Article:

Rockett, P. and Hathway, E.A. (2017) Model-predictive control for non-domestic buildings: a critical review and prospects. Building Research and Information , 45 (5). pp. 556-571. ISSN 0961-3218

https://doi.org/10.1080/09613218.2016.1139885

Model-Predictive Control for Non-Domestic Buildings: A Critical Review and Prospects

Peter Rockett[a] & Elizabeth Abigail Hathway[b]

[a]Dept. of Electronic and Electrical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK; [b]Dept. of Civil and Structural Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK

Abstract

Model-predictive control (MPC) has recently excited a great deal of interest as a new control paradigm for non-domestic buildings. Since it is based on the notion of optimisation, MPC is, in principle, well-placed to deliver significant energy savings and reduction in carbon emissions compared to existing rule-based control systems. In this paper, we critically review the prospects for buildings MPC and, in particular, the central role of the predictive mathematical model that lies at its heart; our clear emphasis is on practical implementation rather than control-theoretic aspects, and covers the role of occupants as well as the form of the predictive model. The most appropriate structure for such a model is still an open question, which we discuss alongside the development of the initial model, and the process of updating the model during the building's operational life. The importance of sensor placement is highlighted alongside the possibility of updating the model with occupants' comfort perception. We conclude that there is an urgent need for research on the automated creation and updating of predictive models if MPC is to become an economically-viable control methodology for non-domestic buildings. Finally, more evidence through operating full scale buildings with MPC is required to demonstrate the viability of this method.

Keywords: Automation; Energy efficiency; Carbon dioxide (CO_2) reduction; Commercial buildings; Control systems; Model predictive control

Introduction

The need to reduce energy consumption and therefore the carbon emissions of buildings is well-rehearsed in the literature. Non-domestic buildings currently account for 18% of the UK's carbon emissions; the situation is made more pressing by the projection that non-domestic floor area will increase by a third by 2050 (Low Carbon Innovation Coordination Group, 2012). Although the processes associated with new build have received great attention, refurbishment is also key to meeting energy-saving targets as the rate of building replacement is low, with 60% of the current building stock likely to still be in use by 2050 (Low Carbon Innovation Coordination Group, 2012).

The energy design gap – the often significant difference between the design projection and actual energy consumption of a building – is an additional factor. In a case study, the Carbon Trust (2012) found that in one building, this discrepancy was a factor of five although the average discrepancy was 16%; 75% of the buildings in their case study used more energy than indicated by their designs. Steps such as: extensive energy monitoring, the Soft Landings framework (Bunn, 2014), seasonal configuration of the control system (Carbon Trust, 2012) and post-occupancy surveys (Bordass, Leaman,& Ruyssevelt, 2001) all have a part to play, and the importance of engaging with occupants should not be underrated. However, in terms of evaluating the energy use and system performance, such resource-intensive processes can only provide a 'snapshot' of a building's energy usage and cannot capture the inevitable drift in its characteristics over time. Frequent repetitions of these procedures would probably be necessary to minimise energy wastage. We question whether this will happen sufficiently frequently is moot; recalibration of the control system would currently seem to be a major undertaking.

Currently, almost all non-domestic buildings employ rule-based controllers, also known as Building Management Systems (BMSs) or Building Energy Management Systems (BEMs) (Levermore, 2000), the fundamental technical basis of which has been criticised by numerous authors, for example, see (Prívara et al., 2013). In essence, the BMS approximates the closed-loop response of the building with a set of handcrafted if-then-else rules. Hathway, Rockett, and Carpenter (2013) have pointed-out that the complexity of this rule set grows exponentially and it is therefore unlikely that anything (near-)optimal operation could be achieved in practice. Furthermore, the commissioning of a BMS tends to be a heuristic affair; a well-performing system is often dependent on reviewing and adjusting its operation over a period of at least a year after handover. Of greater – and more fundamental – concern is the seeming absence of any formal notion of *optimality* in the commissioning of a rule-based BMS. For example, Levermore (2000) notes that it is not uncommon for a BMS to both heat and cool a building simultaneously with obvious waste of energy; Hathway et al. (2013) have even observed such a phenomenon in a recently-completed, award-winning office building.

Reducing the energy consumption of buildings through improved control cannot be discussed without consideration of thermal comfort. Traditionally, the approach to comfort in a highly-engineered building is the specification of temperature setpoints. It is recognised, however, that comfort is a socio-cultural construct which needs to be challenged in order to achieve the optimum reduction in building energy consumption (Chappells & Shove, 2005). Of most interest to control engineering is that when occupants have opportunities for adaptation, they tend to be more positive about the conditions within their environment (CIBSE, 2013). In order to provide the optimum solution for building control, the solutions should not be developed in isolation from developments in thermal comfort and occupant behaviour. To achieve this in large, complex buildings requires a different approach to standard rule-based controllers.

