Report for CLIMA 2000, Brussels 1997, Aug. 31 - Sept. 2.

Control

Jürg Tödtli
Landis & Staefa
Zug, Switzerland

1. Introduction

This report summarises, explains and provides comments on the papers of technical session 5, “Control”, of the International Congress CLIMA 2000 to be held in Brussels from August 31 through September 2, 1997.

The objectives of this report are:

- to explain the subject so that it can also be understood by non-specialists
- to provide information that helps the reader or listener find the papers that he/she wants to study more closely
- to provide impulses to the discussion of the papers.

In order to get to a more unified description of the work described in the papers, the terminology used in this report may not always match the terminology of the original papers.

This report is structured as follows: Chapter 2 summarises the various papers and provides comments on the issues. Each paper is assigned to one of the following five topics: adaptive control, predictive control, various controllers, fault detection and diagnosis, and simulations. Chapter 3 provides some overviews, that are intended to help the reader identify the papers in which he/she may be interested as well as compare the work of the various papers. In addition, this chapter contains a few general comments and questions that may serve as inputs for some of the discussions at the technical session. The remainder of chapter 1 is dedicated to general explanations which may help the reader understand this report.

In this report, the terms *controller*, *fault detector,* etc. are used in a functional sense. This means that the terms stand for a *functional unit* of a control device or BEMS (building energy management system). Such a functional unit can be a piece of hardware or a piece of software. It can be an “atomic” functional unit or a compound of other functional units. A controller, for example, can be a compound of two or more smaller controllers plus some other functional units.

A functional unit has *inputs* and *outputs*. An output of a functional unit can be connected to one or more inputs of other functional units. The *connection lines* are considered signal flows, information flows or causal relations.

ON-OFF controllers, P controllers (proportional), and PI controllers (proportional plus integral), etc. are normally considered functional units characterised by a characteristic curve, a transfer function, a simple formula, or an algorithm. Oftentimes—primarily for more sophisticated controllers—it is sensible or even necessary to either regard controllers as active functional units that calculate the values of the output, or as decision making units that decide the value selected as output based on the current and previous inputs.

A few important terms of control theory are repeated in the following: The output of a controller is called the *control variable*. The system that is controlled is called the *process*. The process has two kinds of input variables: control variables and variables that do not come from the controller, the *disturbance inputs*. There are two typical cases where a controller controls a process to maintain a *process output variable* at a certain *setpoint*. For both cases, the difference between the setpoint and the process output variable is called the *control error*. In the first case—the *feedback controller*—the process output variable is measured and returned to the controller. In the second case—the *feedforward controller*—the process output variable is not returned to the controller, but instead, one or more disturbance inputs are measured and guided to the controller. In the first case we speak of a *closed loop* and, in the second case, of an *open loop* control system.

Many papers describe and investigate model-based control and monitoring methods. It is important to distinguish between two types of models, which in [IEA96a], are called *physical models* and *black-box models*. Both are mathematical models. The first model is derived from physics. Its internal variables and its parameters have a physical meaning. Examples of such parameters are thermal conductivity, flow resistance, and volume. Black-box models simply are mathematical equations, functions, etc., that normally are fitted to data. Typical examples are transfer functions, frequency response, regression models, and ANNs (artificial neural networks). Their parameters such as coefficients of a polynomial, regression coefficients, number of nodes, and weights of an ANN usually have no physical meaning. Some are even nonparametric models.

2. The papers

2.1 Adaptive control

A controller must be tuned in order to be adapted to the process that will be controlled. A PI controller, for example, is tuned by selecting values for the two controller parameters, the proportional gain, and the integral time. Or, for controllers incorporating a process model such as many feedforward controllers, it is necessary to set, i.e. to tune, the parameters of this model. In general, this type of tuning is carried out manually.

Adaptive techniques can be used to automatically adapt a controller to the process. If, on user demand, the controller is tuned automatically—an event that generally occurs during the commissioning phase or at the beginning of the operating phase—we speak of *automatic tuning* or *auto-tuning*. If the controller is continuously adjusted to adapt to changes in process dynamics, we speak of *adaptive control*. Adaptive controllers normally have a special functional unit, the *adaptation unit*, which carries out the adaptation. This unit typically receives the setpoint, the controller output (control variable), and the process output variable in the form of inputs and, in turn, supplies certain controller parameter values as outputs to the controller.

The desire to use adaptive control for HVAC processes is often motivated by the following situation. A nonadaptive feedback controller , e.g. a PI controller, may lead, depending on the operating point, to a very different control performance quality: good performance at or near the operating point at which the controller was tuned, oscillatory or sluggish control performance for other operating points. Oscillatory control performance means a badly dampened oscillatory response of the process output variable to a step change in the setpoint or disturbance input. Sluggish behaviour, on the other hand, means that the process output variable approaches the setpoint very slowly after a step change in the setpoint or disturbance input. A controller that is deliberately tuned so that the control performance will only be good or sluggish, but never oscillatory, is called a *conservatively tuned controller*. Adaptive controllers are designed so that their control performance is good for all operating conditions (as far as possible).

It is sometimes useful to consider why the control performance of a non-adaptive controller differs for different operating points. The reason lies in the non-linear behaviour of the process. As a consequence, the linearised process model that describes the behaviour of the process near an operating point is not constant. It varies with the operating point. The adaptation unit of an adaptive controller adapts the controller to this varying, linearised process model.

In addition to the three papers specifically assigned to this topic, the papers [Duburcq] and [Knabe] include descriptions of adaptive controllers.

The paper of Seem and Haugstad [Seem I]

Seem and Haugstad performed field and laboratory tests for a new type of adaptive PI controller.

The adaptive PI controller is a general-purpose adaptive controller, i.e. an adaptive controller for application in a broad class of processes. It primarily consists of a PI controller and an adaptation unit. The adaptation unit is based on pattern recognition. This unit, to a certain extent, operates similar to a human expert who knows how to modify the two controller parameters of a PI controller when oscillatory or sluggish control performance is recognised. The proposed adaptation unit is also capable of recognising the pattern ‘oscillatory performance’ or the pattern ‘sluggish performance’. Additionally, the adaptation unit is able to determine a measure of the amount of oscillation and a measure of the speed of response. Based on these two measures, the unit determines the two parameters of the PI controller. For a more detailed description of the adaptation unit, refer to an earlier paper of Seem.

The second part of the paper presents the field and laboratory tests; the hardware used for these tests is also described. Tests were performed for two control loops of a central air handling unit.

The first control loop is the duct static pressure control loop. The adaptive PI controller controls a variable speed fan to maintain the static pressure in the supply air duct at a specified setpoint. The paper also shows the static process characteristic which describes the relation between the controller output and the duct pressure in steady-state. Due to its nonlinearity and dependence on the openings of the dampers in the attached duct system, the gain of a linearised process model greatly depends on the operating conditions. This explains why a conservatively tuned nonadaptive controller may lead to very sluggish control behaviour under some operating conditions. The results of field and laboratory tests are shown in the paper.

Additional field tests were performed for the supply air temperature control loop. The adaptive PI controller controls the dampers and the cooling coil valve (in sequence) to maintain the supply air temperature at a specified setpoint.

Comments on the paper of Seem and Haugstad

- The laboratory and field test curves shown in the paper clearly illustrate how adaptive controllers adapt the controller parameters and how control performance is improved. People desiring an introduction to adaptive control should look to the curves.
- If there are no setpoint variations, how well does adaptation work with just natural disturbances? Is it necessary to artificially introduce setpoint variations just to make adaptation work well? (i.e., introduction of test signal)
- How does the proposed adaptive controller compare to other adaptive controllers based on pattern recognition?

- How many “hidden” parameters must be manually tuned if the adaptive controller is applied to a particular plant? Are they difficult to tune? Is it possible to find settings for these parameters for typical types of applications, as for example, the pressure control loop of a central air handling unit?

The paper of Federspiel [Federspiel]

Federspiel describes an adaptive controller specifically designed for air flow control in VAV boxes (VAV = variable air volume). This controller contains a nonadaptive feedback controller that controls the damper by supplying a damper position signal to maintain the air flow at a prescribed setpoint. This feedback controller is conservatively tuned in that it never leads to oscillatory—but very often to sluggish—control performance.

