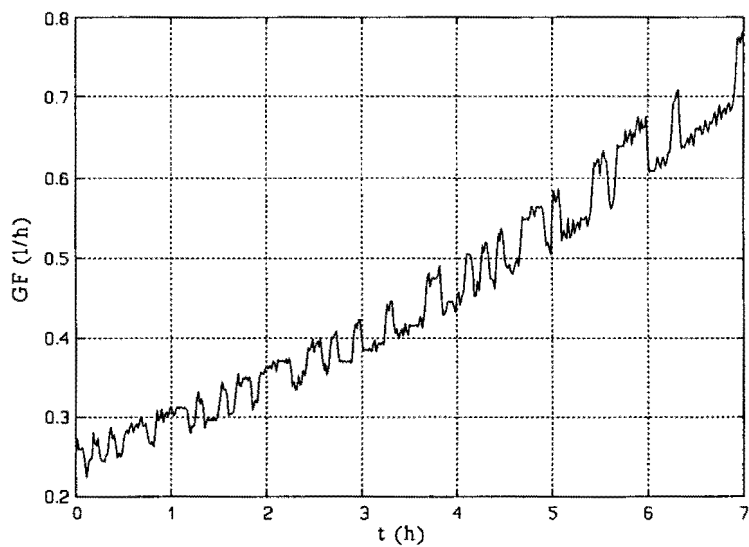figure A.11: *OUR* and *CPR* for identification experiment for predictor