The need for more advanced control is therefore clear, and one highly-promising candidate is model-predictive control (MPC) (Camacho & Bordons, 2004), which has received a great deal of attention for non-domestic buildings in recent years. MPC works by having a predictive mathematical model of the dynamics of the building[1] and repeatedly performing an online optimisation to determine the control law which achieves the desired internal conditions subject to minimising some quantity, for example, energy usage. We discuss MPC in far greater detail below; Afram and Janabi-Sharifi (2014) have recently provided an extensive and wide-ranging technical review of recent work on MPC in buildings.

In this paper, we critically review the state-of-the-art in MPC for non-domestic buildings and attempt to appraise its prospects of delivering the promised energy savings. In particular, we consider the practical difficulties of implementing an MPC controller in a building and the little-discussed issues of evaluating and validating an MPC system. We endeavour to bring a cross-disciplinary approach to the subject and we span: building physics, building services engineering, control engineering and occupant comfort, in order to fully reflect the complexities of buildings and their operation. This is in contrast to the purely control-engineering perspective which has hitherto largely dominated the discussion. Finally, by analysing trends in both building services and control engineering research, we set-out a list of key future research objectives required to make MPC control of buildings a commercial reality.

Model-predictive Control (MPC)

In order to maintain concordance with the conventional control-engineering literature, we will interchangeably refer to the building under control as the "plant".

In general, by far the most popular control methodology is one based on measuring the error between some desired system state and its actual value, and feeding back some linear function of this error to the input to effect a correction to the controlled system state. In practice, the quantity fed back is the sum of a value proportional to the error, its integral with respect to time, and optionally, its time derivative, leading to so-called proportional-integral(-derivative), or PI(D), control (CIBSE, 2009). Because a PI(D) controller employs the instantaneous error, it is not convenient for controlling systems with long time lags between the application of an input and the resulting output response being observed. Long time lags are typical of large reactors in the chemical process industries and this led to the development of model-predictive control as an alternative control paradigm in this sector from the 1970s onwards (Camacho & Bordons, 2004). Long time lags are, of course, typical of heavyweight buildings incorporating exposed thermal mass, a common approach to passive cooling, particularly combined with automated openings.

Starting with a set of system states which are sampled at equally-spaced intervals in time, MPC employs a predictive mathematical model of the plant under control:

$$\hat{\mathbf{y}}_{k+1} = F(\mathbf{y}_k, \mathbf{y}_{k-1}, \dots, \mathbf{u}_k) \tag{1}$$

where $\mathbf{y}_i$ is the vector of system states at time i, $\mathbf{u}_k$ is the vector of control inputs at time k, and $\hat{\mathbf{y}}_{k+1}$ is the one step ahead (OSA) estimate of the system state. In practice, we require $\hat{y}_{k+N}$, the state estimate N steps into the future, which is obtained by repeated application of (1); N is the prediction horizon (i.e. the number of time steps into the future over which the systems aims to effect control.) The only unknown quantity on the right-hand side of (1) is $\mathbf{u}_k$, the system input at time k. One iteration of MPC determines an optimal input sequence $U = \{\mathbf{u}_k, \mathbf{u}_{k+1}, \dots, \mathbf{u}_{k+N-1}\}$ by minimising an objective function:

$$U = argmin \sum_{k=1}^{N} J(y_k, y_{k-1}, \dots, u_k) \tag{2}$$

Typical forms of J are the squared difference between the projected value of a system's state and its target value at that time – the so-called tracking error – possibly combined with a weighted term involving the incremental change in control effort $\Delta\mathbf{u} = \mathbf{u}_k - \mathbf{u}_{k-1}$. (Sometimes the summation of the $\Delta\mathbf{u}$ terms is taken over a shorter period, the control horizon N_c, where $N_c \leq N$(Álvarez et al., 2013) and the set U determined over this control horizon.) Other forms are possible, as we shall discuss below.

In practice, the length of the prediction horizon needs to vary with the building type, a heavyweight building requiring a longer prediction horizon time due its high thermal inertia compared to a lightweight building. In addition, a long prediction horizon tends to make the optimisation step in MPC easier although counter to this, errors in the model's predictions tend to grow with increasing horizon making the predictions less and less reliable with increasing N; a longer prediction horizon also increases the computational effort required to find a solution.

One of the attractive features of MPC is that quite general constraints can be imposed on the solution of the optimisation (2) although physical limitations in the plant may mean that it is not possible to obtain a feasible solution to the optimisation problem that meets some or all of the constraints.

Having performed the optimisation in (2) at time k, and obtained the solution $U = \{\mathbf{u}_k, \mathbf{u}_{k+1}, \ldots, \mathbf{u}_{k+N}\}$, only the first element of U, $\mathbf{u}_k$ is applied to the system and the rest are discarded. The whole optimisation is repeated at the next time step. Since extrapolation into the future inevitably means some uncertainty, which will grow with increasingly-distant projections, the process of only using the first element of U and then updating minimises these uncertainties. This perpetual cycle of controlling over an ever-receding horizon gives the technique its alternative name of receding-horizon control.