To cope with this sluggish behaviour, the feedback controller is not made adaptive, but instead, a feedforward controller is added which continuously calculates the damper position that would be necessary in steady-state to create an air flow which is equal to the setpoint. The calculation is based on a static physical model of the VAV box and the air duct branch to which the VAV box belongs. In order to calculate this damper position, the feedforward controller requires the current setpoint value of the air flow and the current value of the pressure drop across the respective air duct branch. Additionally, it must know the valve characteristic (for a 1 bar pressure drop across the valve) and, as a measure for the air duct branch resistance, the damper authority with respect to this duct branch. Instead of measuring the pressure drop across the air duct branch, the value is estimated by an estimator, which is based on the same static physical model as the feedforward controller. The estimator requires the air flow and the damper position as inputs and needs to know the valve characteristic and damper authority. Due to this estimation, the feedforward controller is called adaptive. In fact, Federspiel’s controller estimates the maximum air flow (i.e., the air flow which would be present if the damper were fully opened and the pressure drop across the air duct branch the same) instead of the pressure drop across the air duct branch; this, in principle, is the same. The damper authority can be received from the plant design or experimentally by a method also described in the paper.

If the damper position signal supplied by the adaptive feedforward controller were directly applied, the following typical situation would occur following abrupt changes in the setpoint or load disturbance: the airflow would initially deviate from the setpoint because the feedforward controller is based on a static only model, and the initially large deviation would quickly decrease. But, a small deviation would generally remain because: a) the model is not accurate, and b) the estimate of the pressure drop across the branch has an error, and c) the valve characteristic as well as the damper authority does not correspond to reality. It is now the task of the feedback controller, whose output is combined with that of the feedforward controller via a special functional unit, to reduce this small deviation. Because the feedback controller must reduce only small deviations, the fact that its control performance may be sluggish can be accepted. This approach to cope with the sluggish performance of a feedback controller, namely by adding an adaptive feedforward controller, differs greatly from the approach in the previous paper in which the feedback controller is made adaptive.

Due to the fact that the model—on which the feedforward controller is based— is a physical model, this controller, in contrast to the general purpose adaptive controller in the previous paper, is a special purpose adaptive controller.

The paper also presents results from computer simulations. They illustrate the effect of a wrongly selected damper authority. The controller seems to be quite robust with regard to such errors.

Comments on the paper of Federspiel

- Controllers–like the one of Federspiel–that contain a non-linear physical model are an interesting alternative to using adaptive linear controllers to control a non-linear process.
- It's useful that robustness has been studied. But, what about the robustness with respect to other parameters, for example, with respect to the damper characteristic?
- The controller contains a physical model. This offers the advantage that controller parameters identical to model parameters have a physical meaning. As a result, finding out which of these parameters should be treated as constant and which as varying parameters is easy.
- There is a controller parameter which has no physical meaning: the weighting factor that determines how the maximum air flow estimates should be smoothed. Is there a value applicable to all applications or is manual tuning required upon each application?
- How does the adaptive controller compare to alternative solutions:
    - Compensation of the nonlinear characteristic, at least partially (what is the conventional solution)?
    - Adaptive feedback controller?

The paper of Jadoenathmisier and van Paasen [Jadoenathmisier]

Jadoenathmisier and van Paasen considered a two pipe fan coil unit whose heating coil has an on/off valve. In order to reduce the variations in the supply air temperature, the authors propose to use cascade control requiring a supply air temperature sensor in addition to the room air temperature.

This paper describes a new type of adaptive controller for the cascade controller's inner loop. Because the valve is an on/off valve, the supply air temperature will oscillate. The design's goal for this secondary controller (inner loop controller) is to maintain the oscillatory supply air temperature within a tolerated range around the supply air setpoint (for example, defined as range from 3^0C below to 3^0C above the setpoint). The controller can operate in different modes. In the first mode, it operates as a simple relay control characterised by two temperature differences: a switch-off and a switch-on temperature difference (typically +0.5^{0}C and -0.5^{0}C). If the valve is open and the supply air temperature rises above the setpoint plus switch-off temperature difference, the valve closes. If afterwards, the supply air temperature drops below the setpoint plus switch-on temperature difference, the valve opens, etc. The secondary controller remains in this first mode as long as the supply air temperature stays within the tolerance range. But, if the range is exceeded, the controller is switched to a second mode which, again, is a relay control but with adapting switch-off and switch-on temperature differences. The adaptation works as follows: if the supply air temperature exceeds the tolerance range by a certain amount, the switch-off temperature difference is reduced by the exact same amount. If it drops below the tolerance range by a certain amount, the switch-on temperature difference is increased by the exact same amount. The controller will change to the third mode if the supply air temperature still cannot be kept within the tolerance range, etc. The paper illustrates the behaviour of the adaptive secondary controller by some results of field tests.

The controller introduced here for a cascade controller's inner loop for fan-coil units, in principle, can be used as a general purpose controller for a broad class of processes with two-position actuators.

Comments on the paper of Jadoenathmisier and van Paasen

- Comparing computer simulations of the room temperature control performance showing the advantages of a cascade controller with an additional supply air temperature sensor to a conventional controller would have been interesting.
- How robust is the adaptive controller? Is it necessary to manually adapt it when applied to another fan coil unit with a different size? (“Hidden” tuning?)
- If we consider the described secondary controller as general purpose adaptive controller for processes with two-position actuators: how does the controller compare to other controllers of this type?
- Adaptive controllers for processes with two-position actuators are attractive for two reasons: they need not cope with actuator nonlinearity and the step changes of the control signal usually sufficiently excite the process to receive sufficient information for continuous adaptation.

2.2 Predictive control

Predictive controllers are controllers that somehow account for future behaviour. The predictive controllers presented in this section belong to a specific class of predictive controllers that are based on optimal control in a deterministic sense. A controller of this class continually, periodically, or at selected time points, looks to the future for a specific time horizon, determines the control variable's optimal behaviour during this time horizon, and applies the determined behaviour until the next time point is reached when the controller again looks to the future. It supports the understanding, to think of the behaviour of the control or other variable over the mentioned time horizon as a “profile” of the variable. The time horizon is called the prediction or optimisation horizon. Predictive controllers usually are special purpose controllers.

The paper of Ren and Wright [Ren]

In this paper, a system consisting of an air handling unit and a building zone with a hollow core ventilated ceiling slab is considered. The supply air from the air handling unit first flows through the hollow core ceiling slab before entering the room. This way, the building's structure is used effectively as a thermal storage element. During the summer season—which is considered in this paper—cool outside air, for example, is used to lower the temperature of the slab at night, thereby allowing the slab to absorb heat during the day.

The overall aim of this work was to design a controller that would control the system so that energy costs would be minimised while still maintaining zone comfort requirements.

It is clearly sensible, if the controller, which determines how much the slab is cooled down at night, looks ahead to the next day and tries to use information indicating somehow the cooling load that must be expected.

This paper proposes a hierarchically structured controller with a high-level controller and several low-level controllers. The high-level controller determines the setpoint values for air flow and supply air temperature. These setpoints are guided to the low level where feedback controllers act directly on fans, chillers, cooling coil valves, heat recovery devices, etc. The paper focuses on summer operation of the high level controller.

The comfort requirements in summer are to maintain during the occupancy period from 0800 until 1600:
- a minimum ventilation rate of 2 air changes per hour
- a minimum level of thermal comfort, expressed as a limit of 10% PPD (predicted percentage of dissatisfied)

The high level controller is a predictive controller. It calculates every midnight an optimal profile for air flow and supply air temperature over the next 24 hours so that energy costs are minimised while maintaining comfort requirements. To perform this optimisation, the high-level controller receives information from the low-level controllers and the plant at midnight, but it acts as an open loop controller for the next 24 hours, when it supplies the calculated optimal control profiles as setpoint values to the low-level controllers.