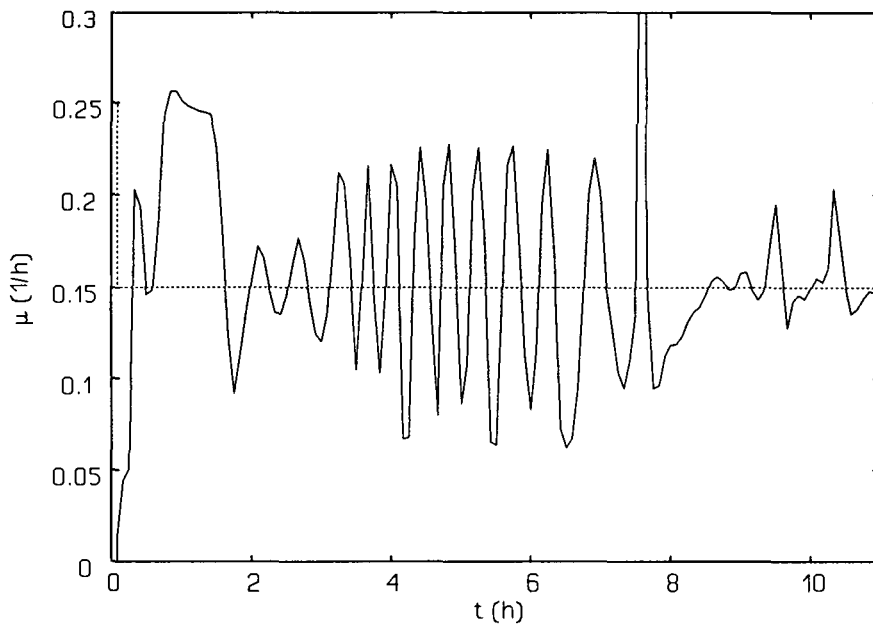figure A.12: *GF* for identification experiment for predictor
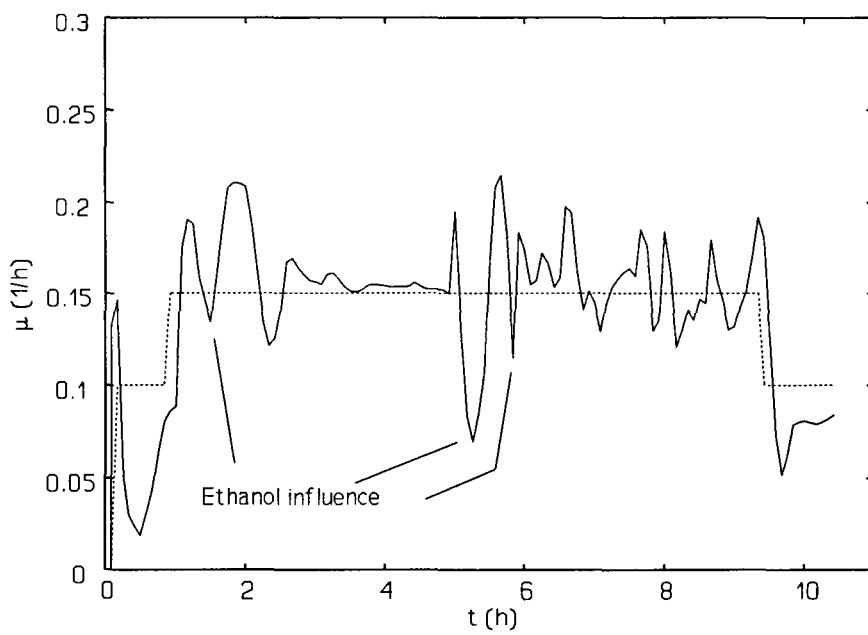
figure A.13: Too agressive PI tuning
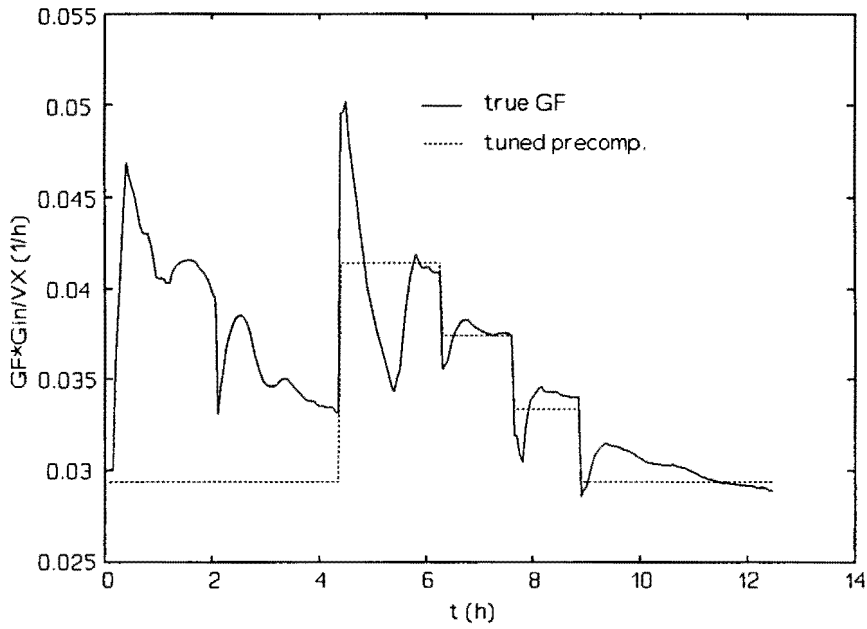