In terms of plant characteristics which lend themselves to MPC, buildings generally display the property of a significant time lag between input and response, especially heavyweight buildings. In fact, deliberately designing a building to have a time lag is becoming increasingly popular in locations where there is a substantial diurnal variation in temperature in order to provide opportunities for passive cooling, thus removing the need for air conditioning. Moreover, it is essential to impose constraints on internal zone temperatures, CO_2 concentrations or other measures of indoor environmental quality (IEQ). Finally, since our objective is to achieve the required indoor conditions using the smallest amount of (non-renewable) energy, this is the obvious quantity to minimise in (2)[2]. Consequently, MPC is able to directly address all the key objectives in advanced building control. The main reasons why the predictive, 'lookahead' nature of MPC yields improved control have been summarised by Váňa, Cigler, Široký, Žáčeková, & Ferkl (2014) as: i) knowledge of the building's dynamics means that MPC starts heating just in time to meet the specified conditions, and ii) less energy is used in overheating and in vacant periods. In principle, it is an ideal control method for buildings, the difficulty, of course, lies in its practical implementation.

MPC Implementation in Buildings

From the preceding section, it should be clear that MPC relies critically on the predictive model used to project system states into the future (1). In this section, we consider a range of issues that arise in connection with the predictive model, not just its accuracy. The details of the objective, or cost, function in (2) are also key to system performance and this too is discussed in greater detail. Finally, the architecture of a proposed MPC system has a bearing on its practicality, and we consider this point below. Although MPC could be used to control individual items of building services equipment (e.g. chillers, boilers, etc.), this will not maximise energy savings, or necessarily realise the desired internal conditions. We thus focus on global control of buildings by MPC.

Creation of the Predictive Model

Of all the processes involved in implementing an MPC system, regardless of application domain, the most challenging is widely accepted as the process of producing the predictive model of the plant. In the context of building MPC, Prívara et al. (2013, p.9) note "It is a well-known fact that modeling and identification are the most difficult and time-consuming parts of the automation process". Reporting a roundtable discussion at the first workshop on MPC in Buildings held in Montréal in June 2011, Henze (2013) records that the attendees offered estimates of 70% of project costs being consumed by model creation and calibration. In fact, this figure of 70% does not appear unique to MPC in buildings – in the wider process-control community, the figure of 75% of project costs being attributable to modelling is widely quoted (Gevers, 2005; Hussain, 1999). The important topic of model generation for building MPC has recently been reviewed by Prívara et al. (2013) from which it is clear that much work remains to be done.

In the chemical and process industries, where MPC has been widely adopted, linear predictive models have usually been employed due to their simplicity; furthermore, linear models require a predictive optimisation step – solving (2) – which can be straightforwardly addressed by standard, direct methods from linear algebra, such as linear and quadratic programming. Buildings, on the other hand, are widely accepted as displaying non-linear characteristics due to factors such as rate and output limited sub-systems (Huang, 2011; Afram & Janabi-Sharifi, 2014), and other causes.

To date, however, much of the reported work on MPC in buildings has used linear models where the optimised variables are heat fluxes; in other words, the **u** vector in equation (1) comprises, or includes, heat flux terms (Sturzenegger et al., 2013; Váňa et al., 2014). A heat flux is not routinely monitored by a BMS and has to be inferred. Considering the high-level system diagram of a controlled building in Figure 1, in most reported work, the MPC process takes place within the dotted box to the right of the diagram. If the objective is only to control the zone temperatures, formulation in terms of heat fluxes is highly attractive since it yields a linear problem since heat flux and zone temperature are linearly related; the objective being minimised is the sum of heat fluxes over the control horizon. In practice, it is necessary to formulate an inverse function connecting heat flux output and the controlled variables, which might be quantities such as water temperature in a wet radiator system.

In the case of direct electric heating, there is an obvious linear relationship between the controlled variable and heat flux output. In other cases, such as water-based radiator systems, this mapping is known to be non-linear (Hazyuk, Ghiaus, and Penhouet, 2012; Váňa, Cigler, Široký, Žáčeková and Ferkl, 2014) and highly variable from system to system (Myhren & Holmberg, 2009). It is, of course, possible to calculate the heat flux from a radiator system by integrating with respect to time the product of the mass flow of heating water and the temperature differential across the radiator, but this leads to a non-convex optimisation (Váňa et al, 2014) and is generally avoided in the buildings MPC community. Hazyuk et al. (2012) give an explicit illustration of the empirical fitting of an inverse transfer function for a radiator system. Váňa et al (2014, p.796), however, allude only to "… heat fluxes are then interpreted by a straightforward function which computes particular mass and supply water temperatures based on short term prediction of optimal heat fluxes" in connection with a thermally-activated building system (TABS). It would thus appear much of the work on heat-flux-based MPC implicitly uses a Hammerstein model (Nelles, 2001) in which a static non-linearity precedes a linear dynamic system. Clearly this static non-linearity has to be accurately identified although comparatively little attention has been paid to this in the buildings literature – indeed it often only receives cursory mention.