In order to calculate the optimal control profiles, the controller requires a process model and must know the initial state (i.e., the state at midnight) as well as the profiles of the disturbance inputs as for example the outdoor temperature over the next 24 hours. Furthermore, the controller must know the comfort requirements and energy tariffs for this time period.

The process model proposed in the paper consists of a dynamic physical building model and a static physical plant model, which also icludes the low-level controllers. The initial state and disturbance input profiles are assumed to be known exactly for the work described in the paper. In reality, a state estimator and predictors that predict outdoor temperature, etc., 24 hours ahead of time would be required.

The paper first considers a relative exact optimisation. The 24 hour profiles of air flow and supply air temperature are parameterised by 24 values, i.e., one value per hour. This results in 48 unknown variables to be optimised. A genetic algorithm (GA) has been used to perform the optimisation. The incorporated optimisation problem of determining optimal low-level controller modes has been solved through exhaustive search.

This paper shows a few optimised profiles for the air flow and supply air temperature. These profiles impressively show how pre-cooling at night increases along with an increase in the day/night electricity tariff ratio. This must of course be expected.

Solving an optimisation problem with 48 unknown variables requires considerable computing capacity. Thus, the authors developed a simplified predictive controller (called “time-stage control”) as their proposal for real applications. The two profiles for air flow and supply air temperature are parametrised so that it results in 15 unknown variables instead of the earlier 48, and so that the sub-optimal control profiles have the main characteristics of the optimal control profiles.

Results shown in the paper illustrate that the control profiles calculated by the simplified predictive controller are similar to the “optimal” profiles. Energy costs, energy used, and comfort achieved with the simplified predictive controller were close to the “optimal” solution and considerably better than with a conventional controller switching to night ventilation if a few simple temperature conditions are fulfilled. The genetic algorithm (GA) used also in the simplified predictive controller proved to be efficient and robust.

Comments on the paper of Ren and Wright

- The paper provides first elements of a predictive controller. What remains to be developed is an estimator that estimates the initial state as well as predictors that are able to produce 24 hour forecasts of outside temperature, etc. Additionally, the question remains unanswered as to how much improvement can be achieved compared to the conventional controller if the state estimation and the predictions of the outside air etc. were not ideal but supplied by (real) estimators and predictors.
- Question: The controller incorporate a physical building and plant model. The model parameter values had to be set individually for each application. How is this done in reality? Who knows the values? Is it possible to design

an auto-tuner? Or is it possible to get these values efficiently by using information derived from the building's and HVAC plant’s design or from the component manufacturers? I will return to this question in section 3.2.

The paper of Kummert, André and Nicolas [Kummert]

The work of Kummert et al. is motivated by the desire to allow not the operator to find a good trade-off between comfort and energy consumption but to apply optimal control theory to best solve the trade-off. In this paper, this principle is applied to a simplified problem: the temperature control of a building zone with a heater and a cooler whereby both are assumed to work by convection. The control variable is the heating/cooling power transferred from the HVAC system to the zone. Converting this control variable to real control actions by means of a lower level controller such as a valve controller is the subject of future work.

The proposed controller is a predictive controller working similarly to the high level controller of the previous paper. It also calculates an optimal control profile of the control variable at midnight for the next 24 hours. Energy consumption, not costs, is minimised while maintaining the zone temperature within a defined, acceptable comfort range during the occupancy hours. At the same time, the aim is to prevent the zone temperature from exceeding specific absolute lower and upper limits at any time. The process model is a simple, dynamic, physical model of the building zone. No plant is considered in the model. The initial state of the process model (i.e., state at midnight) is determined by a state estimator of the Kalman filter type. The profiles of the input variables, for example the outdoor temperature, will be determined by an ANN based predictor (in development); but, for the computer simulations presented in the paper, they are assumed to be accurately known. The proposed controller also contains a feedback controller of the PID type which continuously modifies the pre-calculated control variable to maintain the process output variable near the pre-calculated profile.

It was possible to formulate the entire optimisation problem as a linear quadratic optimisation problem, to which a projected gradient method could be applied.

Results of computer simulations—performed with a simulation model more complex than the process model used in the predictive controller—are presented in the paper. A comparison to two conventional controllers shows improvements with regard to energy consumption, energy costs, and comfort.

Future announcements pertain to controller modifications to allow for operation with a moving (receding) horizon, i.e., calculation of the optimal profile is repeated hourly, but only the first hour of this profile is actually applied. But, this approach, which is typical of predictive control in automatic control theory, is not evaluated by means of computer simulations.

Comments on the paper of Kummert, André and Nicolas

- The main deficiency of the conventional controllers compared to the proposed controller is that comfort conditions are reached too early or too late, but not at the beginning of the occupancy period. The reason for this is a missing STO (start time optimisation) function. About 20 years ago, STO functions were introduced in commercial control devices and are quite common today. Why did the authors not use controllers with this function for their comparison? It would be interesting to see the advantages of the new method compared to conventional controllers with STO functions.
- The method which has been presented in the paper appears to go towards an integrated solution of what, today, is often solved via two separate functions: a STO and a heating curve. What are the advantages of this integrated solution? What would the advantages be if the authors were successful in their attempt to create an auto-tuning version of their method?
- And, the question asked at the end of the comments on the previous paper regarding model parameter determination applies again.

The paper of Duburcq and Guillerminet [Duburcq]

Duburcq and Guillerminet attempted to develop advanced control of intermittent heating. They considered buildings that are occupied only at daytime during weekdays; in this case, there is broad consensus on the fact that energy can be saved by intermittent heating, provided thermal inertia is not too high. The considered heating system is a hot water heating system with boilers and radiators without thermostatic valves. To keep it simple, we assume that there is a flow temperature controller (the flow temperature is the temperature of the water flowing to the radiators) which acts on a three-way mixing valve and that the proposed predictive controller supplies the flow temperature setpoint.

The proposed predictive controller contains an occupancy scheduler programmable by the operator. The predictive controller works, like many other heating controllers, alternatively in four different operating modes in the following sequence: the unoccupied period starts with mode 0 (numbering of the mode according to the paper) during which the heat distribution system is switched off; if the room temperature reaches a lower setpoint of for example 12 ^{0}C, the controller switches to mode 1 and maintains this setpoint; in the final phase of the unoccupied period, the controller is in mode 2 during which the room temperature is raised with almost maximum heat power to the upper setpoint of for example 19 ^{0}C; once this setpoint has been reached, the controller changes to mode 3 and maintains this setpoint to the end of the occupancy period.

The proposed controller uses predictive control to determine the time point at which it switches to mode 2, so that the room temperature reaches the setpoint of 19 ^{0}C as close as possible to the beginning of the occupancy period. The controller performs this by periodically predicting the time point at which the room temperature would reach the 19 ^{0}C setpoint if it immediately changed to mode 2. In fact, the controller will actually switch to mode 2 if the predicted time point coincides with the beginning of the occupancy period. This is, by the way, how most STO functions work in practice. However, it is more unusual that the proposed controller uses predictive control to keep the control performance close to the pre-calculated performance during mode 2 and to maintain setpoints of 12 ^{0}C and 19 ^{0}C during mode 1 and 3.

The internal process model used by the predictive controller is a simple transfer function model with the flow temperature setpoint and the outside temperature as the inputs and the room temperature as the output. Contrary to the process models used in the previous two papers, the model is not a physical but a black box model. The process model has 4 parameters: two describing how the room temperature depends on the flow temperature setpoint and two describing how it depends on the outside temperature. Two sets of such four parameters are used, one for weekdays and one for weekends. An adaptive unit was designed that automatically adapts the parameters that characterise how the room temperature depends on the flow temperature setpoint. The adaptation unit is based on a least square algorithm. The other parameters must be tuned manually. As a predictor for the outside temperature, a simple extrapolator is used. This is sufficient because the outside temperature will not vary greatly during the relatively short optimisation period.