figure A.14: More sluggish PI tuning
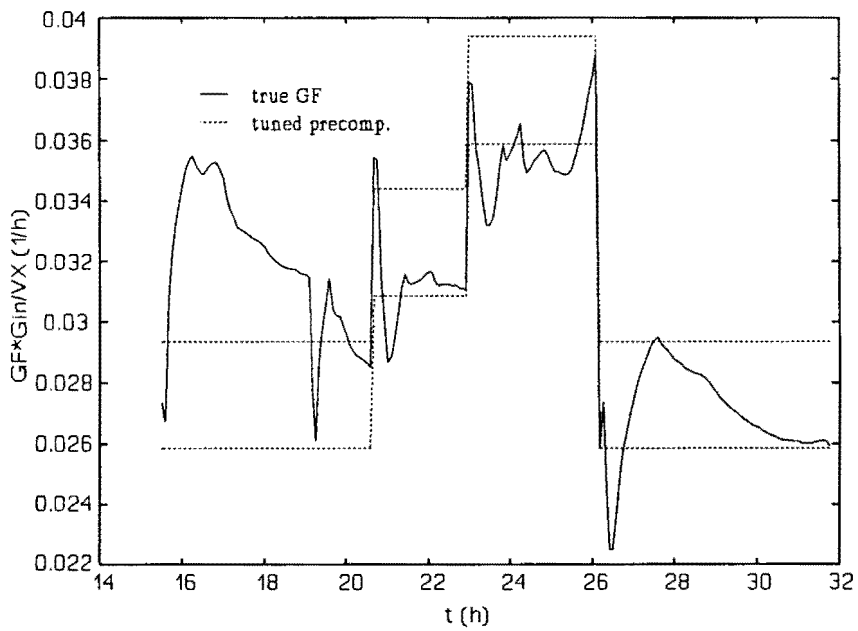
figure A.15: Precompensation tuning, first experiment

figure A.16: Precompensation tuning, second experiment

Appendix 6: Evaluation of the control method

In this appendix, the control structure is compared with other control techniques used for fed-batch fermentation. Also some inherent problems and future changes are discussed. The issues presented reflect my opinion instead of generally accepted facts. They are open for discussion, more than any other part of this report.

The implemented control uses no explicit adaptation scheme. Other control structures use adaptive schemes to cover for the different dynamics for different metabolic pathways. A major drawback of this approach is, that control is deteriorated immediately after a change in the used pathways: the adaptation requires several samples to converge to a new model. As the change can be detected using RQ, it is possible to eliminate this by changing the model used on a feed-forward basis. This is what the precompensations in our control structure try to do. Further, the on-line identification suffers from the fact, that the input signal of the process is not very rich. External disturbances are either non-stationary peaks or slowly changing offsets. These disturbances certainly do not excite the process persistently, so that the on-line identification results may be unreliable. For our particular case, there is another argument not to use adaptive schemes. The sampling interval of our equipment is about 5 minutes. This means, that it can take 2 hours and more before the adaptation algorithm has converged. It is unlikely however, that ethanol consumption or production occurs uninterruptedly for such a period of time. Consequently, these special cases will never be modelled accurately. The model for the normal situation, oxidative growth on glucose, only gets disturbed by occasional ethanol production or consumption.

The two concepts mentioned above could be mixed: a state recognition algorithm changes the models to be used on a feed-forward basis, an adaptation scheme tunes these models to the actual situation. In this case, it is again doubtful whether the models for ethanol production and consumption ever get a chance to converge.

The current control suffers from inaccurate ethanol estimates. A relatively large sensitivity to errors in these estimates can be expected for sub-optimal growth based on physical reasoning: a small amount of ethanol can make a big difference in the overall growth rate. This is probably best demonstrated at the end of the batch phase. When all ethanol has been consumed the oxygen consumption rate falls to a low level within a minute or so (this can be seen from the oxygen concentration in the broth). There is no gradual changeover when the ethanol concentration decreases. Other control techniques will have to cope with this problem as well.

The PI-controller used in the μ controller should eventually be replaced by a more sophisticated robust controller. This controller should be based on a model for the transfer from the input of the precompensation to the output of the process. According to the simulation model, this is a pure delay, but the experiments have shown, that other dynamics are involved in this transfer as well. This controller is a better place to account for these dynamics than the feed-forward dynamic precompensation used in our later experiments. One should be aware of possible complications in this approach however: it may interfere with the static precompensation used so far. Also the dynamics for an increase in μ differ from those involved for a decrease. The size and the starting point of the change may also be important. The dynamics have not been thoroughly identified yet. This has to be done before a new controller can be designed.

The current control does not aim for optimal growth. If optimal growth is desired, RQ-control and ethanol control should be considered, as these schemes are designed especially for optimal growth. They are not suited for sub-optimal growth, however. If both optimal and sub-optimal growth are desired, one should switch between techniques.