If the MPC objective is minimising the non-renewable primary energy (NRPE), formulation in terms of heat fluxes is somewhat more problematic. In particular, this requires knowledge of the mapping from NRPE-to-heat flux. One approach that has been employed (Sturzenegger et al., 2013; Váňa et al., 2014) appears to be to argue that total heat flux is proportional to NRPE and therefore minimising total heat flux is equivalent to minimising NRPE. Although heat flux and NRPE are undoubtedly *monotonically* related, the universal validity of this equivalence is unclear.

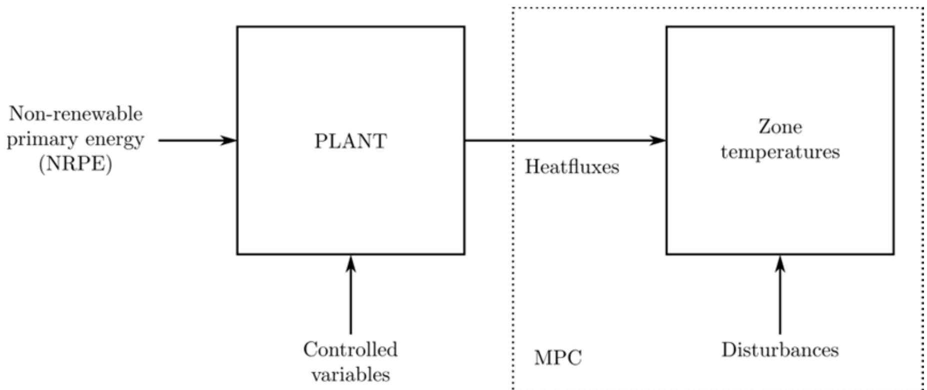

Figure *1*: Typical system architecture for an MPC-controlled building

The fundamental misgiving over the standard heat flux formulation lies in its generality – can it be readily applied to a wide variety of buildings without significant (and expensive) customisation? For example, it may be more energy-efficient in terms of NRPE to intermittently operate a piece of equipment at 100% rating and rely on the thermal inertia of a heavyweight building to store the heat, compared to continuously running the equipment at low efficiency to generate the same total heat flux output over time. An MPC formulation in terms of minimising total heat flux over a control horizon would not seem to differentiate between these two options. Further, the integration of multiple heat sources (district heating, gas boiler, ground-source heat pumps, on-site solar-thermal, etc.) may depend critically on the accuracies of NRPE-to-heat flux mappings.

As observed above, the use of a non-linear model for the whole building system inevitably means that the MPC optimisation step is more complicated, iterative and may not ultimately yield a (global) minimum due to non-convexity. Such an approach has been successfully demonstrated by, among others, Bengea et al. (2014) although again, these authors give little detail on the mapping between input power and zone temperatures. We suggest that much greater attention needs to be paid to the overall system architecture in Figure 1, in particular, i) the identification of any static non-linearities and the ramifications of the Hammerstein model they imply, and ii) whether identification of an overarching non-linear model (and the attendant non-convex optimisation) has greater applicability to a wider range of building types. This has yet to be established, or indeed even investigated, in the literature, as far as we are aware.

In terms of the methodology for producing a predictive model for MPC, there are two basic approaches: white-box modelling, black-box modelling, as well as a third hybrid approach known as grey-box modelling. The white box, or physics-based, approach starts from a fundamental, physical model of the plant. The drawbacks of such an approach are:

- It is prone to improper or inadequate specification; the modelling process may not fully capture the complexities of the plant.
- The process requires a very high level of expertise.
- The engineer formulating the model requires great familiarity with the target building and its proposed operation.
- Many of the necessary model parameters may be unknown, uncertain or even unobservable.

Ultimately, non-domestic buildings are complex systems that strain the comprehension of engineers. As Nishiguchi, Konda, and Dazai (2010, p.118) remark, "... it has been difficult to understand the complicated characteristics of chillers, coils, and pumps, which are widely varied among buildings". Formulating white-box models in a research environment may be possible given enough effort but we question its feasibility in a commercial environment, especially when, in the case of retrofitting, detailed drawings of the building may not even exist. Further, the notorious problem of quality assurance in the construction industry means that what is actually built sometimes differs from what the designers have specified. Prívara et al. (2013) suggest that methods which rely on a physical description of the building are suitable only for simpler structures.