Simulation studies were carried out in which the proposed predictive controller is compared to the two conventional controllers of which only the better one will be discussed in this report. This conventional controller features the same four control modes, and it also uses predictive control to determine the time point when mode 2 is activated. The difference, however, lies in how the controller predicts the time required to raise the room temperature from the current value to the setpoint value of 19 ^{0}C . The conventional controller uses a function in the mathematical sense (geometrically: a surface) that describes how this predicted heat-up time depends on the current room temperature and the current outside temperature. The function is nonadaptive, i.e., its parameters must be tuned manually. In addition, the conventional controller differs from the proposed controller in that no predictive control is used to keep the control performance close to the precalculation during mode 2 or to maintain the room temperature setpoint during mode 1 and 3. In fact, no feedback control is used during mode 1, 2, and 3, only feedforward control. The simulations show the superiority of the proposed predictive controller over the conventional controller in terms of comfort, energy consumption and robustness with respect to the manually tunable controller (incl. model) parameters. The authors concluded the following from this comparison: “For the heating of buildings occupied on a discontinuous basis, advanced control provides a more effective solution than the intermittent heating control devices currently available.”

Comments on the paper of Duburcq and Guillerminet

- There are three main deficiencies of the considered conventional controllers compared to the proposed predictive controller:

- Comfort conditions on Monday morning are normally reached too late (sometimes referred to as the *Monday morning effect*). The question arises as to why no conventional controllers were considered which do not lead to this Monday morning effect. At least two such controllers have already been introduced in the market. They provide STO functions based also on a mathematical function (surface) for the predicted heat-up time, but, in addition to the actual room and outside temperature, they contain a third independent variable: either the wall temperature (requiring an additional sensor) or the time where the controller is in mode 1.
- The room temperatures during the occupancy period are farther from the setpoint. Why were conventional controllers with room temperature feedback control during mode 3 not considered, as they are available on today's market?
- The conventional controller is less robust. This means that control performance is worse if the manually tuned parameters are not adapted to the actual system. Controllers that are available on the market, with, as the primary feature, an adaptive STO function and, as a secondary feature, a room temperature feedback control or adaptive heating curve would have performed much better. Why were these controllers not considered?

An answer to all the questions can perhaps be found in the above cited conclusion. The authors were not fully aware of the type of controllers available on the market. However, the critique should perhaps not only be addressed to the authors, but also to industrial companies who do not publish new solutions broadly enough.

The paper of Datta, Tassou and Marriott [Datta]

The three previous papers on predictive control did not contain any explanations as to how predictors—which would be required for real applications—work.

The paper of Datta et al. describes and compares predictors which, according to the authors, could be applied in certain predictive controllers, but also for other applications such as automatic diagnosis. They all predict the electricity consumption of a supermarket half an hour ahead of time. They all are based on a mathematical function describing how predictions for electricity consumption depend on time, on current or previous measurements of electricity consumption, and on current measurements of internal and external temperatures, and of internal and external humidity. The various predictors differ with regard to the measurements used. Furthermore, they differ in the overall approach: artificial neural networks (ANN) or nonlinear regression. Both approaches allow for a nonlinear mathematical function.

The authors applied the different predictors to measured data from a supermarket in the UK. One of the conclusions is that the selected type of ANN could better predict electricity consumption than the selected type of nonlinear regression method. Additionally, the authors state that less expertise is required to apply ANN techniques than traditional statistical techniques.

Comments on the paper of Datta, Tassou and Marriott

- Half hour ahead predictions of electricity consumption could be applied in a peak demand limiting (PDL) function, which can be interpreted as a predictive controller. With regard to such applications, however, other evaluation criteria, which are asymmetric, were of interest.
- 24 hour ahead predictions of electricity consumption as would be required to predict internal heat gain in predictive controllers of [Ren] or [Kummert] are not considered in the paper. But, similar approaches could be applied to develop such predictors, also to predict other input variables of the process model contained in the predictive controller.
- Many users of ANN have previously stated that less expertise is required to apply ANN techniques compared to traditional statistical techniques. But, is the question of which technique requires less expertise not rather a matter of experience?

2.3 Various controllers

The paper of Sauer and Utterson [Sauer]

Sauer and Utterson compared two control strategies for a VAV HVAC installation with respect to comfort criteria, in particular IAQ (indoor air quality), by carrying out field studies in a building at the University of Missouri in Rolla, United States. The first control strategy was a normal strategy. The second strategy applies measurements of the outside air flow and has an additional control loop that controls the dampers to maintain the outside air flow at a desired minimum value, provided no free-cooling (economizer cycle) is required. The following measurements were monitored for evaluation: indoor temperature, relative humidity, CO_2, VOC, particles, outdoor air flow, and “building pressure”. The authors conclude that the second control strategy is better with respect to comfort criteria, and, when taking into account the results of computer simulations of a previous publication, saves significant amounts of energy.

The paper of Knabe et al. [Knabe]

Knabe et al. describe a controller for residential buildings that controls a conventional hot water heating system with radiators, the fans of a local ventilation system and the motor-driven outside blinds to maintain the room temperature, the room humidity and the air exchange rate in a prescribed comfort region. The controller is hierarchically structured into two levels: The upper level contains a rule-based supervisor that determines the controller modes—for example “night mode” or “reduced temperature”. At the same time, it determines which of the lower level controllers is active as well as the respective setpoints. One of the controllers on the lower level simultaneously controls the radiator valves and the fan speed to maintain room temperature and room humidity at the setpoints issued by the supervisor. An adaptive controller is selected under the argument that the process is nonlinear and time-varying. Among the various adaptive controller types, a model reference adaptive (MRAC) controller was selected.

The results of some computer simulations are shown in the paper. The paper provides a detailed description of the simulation model of the room with its heating and ventilation equipment. The model was implemented in MATLAB/SIMULINK.

Comments on the paper of Knabe et al.

- The supervisor is rule-based. Rule-based controllers, no matter whether they are based on sharp logic rules as in the previous paper or on fuzzy-logic rules, have the potential to be self-documenting and easily understandable. An operator or service engineer can easily check if the controller is working correctly. But, it appears that this potential has not yet been exhausted.
- The rule-based supervisor comprise two sets of rules: one for winter and the other for summer. It is not clear how the system is controlled in spring and autumn, the most difficult periods in our experience. A presentation of the rules describing how the supervisor switches between winter and summer and vice versa, or special sets of rules for spring and autumn would have been interesting.
- The authors conclude that the controller is easy to commission and has a simple operator interface. These are important points. But unfortunately, these points are not explained in the paper. For example, which parameters must be set upon commissioning and what information is required?
- It seems that an unusual amount of sensors has been used for this controller (the paper does not provide information on the number and type of the required sensors). What is the cost-benefit ratio for these sensors?
- The model reference adaptive controller appears to be quite complex. It would be interesting to see a simulation result comparing its control performance with a robustly tuned conventional controller under the same conditions. This would help to assess the benefits of the MRAC controller. Does the MRAC controller lead to better control performance or is the MRAC controller easier to tune?

2.4 Fault detection and diagnosis

Six papers discuss the development and evaluation of FDD (fault detection and diagnosis) functions. The purpose of such functions is to automatically detect and diagnose faults or to support the operator or service engineer in detecting or diagnosing them. Such fault detection and diagnosis functions can be implemented in a control device, in a BEMS, or in computers connected to the control system. The kind of faults which may be detected and diagnosed are: faults in the HVAC plant or in the control system, but also, even though less often, faults in the building structure, operation or user behaviour. There are faults which occur at the design phase, the installation phase, the commissioning phase or in the operation phase. Some fault detection and diagnosis methods are developed specifically for application during the commissioning phase where the building is usually unoccupied so that special tests can be carried out. Others are specifically developed for continuous monitoring during the operation phase. There is a broad consensus today on how to use the terms fault detection and diagnosis. *Fault detection* means detecting that there is a fault. To further localise the fault or to find out the cause of a fault is called *fault diagnosis*.

There are several faults which, without a fault detector, would not be detected over a long period of time or even during the entire life of the plant because they do not perceptibly impact the indoor climate. The use of a fault detector, in this case, can result in reduced energy consumption, reduced energy costs, reduced wear, or, if a fault that potentially causes a breakdown can be detected and removed, the benefit is increased availability of the plant. Advantages of diagnostic functions usually lie in reduced maintenance and service costs.