Black box, or data-driven, approaches construct an empirical model of the plant; appropriate models include:

- Subspace methods (Ferkl & Široký, 2010)
- Autoregressive moving average with exogenous inputs (ARMAX) models (Ferkl & Široký, 2010)
- Autoregressive models with exogenous inputs (ARX) (Y. Ma, Kelman, Daly, & Borrelli, 2012)
- Non-linear ARMAX models (Billings, 2013)
- Box-Jenkins models (Box, Jenkins, & Reinsel, 1994)
- Artifical neural networks (Ferreira, Ruano, Silva, & Conceição, 2012; Nelles, 2001)
- Fuzzy logic
- Thin-plate splines (Nishiguchi et al., 2010)

See Afram and Janabi-Sharifi (2014) for a detailed review of MPC-relevant models. The empirical learning of models from data is often described by the umbrella term of machine learning in the statistics literature (Hastie, Tibshirani, & Friedman, 2009), or system identification (SID) in the control-engineering field (Ljung, 1999). The structure of the (simplified) model usually has to be determined by trial-and-error, a statistically-founded procedure termed model selection (Hastie et al., 2009), which requires significant expertise. For example, Ferkl and Široký (2010) have explored the use of ARMAX models for a complex building, describing the search for appropriate model orders as "tedious", and estimating the computation time to identify the ARMAX model by 'brute force' enumeration as around 209 days, necessitating the use of an "engineering approach" and experience. Consequently, black-box approaches, like their white-box counterparts, also require significant levels of expertise but in a different area from white-box modelling. Another key aspect of data-driven models is missing data – Váňa, Kubeček, and Ferkl (2010) note the significant complications for calibrating a black-box predictive model caused by the 'drop-out' of monitored data.

From a classical system identification standpoint, producing an accurate black-box plant model requires the excitation of all the modes of the system (Söderström & Stoica, 1989) and this is typically accomplished by applying a pseudo-random binary sequence (PRBS) to the system in open-loop. (A PRBS is a bounded-amplitude signal which approximates the spectral properties of white noise, and thus excites all possible system modes (Söderström & Stoica, 1989).) Such an approach has been employed in buildings by, among others, Aswani, Master, Taneja, Culler, and Tomlin (2012), Ferreira et al. (2012) and Váňa et al. (2010). Although feasible in the process industries, such an open-loop system identification experiment is very inconvenient in an occupied building due to high energy costs and the discomfort it may cause. Open-loop SID often drives the systems states to extremes so there is a danger that occupants, when faced with uncomfortable conditions, will react by opening windows, using personal heaters, etc. which may well subvert the system-identification experiment. The duration of these experiments may also be problematic with Rivera, Lee, Braun, and Mittelmann (2003) citing identification times up to almost one month in the chemical process industries. Ultimately, it is not clear whether such an exact model is even necessary. In essence, the most important feature of a predictive model for MPC is that it only needs to be accurate over the prediction horizon and effort should be expended to this end; its accuracy beyond the prediction horizon is of far less concern.

Since modelling effort is so important in the general control field, and complex models are costly to calibrate accurately, there has been much consideration of "identification-for-control" where the objective is to produce a (simpler) model which is sufficient for its intended purpose, rather than completely identified (Hjalmarsson, 2009). The subject of identification-for-control, or control-relevant identification, has been reviewed by Gevers (2005); Zhu and Butoyi (2002), for example, have investigated the closed-loop identification of some ill-conditioned chemical process plants. Given the high total cost of model identification, it is often cost-effective to implement an approximate system model, operate the plant (sub-optimally) under closed-loop control and use the data gathered under closed-loop conditions to perform a second identification procedure to produce an updated controller; further, subsequent updates are, of course, possible. A key issue is that under closed-loop control, the system responses will be dominated by those modes that are important for control, thereby facilitating the derivation of simpler but sufficiently-accurate plant models. Crucially, during such a closed-loop SID experiment, the plant remains under feedback control, albeit not necessarily optimal. Nonetheless, re-identification under closed-loop has its technical challenges since the nature of feedback control will treat any input stimulus designed to elicit the system response as a disturbance, which the control system will try to eliminate. The basic approach is to inject a 'dither' signal uncorrelated to the disturbance (Genceli & Nikolaou, 1996; Rathouský & Havlena, 2013; Marafioti, Bitmead, & Hovd, 2014). Typically, this dither signal is used to perturb the system setpoints (Zhu & Butoyi, 2002) although other approaches have been explored (Sotomayor, Odloak, and Moro, 2009). See Sotomayor et al. (2009), and Zhu & Butoyi (2002) for additional details.

A further, hybrid approach to constructing predictive models uses an electrical circuit (resistor-capacitor, or RC) analogy to associate every first-order linear process with resistors and capacitors (Levermore, 2000; Coley & Penman, 1992; Prívara et al., 2013). Such an approach by analogy is often termed a grey-box model and is restricted to linear models. Prívara et al. (2013) remark that this approach performed better than the more principled 4SID subspace method (Cigler & Prívara, 2010). Since it requires detailed knowledge of the building structure, the grey-box approach can be expected to suffer from similar shortcomings as white-box modelling. In a rare study, Behl, Nghiem, and Mangharam (2014) have analysed the sensitivity of a fitted RC model to random errors in the training data.

The practical difficulties in constructing a predictive model should not be underestimated. There is clear evidence that model formulation is the most challenging aspect of deploying MPC even in its traditional enclaves of the chemical process industries. Unless the predictive model is sufficiently accurate, overall system performance cannot be expected to be acceptable.

Cost Function

MPC calculates a control law by minimising an objective function over some control horizon, possibly subject to constraints. The efficacy of MPC is critically influenced by the objective that is minimised. Afram and Janabi-Sharifi (2014) have listed five objectives that have been used previously in the buildings literature: tracking error, control effort, energy cost, demand cost, power consumption, or a combination of these. In addition, both Ferreira et al. (2012) and Cigler, Prívara, Váňa, Žáčeková, and Ferkl (2012) have optimised a measure of thermal comfort, predicted mean vote (PMV). PMV is a commonly used representation of thermal comfort. However, it was developed for fully air-conditioned spaces, and there is evidence it is not appropriate in free-running, passively-conditioned spaces where occupants expect daily and seasonal variation (de Dear & Brager, 1998). Further, it does not account for psychological and social effects on an occupant's feelings of comfort and acceptance of internal conditions. In order to reduce the energy use of our buildings, it has been argued that there needs to be more focus on this factor (Cole, Robinson, & O'Shea, 2008). West, Ward, and Wall (2014) addressed occupant comfort by incorporating percentage of people dissatisfied (PPD) with their comfort – a measure related to PMV – into a more conventional cost function comprising energy costs, and carbon emissions. It is interesting to note that Cigler et al. (2012) were able to save an additional 10-15% of energy over conventional setpoint tracking due to the more relaxed tolerance on maintaining a given zone temperature.

The form of the cost function for linear MPC based on a heat flux (implicit Hammerstein) model has been reviewed by Cigler, Široký, Korda, and Jones (2013) who considered a general cost function of the form:

$$a^T u + (y - r)^T Q (y - r) \tag{3}$$

where $u \in R^{N_c}$, $y, r \in R^{N_p}$, $N_{c,p}$ are the lengths of the control and prediction horizons, respectively and $N_c \leq N_p$, a is a weight vector such that $a_i \geq 0 \forall i \in [1 \dots N_c]$, and Q is a matrix of weights. The vector r has two possible interpretations. Firstly, as a vector of temperature setpoint values, or second, as a vector of slack variables designed to enforce constraint bands on the zone temperatures.

The elements of vector u are heat flux inputs, so minimising only the $a^T u$ term minimises total energy input. Cigler, et al. (2013) have pointed-out, however, where there is a plant-model mismatch, the linear programming optimiser typically used can exert `bang-bang' control in which the heat flux switches in alternate sampling intervals from fully on to completely off, with an attendant waste of energy and poor control. Using an objective of the form $u^T S u$ can be solved with quadratic programming and does not exhibit this problem, but has no physical basis as energy use is not proportional to the square of the heat flux (Cigler et al., 2013).

Considering both terms in (3) where r is a vector of temperature setpoints introduces temperature tracking as an objective (Cigler, et al., 2013); the emphasis placed on the two terms depends on the relative values of a and Q. Although common in the chemical process industries, such a control method is unduly restrictive – precisely controlled temperatures are not necessary in buildings – as well as using more energy than necessary.

As a final option, r can be interpreted as a vector of slack variables in which case (3) prescribes a constrained optimisation which, while potentially using less energy, is not free from highly oscillatory control if the $a^T u$ term in (3) dominates. Cigler, et al. (2013) go on to discuss a number of ways of smoothing this undesirable 'bang-bang' control action.

We reiterate, the non-renewable primary energy (NRPE) and heat fluxes are only monotonically related. In principle, heat flux and control variables, such as temperatures in a water-based heating system, can be related precisely by considering the summation (or integration) over some period of the mass flow of the heating/cooling medium and the differential between supply and return temperatures. Such a formulation, however, yields a non-convex optimisation.

Non-linear cost functions have been employed in buildings MPC (Bengea, et al., 2104) although we are unaware of any comparative studies of linear versus non-linear cost functions for buildings MPC.

System Architecture

There are three main architectures for MPC: centralised, distributed and decentralised. In centralised MPC, the whole optimisation is performed within a single computational process which, while simple to implement, may not scale well for large buildings with large numbers of state variables, constraints, etc. For this reason, distributed MPC has been much studied and where the optimisation is performed with a set of simpler, 'local' optimisations using only a subset of the variables and constraints, and exchanging these with 'neighbouring' processes; the objective is therefore to decompose the large optimisation into a number of sub-optimisations where the degree of coupling between sub- optimisations is limited. The design of such systems is much more complicated than for centralised MPC and depends on whether the locally-minimised objectives can be combined into a single sum, in which case a global optimisation is feasible. If they cannot, then the problem is multi-objective representing a trade-off between competing solutions. Distributed MPC has been reviewed by Camponogara, Jia, Krogh, and Talukdar (2002); application to buildings has been reported by Lamoudi, Alamir, and Béguery (2012), Álvarez et al. (2013) and others.