A common principle of most fault detection and diagnosis methods is the comparison of an *actual behaviour* of a system with a *reference behaviour*. This comparison is carried out by the FDD functional unit, by the developing engineer during the design (or programming) of the FDD functional unit, or during the operation phase by the operator or service engineer supported by an FDD function or system. To identify where the principle is applied for a certain FDD function, helps to understand it. The actual behaviour can be characterised, for example, just by a set of measured variables, by performance indices, by characteristic curves and surfaces, or by parameter estimates. In the case of fault detectors, the reference behaviour normally represents the behaviour which would be expected were there no faults.

The first two papers, [Seem II] and [Visier], are quite pragmatic in their approach and prototypes are available. Field tests were performed or started. The following three papers, [Breuker], [Grob], and [Buswell], describe the concept of an FDD function and report on the work progress toward concept realisation. All three FFD functions are based on a static process model, i.e., a model describing the process behaviour in steady state. The last paper, [Soethout], is devoted to energy consumption monitoring of buildings. Most of this work has its origin in an international research project devoted to FDD [IEA96a] [IEA96b].

The paper of Seem, House and Monroe [Seem II]

Functional units for the control devices and the operator station of a BEMS that help the operator detect plant or control system faults during operation have been developed. The functional units of the control devices continuously calculate particular performance indices, as for example, the duty cycle of a motor, the number of starts, stops and reversals of an actuator, or the absolute value of a feedback controller's control error. The performance indices are then smoothed (by exponentially weighted moving averaging), before they are communicated to the operator station. Thus, the data are compressed on the control level which results in reduced data traffic on the BEMS communication network. At the operator station, the smoothed performance indices are visualised enabling the operator to quickly assess the performance of a large number of controllers.

Laboratory tests were carried out on different controllers of a VAV air handling unit and of a VAV box. Furthermore, the fault detection functional units have been implemented in over 100,000 digital controllers of VAV boxes in the field. The paper shows the results of a field test with 24 VAV boxes. It demonstrates in an impressive way how faults could be detected by using as performance indices the absolute values of the air flow control errors. These performance indices were considerably greater for the faulty VAV boxes than for the others. One fault was a defect capacitor and the other an incorrectly installed damper actuator. Before the faults had been detected, both the subcontractor and the building engineer were confident that the VAV boxes were operating properly.

The comparison between actual and reference behaviour is, in this case, not performed by automatic limit value checking but by the operator. He/she is supported by the operator station which visualises the actual behaviour in the form of a bar chart. The operator knows the reference behaviour from experience or he can derive it by looking to the performance indices of the majority of VAV boxes, which he assumes to operate correctly. If the performance indices are the smoothed absolute values of a control error, the reference behaviour is zero or “close to zero”.

Comments on the paper of Seem, House and Monroe

- Fault detection based on monitoring the absolute value of control errors and similar performance indices was introduced in commercial controllers a long time ago [IEA91]. What is new and an important scientific contribution to the knowledge about the benefit of such methods is the evaluation of laboratory and field tests as well as the detailed description of cases where faults have been detected.

The paper of Visier et al. [Visier]

A prototype was developed for an FDD tool allowing the responsible administration of a town to centrally monitor the hot water heating systems of a large number of school houses. The tool, connected to the BEMS, collects hourly means of the outdoor temperature, indoor temperature, and flow temperature for each heating circuit. A first functional unit of the FDD tool receives the hourly means and calculates daily values from them. For example, the daily means of the outdoor temperature, the indoor temperature two hours prior to occupancy, or the indoor temperature at the beginning of occupancy. In order to do this, the occupancy schedules, as programmed in the BEMS, must be communicated to the FDD tool. A second functional unit, which is rule-based, produces “diagnostic messages”, such as “under-heating during occupancy”, “boost too early”, or “uncomfortable at beginning of occupancy”. This diagnostic unit requires seven rules. The comparison between actual and reference behaviour is not carried out in the FDD tool, but instead, guided the rules' derivation. Once a week, the tool produces a synthesis report presenting the diagnostic messages on a daily basis for each heating circuit of each school building. The user can induce trend plots from the tool. This serves as a support to locate the cause of a fault.

Two measures were taken in an attempt to keep the tool simple: not the ultimate causes of faults are “diagnosed”, but the symptoms. Secondly, no attempt was made to reduce false alarms through sophisticated determination of thresholds or through a large number of additional rules. It is up to the user to recognise false alarms.

The work started with an analysis of the organisation where the tool was supposed to be used and of the knowledge of its various members. The paper includes pictures used at the man-machine-interface (MMI).

First field tests were carried out in the 1996/97 heating season in Montpellier, France.

Comments on the paper of Visier et al.

- The tool is interesting because of its simplicity. It will be interesting to see the results of the field tests. How about user acceptance? How many relevant faults will have been detected with the tool? It would be interesting to compare the costs, benefits, and user acceptance of this tool to a more complicated FDD tool designed to produce few false alarms and to diagnose causes of faults.
- The work is interesting due to the analysis of the organisation and the members' knowledge.

The paper of Breuker and Braun [Breuker]

This paper first describes the overall structure of an FDD function specifically designed for vapour compression equipment of rooftop units. Ten measurements are fed to the FDD functional unit. Seven of them characterise the actual system behaviour. The reference behaviour, i.e., what is expected in a faultfree case, is represented by the seven outputs of a static process model, which is driven by the other three measurements. Seven residuals, i.e., seven differences between actual and reference variables are calculated and fed to a fault detection classifier as well as to a fault diagnosis classifier; both of them issue output messages when a steady-state detector informs them that the system is in the “near steady-state”.

The fault detection classifier outputs “no fault detected” if all residuals are small. If at least one of the residuals is not small, the fault detection classifier reports “a fault detected” and the fault diagnosis classifier reports one of five possible faults, provided sufficient evidence exists. The five possible faults are “refrigerant leak”, “compressor valve leak”, “liquid restrictions”, “condenser fouling”, and “evaporator fouling”. The selection of one of five faults is conducted automatically by applying rules of the following type: “If the first residual is negative and the second residual is positive and .... then the fault is a ‘refrigerant leak’”. These rules were derived during the development of the fault diagnosis classifier by inverting a set of fault-symptom causal relations of the form “If there is the fault ‘refrigerant leak’ then the first residual is negative and the second residual is positive ....”. Probability analysis has been performed to find good threshold parameter values for the fault detection and fault diagnosis classifiers.

The paper describes a few laboratory tests carried out to continue the previously begun evaluations of the described FDD method. All five faults could be introduced in the laboratory rooftop unit via reproducible means. Because until today no process model has been developed, measurements of the laboratory rooftop unit in the faultfree case have been used as reference values. Noise was added to the measurements to account for model inaccuracy and measurement errors. The results presented in the paper show that the FDD method was able to detect and diagnose all five commonly occurring faults in vapour compression equipment before significant impact on equipment operation would be exerted.

Comments on the paper of Breuker and Braun

- It will be interesting to follow up on future progress of the work which will hopefully be reported in future papers:
  - What kind of process model will be selected and developed: a purely physical model or a mixture between physical and black-box model? Is it possible to achieve the accuracy as assumed for the evaluation described in this paper? What are the costs to adapt the model to a specific type of rooftop unit?
  - The laboratory tests were conducted at almost ideal steady-state conditions. But, what if the FDD method is applied in real operation? In

    this case, fault detection and diagnosis are performed on a system in near steady-state detected by a steady-state detector. Is it necessary to retune the parameters of the fault detection and fault diagnosis classifiers?
  - How robust is the FDD method when applied to different samples of the same type of rooftop units?
- Ten sensors are used for FDD. This is much more than what is required for control. By how much is the risk of faults in the overall system increased through the additional sensors which themselves can be faulty? Is there a reasonable cost-benefit ratio for the additional sensors?

The paper of Grob and Madjidi [Grob]

In the first part of the paper, Grob and Madjidi describe an approach and a corresponding tool which, in the commissioning phase, allows to ensure that an HVAC plant is operating as intended during design. Contrary to current commissioning practice, commissioning is not completely carried out on-site. What remains to be done on-site is to measure a few steady-state operating points for each component allowing to calculate the parameters of characteristic curves and surfaces. The rest can be done in the consulting engineer's office or at the BEMS manufacturer’s office. Many steady-state operating points representing “actual behaviour” can quickly be calculated on the basis of characteristic curves and surfaces. The corresponding steady-state operating points representing the reference behaviour are calculated by applying a physical plant model. The parameter values of the physical plant model were taken from the HVAC plant's design as well as from data of the HVAC component manufacturer. The residuals, i.e., the differences between the “actual” and the “reference” steady-state operating points are fed to a fault detection and to a fault diagnosis unit. These two units work similar to those used in the FDD method of Breuker and Braun. This approach will allow for testing the system across the entire operating range.

The second part of the paper describes the work performed to test and evaluate parts of the approach and the tool. Measurement series of a VAV air handling unit in a faultfree as well as in some faulty cases— produced for a common exercise among participants of the IEA ECB Annex 25 project [IEA96a] [IEA96b]— were used to check if the tool could detect and diagnose the faults. By applying the measurement series of the faultfree case as the “reference behaviour”, the physical model and its parameters were excluded from the tests and evaluations.

Comments on the paper of Grob and Madjidi

- The authors claim that this approach even serves to check sophisticated control strategies. In order to do that, the actual behaviour would have to be measured on-site (from how I understand the approach).
- An application of the method would require a drastic change of current practice in engineering and commissioning (the approach thus appears very futuristic).
- The approach is based on the assumption that the parameters of the physical model come from design information and manufacturers' data. I will return to this point in section 3.2.
- Generally, applying this approach to recommission old plants will not be possible due to a lack of design information and manufacturers' data.
- The authors claim that not much additional work is needed for this new commissioning approach (compared to traditional commissioning). Is it possible to develop the tool and to organise the process so far that this is in fact true?

The paper of Buswell, Haves and Salsbury [Buswell]

The first part of the paper describes an integrated approach and tool for performance validation of HVAC systems during the commissioning phase as well as condition monitoring during the operating phase. The tool will first be used in the commissioning phase for some on-site tests and, based on the resulting measurements, it estimates the parameters of a static physical model. The resulting parameter estimates, which represent the actual behaviour, are compared to the parameter values characterising the faultfree reference behaviour. These reference values are derived, as is the case in the method of Grob and Madjidi, from design information and the manufacturers' data. The residuals, i.e., the difference between actual and reference values of the parameters, are used by the fault detection unit. Faults detected in this way are assumed to stem from equipment not meeting the design specification, from incorrect installation or from inadequate commissioning. The tool also contains a fault diagnosis unit, based on the assumption that specific abnormal values of the parameter estimates can be associated with the presence of particular faults. After the service engineer has removed (as far as possible) the detected and diagnosed faults, the tool will again be used on-site to repeat the parameter estimation. These new parameter estimates are used by the tool as the basis for reference behaviour during the operating phase. This second parameter estimation is carried out because there will always remain a difference between an even well commissioned plant and what is expected and calculated at the planning phase, due to uncertain knowledge during the plant design phase.

The second part of the paper describes preliminary experiments carried out at a laboratory air handling unit. It shows, for example, the measured temperature behaviour of a heating coil in steady-state compared to the behaviour as calculated by applying a physical model whose parameters stem from design information and manufacturer's data.

Comments on the paper of Buswell, Haves and Salsbury

- There is no doubt that system quality could generally be greatly improved if the described method were applied. But, developing the method and tool to a state where in can be applied at reasonable costs (so that not a research project must be started for each application as it may appear to be in some cases!) is very challenging.
- Again, the approach is based, as is the case for Grob and Madjidi, on the assumption that the parameters of the physical model come from design information and manufacturers' data.
- Among the three previously discussed papers describing a model-based fault detection approach, this approach is the only one where a model was used in the experimental part.
- The experimental results are interesting because they provide understanding of the real behaviour of HVAC plants as well as of the accuracy of models and data used to design HVAC plants. It appears that the research carried out to realise the described integrated approach is very useful, and would be meaningful even if the very ambitious and perhaps futuristic goal were never reached.

The paper of Soethout, Honsclaar and Peitsman [Soethout]

The goal of the work described in this paper is to develop a method that allows to identify the households in a municipality which have potential for energy savings. Until now, this has been restricted to just gas consumption of residential buildings. A statistical model for the average specific gas consumption of a household (in m^3 gas per year and m^3 volume of room) and its standard deviation— in dependence of the type of ownership (rented or owned), building type (single family, multi-family, apartment), building position (4 classes), year of construction (6 classes), and number of inhabitants— has been derived. The data from over 16,000 buildings in Schiedam, Netherlands, was used to calculate the parameters of this model. In order to conclude if an individual household has a gas savings potential, the model is used to calculate the average specific gas consumption and its standard deviation for the specific type of household. If the actual specific gas consumption is greater than the sum of the calculated average value and the standard deviation, then a savings potential is identified.

Comments on the paper of Soethout, Honsclaar and Peitsman

- Assuming that the values for gas consumption of the households of a specific type are normally distributed, one can determine if an individual household belongs to the worst 16% of its type by applying this method. What are the consequences? The percentage of households with an energy savings potential is the same for each category—a fact that most likely does not correspond with reality. And, assuming that the model fitting procedure is repeated from time to time, the percentage of households with a savings potential will always remain at 16% even if the energy savings campaign is successful and an increasing number of households will reduce gas consumption. Is it not better to determine the reference values for a specific household type on the basis of households with a proven low energy consumption as is the case, for example, in Switzerland?

2.5 Simulations

The paper of Ahmed, Mitchell and Klein [Ahmed]

This paper describes a simulator for use in design and evaluation of VAV HVAC systems for laboratories. HVAC systems for laboratories differ from those for commercial buildings in that safety requirements, usually in terms of pressure differences between spaces, are added to the comfort requirements. Because the dynamic interaction between laboratory and HVAC system and control must be considered to evaluate comfort, energy consumption, and fulfilment of safety requirements, the simulator model includes laboratory envelope, air flow system, heating and cooling coils, and control functional units. As a special feature, the simulator model considers the compressibility of air allowing for accurately simulating the dynamic response of pressure control loops in addition to the dynamic response of the temperature control loops.

The second part of the paper illustrates the simulator capabilities through the following examples: pressure control, temperature control during cooling, and temperature control during heating. These simulation examples also demonstrate the use of two recently developed functional units for building management systems. The first one calculates the setpoint for the airflow control loop. It contains a static physical model that also considers the compressibility of air: the setpoint predictor (setpoint calculator would probably be a better term). This unit is a feedforward controller. The second unit continuously estimates the current cooling load and is called the load predictor (load estimator would probably be a better term). The simulation results show that a controller with these functional units is capable of maintaining the laboratory pressure and temperature within the desired limits. Also, this controller reportedly simplifies tuning.

Comments on the paper of Ahmed, Mitchell and Klein

- It has been shown that the simulator fulfils the technical requirements. It would be interesting to evaluate the simulator's application to design HVAC systems of laboratories in terms of costs and personnel qualification requirements.

The paper of Murphy, Rémond and Déqué [Murphy]

Murphy et al. report on a extension to the previously developed CLIMA 2000 building simulation software. A connection to the general mathematical program MATLAB was developed. The connection with MATLAB allows the CLIMA 2000 user to access many general evaluation and controller design tools. Such an environment is aimed at being used in research laboratories and in development departments of BEMS manufacturers.

The second part of the paper illustrates the potential of this connection with the aid of two room temperature control examples. In the first example, a fuzzy control tool was linked to the simulator allowing to evaluate the control performance of a fuzzy room temperature controller. The second example shows the computer simulation results for control performance with a predictive room temperature controller enabled via a link to a generic model predictive control toolbox.

Comments on the paper of Murphy, Rémond and Déqué

- How does the extension compare to a building and HVAC simulation environment which is directly implemented in MATLAB/SIMULINK and which, as a result, offers access to all the mentioned tools?