Whereas distributed MPC specifically accounts for coupling between local agents representing, for example, zones within a building, decentralised MPC divides the control system across a set of uncoupled processes. Due to each sub-optimisation working with incomplete information, performance is unsurprisingly not as good as for centralised or distributed MPC (Afram & Janabi-Sharifi, 2014).

Incorporating weather forecasts into building MPC can improve its performance (Afram & Janabi-Sharifi, 2014). Moreover, also accounting for the errors in the weather forecast can improve matters further still (Oldewurtel et al., 2012) since it provides direct information on a key system disturbance. This is to be expected as additional (pertinent) information can only improve system performance. Florita and Henze (2009) have compared a number of local weather forecasting approaches. In terms of current experimental implementations of MPC for real buildings, we consider the system described by Sturzenegger et al. (2013) to be an exemplary engineering solution. In a trial on a Swiss office building of 6,000 m^2 and five floors as part of the OptiControl project[3], these authors implemented centralised supervisory control retaining the existing low-level control of the BMS. The software was implemented in (interpreted) Matlab on a standard PC with state sampling every 15 minutes and Kalman filtering of the building's states; the control law required less than 2 minutes to calculate using a bilinear plant model. The MPC optimisation minimised total energy supplied to the building while imposing constraints on the internal temperatures and minimum air flow. Local weather forecasts were used and the system implemented a fallback strategy in the event of the MPC algorithm failing (which it never did over the period of the trial).

Evaluation of MPC: The Prospects for Energy Saving

By far the most frequent demonstration of MPC for buildings has been in simulation, probably for the highly-understandable reasons of cost, and that building operators are reluctant to allow their buildings to be used to trial an as-yet unproven control technology. Simulation, however, is not without its limitations. Firstly, simulation is a white-box approach and therefore has all the shortcomings of inadequate specification and uncertain parameters. As an example, J. Ma, Qin, Salsbury, and Xu (2012) performed a full system identification procedure (i.e. model derivation) and energy-saving appraisal entirely within EnergyPlus, suggesting energy savings of 25%, although the approach was not subsequently validated on a real building.

Second, building simulators, such as EnergyPlus (Crawley, Pedersen, Lawrie, & Winkelmann, 2000), do not include stochastic effects due to occupants by default; the effect of occupants on a building's performance is regarded as highly significant (Carbon Trust, 2012). For example, in a simulation study of a mixed-mode building, Tanner and Henze (2014) concluded that the potential energy savings of MPC were roughly halved by the inclusion of a stochastic model of occupant window opening. We consider the role of the stochastic behaviour of occupants in a following section.

In consequence, we consider that simulation studies, while often invaluable, should be treated with some caution since it is unclear how well computer models capture the behaviour of real buildings; this issue has been discussed by Sturzenegger, Gyalistras, Morari, and Smith (2012), and J. Ma et al. (2012).

Although there has been a number of reports of the control of single rooms, comparatively few results have been reported for MPC on real buildings of significant scale. Nishiguchi et al. (2010) have reported results for a 15,000 m^2 hotel and estimated energy savings of 9%. Ferreira et al. (2012) reported results for a group of four rooms in a university building and estimate the energy savings as "probably above 50%". It is noteworthy that many of the reports of MPC in real buildings have been obtained in university buildings which, arguably, have atypical, highly-seasonal and unpredictable occupancy patterns. It is also noteworthy that the periods over which many MPC results have been obtained are often on the scale of days. Notable longer-scale trials include:

- Prívara, Široký, Ferkl, and Cigler (2011) reported operation of the whole of a building at the Czech Technical University, Prague comprising seven "blocks" for "February 2010". Experimental details are sparse but energy savings of "17-24%, depending on the particular building block" (p.569) and "29%" were claimed.
- Bengea, Kelman, Borrelli, Taylor, and Narayanan (2014) conducted a trial over 20 days in a mixed-use 650 m^2 building zone (although this was an intermediate trial and the subsequent experiment on the refined MPC system lasted only five days).
- Sturzenegger et al. (2013) successfully operated an office building of 6,000 m^2 over five floors for a total of 92 days.
- West, et al. (2014) operated a three-floor, 3,322 m^2 office building in Newcastle, New South Wales, Australia for around 25 days.
- Váňa et al. (2014) have operated a five-storey office block of 1,500 m^2 per floor for 25 days; an average saving of 17% was claimed.