The paper of Jandon et al. [Jandon]

This paper describes a method to assess BEMS. The central tool is a simulator allowing to simulate a building's and an HVAC plant's thermal behaviour and allowing to connect a real BEMS to the simulator (this type of simulator usually is called an emulator). This way, evaluation of closed loop performance of a BEMS or of a BEMS's control function is possible. The assessment method as described in the paper is intended to be applied by laboratories or consultants. They will use the method to provide a service to either BEMS manufacturers by helping them check and improve the quality of their products, or to customers of BEMS manufacturers by providing them with “objective” information on a BEMS.

At the end of the paper, three simulation examples, carried out with the simulator, are presented: a) zone temperature control in heating mode and in cooling mode (the temperature control accuracy classification scheme as proposed in CEN regulation draft is applied (CEN = Comité Européen de Normalisation)), and b) the optimal start controller, and c) load scheduling as a part of peak demand limiting.

Comments on the paper of Jandon et al.

- What is new in this paper is not the idea of building a simulator to which real control equipment can be connected in order to evaluate the closed loop behaviour. Such simulators were developed and used in the development departments of BEMS manufacturers more than 10 years ago. What is new is the assessment method to be performed by laboratories or consultants for BEMS manufactures and customers.
- In the case where the method is used to supply customers of BEMS manufacturers with “objective” information on a BEMS, a transparent description of the test procedures and the simulation models is very important to remain open to honest assessments and constructive criticism.
- According to the paper, the user-friendliness of BEMS was also assessed. Indeed, user-friendliness is an important criterion. Consequently, it would be interesting to know the specific criteria used to conduct the assessment.

3. Final issues

3.1 Overviews

The following overviews are intended to help the reader identify the papers in which he/she may be interested and to compare the different papers.

All papers

	Focus of Paper						Building Type				HVAC				
	adaptive control .	predictive controller .	other controllers .	FDD .	simulation/emulation	assigned to section .	residential building .	commercial building .	school .	laboratory .	hot water heating .	AHU .	VAV .	fan-coil unit .	chiller .
[Seem I]	x					2.1						x	x		
[Federspiel]	x												x		
[Jadoenathmisier]	x													x	
[Ren]		x				2.2		x				x	x		
[Kummert]		x						x			(x)	(x)	x		
[Duburcq]	x	x						x			x				
[Datta]		(x)		(x)				x							
[Sauer]			x			2.3							x		
[Knabe]	x		x				x				x				
[Seem II]				x		2.4						x	x		
[Visier]				x					x		x				
[Breuker]				x											x
[Grob]				x								x	x		
[Buswell]				x								x			
[Soethout]				x			x								
[Ahmed]			(x)		x	2.5				x			x		
[Murphy]					x		x								
[Jandon]					x			x						x	

Controllers and FDD functions

	Function Type		How general?		Model based?			Model Type		Model Type		Evaluated by		
	controller .	FDD function .	gerneral purpose .	special purpose .	no .	model based design .	incorporated model .	black-box model .	physical model .	static .	dynamic .	computer simulations .	laboratory tests .	field tests .
[Seem I]	x		x		x								x	x
[Federspiel]	x			x			x		x	x		x	x	
[Jadoenathmisier]	x		(x)	x	x									x
[Ren]	x			x			x		x	x	x	x		
[Kummert]	x			x			x		x		x	x		
[Duburcq]	x			x			x	x			x	x		
[Datta]	(x)	(x)		x			x	x			"x"	(x)		(x)
[Sauer]	x			x	x							(x)		x
[Knabe]	x			x	x							x		
[Knabe] MRAC	x			x		x	x	x			x			
[Seem II]		x	x		x								x	x
[Visier]		x		x	x									x
[Breuker]		x		x			x	?	?	x			x	
[Grob]		x		x			x	x	x	x		x		
[Buswell]		x		x			x		x	x	?		x	
[Soethout]		x		x			x	x		x				(x)
[Ahmed] controller	x			x					x	x				

Controllers

	Information Flow			Adaptive Techn.		For Predictive Controllers			Various		Various Technologies			
	feedback (closed loop) .	feedforward (open loop) .	predictive control .	auto-tuning .	adaptive control .	based on optimal control .	predictor required .	predictor developed .	comparison of controllers .	hierarchical control .	rule-based, crisp logic .	rule-based, fuzzy logic .	ANN .	pattern recognition .
[Seem I]	x				x									x
[Federspiel]		x			x									
[Jadoenathmisier]	x				x						x			
[Ren]	(x)		x			x	x		x	x				
[Kummert]	x		x			x	x		x	x				
[Duburcq]	(x)		x	x	x	x	x	x	x	x				
[Datta]			(x)				x	x					x	
[Sauer]	x								x					
[Knabe]	x									x	x			
[Knabe]MRAC	x				x									
[Ahmed]controller	x	x												
[Murphy]example			x			x			x			x		

FDD functions

	Purpose		At hich phase?		Method Type			How far to appl.?	
	condition monitoring .	energy consumption monitoring .	at commissioning time .	at operation time .	quantitative .	qualitative .	contains rule-based elements .	close (pragmatic approach) .	far (futuristic approach?) .
[Seem II]	x			x		x		x	
[Visier]	x			x		x	x	x	
[Breuker]	x			x	x		x		x
[Grob]	x		x		x		x		x
[Buswell]	x		x	x	x				x
[Soethout]		x		x	x			x	

3.2 Three discussion topics

Three topics are discussed in this section which may provide impulses to the conference discussion. We start with a more specific topic and proceed in two steps to more general topics. The presentation of each topic is structured into explanations, questions, and finally statements reflecting the author's personal opinion.

1st topic: Physical models.

Explanations:

Six of the papers describe a controller or FDD function containing a physical model: two in feedforward controllers [Federspiel] [Ahmed], two in predictive controllers [Ren] [Kummert], and two in FDD functions [Grob] [Buswell]. Additionally, one paper [Breuker] may contain a physical model: but the authors were not clear on that issue. Physical models are characterised by parameters with physical meaning.

Questions:

- Is there a chance for a broader application of physical models in practice?
- Is there a chance that the values of the parameters with a physical meaning can be determined with the necessary accuracy and with reasonable effort in practice?

To discuss this question, distinguishing between two application cases may prove helpful:

- *Application case 1: Series-produced components or plants.*
  A large number of identical HVAC plant components or plants with integrated control and perhaps an integrated FDD function are produced. The physical model and its parameter values have to be determined only for one component or plant and can be copied for the others. This is conducted in the R&D department of the respective industrial company. The costs can be distributed across many components or plants. A typical example is the FDD function for the roof top unit in the paper by Breuker and Braun [Breuker].
- *Application case 2: Customised systems.*
  The controller or FDD function is applied to a customised system with a standard structure for which a physical model is available. The parameter values of the physical model must be individually determined for each system during the engineering or commissioning phase. The costs cannot be shared. Typical cases are described in [Ren], [Kummert], [Grob] and [Buswell].

In the 1st application case the questions are directed at R&D engineers of industrial companies. For the 2nd application case the question addresses those conference participants who are involved in the engineering and commissioning process of controllers and FDD functions; or, if the parameter values are obtained from design information or manufacturers' data, the question is directed at those involved in the plant design process and at component manufacturers.

Statements:

- The introduction of physical models in commercial products will be faster for the 1st than for the 2nd application case.
- With respect to the 2nd application case, research should be conducted in two directions:
  - Research in the first direction—with a shorter time horizon—should focus on using physical models with easily available parameters.
  - Research in the second direction—with a long time horizon and bearing greater risk—should focus on solutions requiring new information flows from the plant design process and component manufacturers to the engineering and commissioning of controllers and FDD functions; perhaps by using a common database throughout the building life cycle.

2nd topic: The tuning problem.

Explanations:

The topic is the tuning problem in a broad sense, i.e., the setting of all controller or FDD function parameters. For example: parameters of P and PI controllers, threshold values in FDD functions, parameters of black-box and physical models in model-based methods, weights in cost functions of predictive controllers based on optimal control, etc.. Even adaptive and auto-tuning methods have parameters that must be tuned prior to starting, for example, parameters of reference models in model reference adaptive control. When reading the various papers, the question as to how to choose a parameter of a controller or FDD function often surfaced and remained unanswered.

Questions:

- Can the tuning problem be left to those who really apply the controller or FDD functions, e.g., to the commissioning engineer?
- If not: how can research contribute to solving this problem?

Distinguishing between the two application cases defined above will again be meaningful.

Statements:

- More research work should be devoted to the tuning problem in the 2nd application case. Possibilities are:
  - The tuning problem is solved at the end of controller or FDD function development:

        - Developing tuning rules or determining fixed parameter settings ("factory settings") covering a large range of applications.
        - Conducting sensitivity analyses (robustness) for the parameters received this way (examples are provided in [Federspiel], [Breuker] and [Duburcq]).

    - Auto-tuning or adaptive controllers and FDD functions are developed (examples in sections 2.1). But also auto-tuners and adaptive controllers have parameters that must be selected. It should not be omitted to tell how these parameters should be selected (no hidden tuning).
    - Instead of first developing a controller or FDD function and then asking how the parameters can be selected in practice, the desire to solve the tuning problem should guide the entire development (from the beginning on). Examples:
        - The use of physical models may simplify the problem of parameter selection (examples are mentioned above in the subsection on the 1st topic).
        - Optimal control based on a physical model and cost function. This will help to understand the trade-offs between conflicting goals made by tuning the controllers. (Typical examples are in [Ren] and [Kummert].)
        - Methods with no (or only a few) difficult-to-tune parameters. (A typical example is the fault detector method in [Seem II].)
- To summarise the situation for the 2nd application case: a controller for which it is known how to tune it in practice is better than a controller which would have better control performance if it were tuned correctly, but for which it is not known how a service engineer can perform this tuning with reasonable effort.

3rd topic: What about the preference for simple understandable solutions? Or the acceptance problem.

Explanations:

Several sophisticated and complex control functions have been introduced in commercial products. But, it appears that most service engineers and plant operators prefer simple, understandable controllers. The following are the reasons stated in favour of simple, understandable controllers:

- Service engineers know where to successfully apply the controllers.
- They know where to install the sensors.
- They know, for example from experience, how to tune the controllers.
- Plant operators and other building occupants know how to manually operate the controllers.
- Service engineers may be able to explain to end users unusual control behaviour and are able to judge whether that behaviour is caused by a fault.

- They know where re-tuning is useful and how to perform it.
- They understand the interaction of control loops with other parts of the system.

This list of reasons also reflects why sophisticated and complex functions are often not accepted. The tuning problem discussed in the 2nd topic represents just one but certainly an important factor of the acceptance problem.

Questions:

- How widespread is this preference for simple, understandable solutions?
- What is a simple, understandable solution? Or better, what is a simple and what is an understandable solution?
- Is there a potential risk that most control and FDD functions developed by researchers will not be used in practice?
- If yes: what must be done?
- Should the problem be left to those writing documentation? (They should explain the controllers and the FDD functions so that service engineers, plant operators and other users do understand and accept them!)

Statements:

- The risk is real that control and FDD functions developed in research will not be used in practice. And admittedly, a simple, understandable controller that is applied and operated correctly is better in terms of comfort and energy costs than a sophisticated and complex controller applied and operated incorrectly. To counteract these risks, research should proceed in two directions; and both represent a challenge (also for researchers!):
    - Development and evaluation of simple and understandable algorithms. There is still a potential for such solutions.
    - There are different possibilities to increase acceptability of sophisticated and complex algorithms:
        - Develop good, if necessary user-group-specific, explanations of the control and FDD functions. Although the service engineers' and plant operators' understanding is generally based on causal explanations (i. e., cause-effect relations), goal oriented explanations of complex sub-functions would in some cases enhance understanding.
        - Develop concepts for operator interfaces that support understanding.
        - As basis for developing an explanation or an operator interface it will often be useful to explicitly design a user’s model, i.e., a mental model that helps the user understand the function.
        - Instead of first developing a controller or FDD function and then asking how to explain it and how to create a good operator interface, the desire to produce an understandable and easy-to-use function should guide the entire development process from the beginning on. This can mean, for example, that the user’s model is designed together with the algorithm.

- To be more general, development of control or FDD functions should be guided (at least for the 2^{nd} application case) on the one hand, by a stronger process-oriented viewpoint – or in other words, a life-cycle viewpoint – including engineering, commissioning, tuning, operation, and service, instead of a purely functional viewpoint. On the other hand, it should be guided by all criteria applicable to a good controller or FDD function including the following in addition to ‘high comfort for little energy’: ‘easy to engineer’, ‘easy to tune’, ‘easy to manually operate’, ‘easy to maintain’ and ‘easy to explain’. The papers of this session describe many attempts and solutions moving in the directions indicated by these suggestions. But, as a final point, increased effort is still required.

References

Papers of CLIMA 2000, 1997, technical session 5 on control and management:

[Seem I]: J. E. Seem and H. J. Haugstad, "Field and laboratory results for a new pattern recognition adaptive controller".
[Federspiel]: C. C. Federspiel, "Damper authority estimation and adaptive flow control".
[Jadoenathmisier]: W. I. A. Jadoenathmisier and A. H. C. van Paasen, "Rule based control of fan coil units".

[Ren]: M. J. Ren and J. A. Wright, "Optimal control of fabric thermal storage systems".
[Kummert]: M. Kummert, P. André and J. Nicolas, "Optimised thermal zone controller for integration within a building energy management system".
[Duburcq]: S. Duburcq and S. Guillerminet, "Advanced control for intermittent heating".
[Datta]: D. Datta and S. A Tassou, "Application of a neural network for the prediction of the energy consumption in a supermarket".

[Sauer]: H. J. Sauer and E. G. Utterson, "Instrumentation and data acquisition for monitoring effect of system operation on indoor air quality".
[Knabe]: G. Knabe, H. Werdin, M. Klingner, G. Lauckner, G. Billerbeck and J. Kunze, "Demand guided control of the indoor air quality in domestic buildings".

[Seem II]: J. E. Seem, J. M. House and R. H. Monroe, "On-line monitoring and fault detection of control system performance".

[Visier]: J. C. Visier, X. Li, P. Coralles, M. Irigoin, I. Le Vannier, J. Lovetri, M. Le Men and P. Picard, "Fault detection and diagnosis tool for schools heating systems".

[Breuker]: M. S. Breuker and J. E. Braun, "Demonstration of a statistical, rule-based fault detection and diagnostic method on a rooftop air conditioning unit".

[Grob]: R. F. Grob and M. Madjidi, "Commissioning and fault detection of HVAC systems by using simulation models based on characteristic curves".

[Buswell]: R. Buswell, P. Haves and T. I. Salsbury, "An integrated approach to HVAC testing and condition monitoring based on first principles models".

[Soethout]: L. L. Soethout, B. C. Honselaar and H. C. Peitsman, "An analysis and diagnosis method for the energy consumption in existing residential buildings on municipal level".

[Ahmed], O. Ahmed, J. W. Mitchell and S. A. Klein, "Development of a simulator for laboratory HVAC system".

[Murphy], K. M. Murphy, B. Rémond and F. Déqué, "Controller design for climate control in buildings using CLIM 2000".

[Jandon], M. Jandon, H. Vaezi Nejad, J. C. Visier, B. Clemençon, L. Halleur and O. Lusson, "Assessing BEMS systems".

Other references:

[IEA96a]: J. Hyvärinen (ed.), “Building Optimization and Fault Diagnosis Source Book”, International Energy Agency (IEA), Energy Conservation in Buildings and Community Systems Programme. Annex 25, Real Time Simulation of HVAV Systems for Building Optimization, Fault Detection and Diagnosis. Espoo: VTT Building Technology, 1996.

[IEA96b]: J. Hyvärinen (ed.), “Technical Papers of IEA Annex 25”, International Energy Agency (IEA), Energy Conservation in Buildings and Community Systems Programme. Annex 25, Real Time Simulation of HVAV Systems for Building Optimization, Fault Detection and Diagnosis. Espoo: VTT Building Technology, 1996.

[IEA91]: P. Sprecher, “Fault detection in commercial products for control of HVAC systems”, IEA Annex 25 Meeting, October 28 and 29, 1991, Delft, Netherlands.