Although existing reports are encouraging, there is inevitably some element of inference about published energy saving results to date. Sturzenegger et al. (2013), despite successfully operating a large office building with MPC as part of the OptiControl project, terminated the trial after three months and resorted to simulation with EnergyPlus for a performance appraisal for a full year. As Sturzenegger et al. (2013, p.3233) remark, "The sequential nature of on-site experiments and the varying operating conditions make experimental comparison of different controllers very difficult". We suggest that the ideal testbed should comprise two identical, adjacent buildings with identically-behaving occupants and subject to the same weather and solar gain. MPC could be implemented in one and the second allowed to continue with a conventional BMS, and the results compared over a fairly lengthy period of time. The opportunity to conduct such a controlled experiment, however, is probably rare although such situations can present. Váňa et al. (2010) describe a building at the Czech Technical University in Prague which comprises several identical blocks although it is unclear if the comparisons presented by these authors actually used one of these blocks as an experimental control.

Comparisons using a single building are more problematic and require periods of (almost-) identical weather/occupancy, etc. although there are methods to normalise for weather effects and occupancy when calculating operational ratings over a year (CIBSE, 2008) which may assist in evaluating the performance of MPC. In a comment on their study, Bengea et al. (2014, p.128) note, ". . . identical conditions were very difficult to establish due to the intrinsic highly variable weather and indoor usage patterns…". More generally, we suspect that rather more indirect and inferential methods will have to be used to validate energy savings due to MPC with consequent uncertainty in the savings estimates. The projections for energy saving with MPC also need to be viewed critically. For example, Bengea et al. (2014) report some spectacular energy savings (60-85%) in a study on a real building but these authors appear to be comparing between a baseline system where the schedules were ". . . heuristics . . . generated by [an] installation engineer" (p.127), and an MPC system configured by highly-skilled control engineers. As Bengea et al. (2014, p.134) concede, ". . . a more fine-tuned [comparator system] . . . would have captured some of the savings achieved by the MPC scheme". The work of these authors is noteworthy, however, because they investigated *why* their MPC system produced more efficient operation and pointed-out that the strategies it was adopting, while commonsensical, were being implemented automatically.

Many of the reports, both simulation and for real buildings, estimate energy savings in the range 15-30% while maintaining or even improving internal comfort, which, at face value, is an impressive saving. Many of these reports, however, have implemented MPC by minimising deviation from some temperature setpoint. As a fundamental principle, we cannot see why accurately controlling temperature would necessarily yield a reduction in energy consumption – we can easily understand why a technically-superior method of controlling temperature would deliver better temperature regulation but not why this should necessarily be accompanied by a reduction in energy usage. Indeed, Váňa et al. (2010, p.1020) remark that ". . . the high (energy) savings achieved by MPC could be mere coincidence". We sound what we believe to be an important note of caution in assessing the prospects of MPC. We conjecture that the reported energy savings with temperature-controlled MPC may be due to the fact the existing BMSs in the test buildings were poorly calibrated, and so almost any 'care and attention' lavished on such a building would have produced some energy savings. As further (circumstantial) justification for this suspicion, we are aware of a number of building-energy consultancies in the UK currently advertising that they can achieve 15-30% energy savings for their clients by careful scrutiny of building operation and 'fine-tuning' of the existing BMS. The existence, and indeed apparent commercial success, of such operations[4] suggests that a great many non-domestic buildings in the UK (at least) have very poorly calibrated BMSs. If our supposition is correct, however, this does not diminish but critically changes the arguments surrounding MPC.

Firstly, MPC has, at its core, the notion of *optimality*. Hence a control system that continuously optimises can be expected to deliver consistently greater energy savings over the longer term than periodic 'fine-tuning' by consultants.

If the objective of implementing MPC is to achieve acceptable comfort with the minimum energy usage, the objective function to be minimised should be energy consumption and not temperature deviation from a setpoint; we argue above that temperature regulation will not necessarily minimise energy use. Thermal comfort should be achieved by imposing constraints on the MPC minimisation of energy use. In fact, a number of researchers have already followed this approach. For example, both Bengea et al. (2014) and Y. Ma et al. (2012) minimised energy cost (which is some proxy for energy consumption) under dynamic energy-pricing environments, and Sturzenegger et al. (2013) minimised a measure monotonically related to non-renewable primary energy.

Even if we can assume that the BMSs in existing buildings were properly calibrated on hand over, the inference is that the BMSs of many buildings are now poorly calibrated and wasting significant amounts of energy. This suggests far greater attention should be paid to frequent updating of control strategies. We expand on this theme below.

Controller Updating

If we proceed from the (actually highly questionable) assumption that a newly-commissioned building is set-up optimally, changes in the way the building operates and its environment over its lifetime are almost inevitable. Aside from degradation of the HVAC plant, organisations expand or contract, and adapt their operations. Building tenants may change and a building may be subject to modification of internal partitions, all of which may significantly modify the thermal characteristics of the building. Even more subtle than internal changes, Hathway et al. (2013) raise the intriguing case of external changes where, for example, a building across the street from the building under control is demolished, or a new one erected, either of which can change the solar gain and/or wind pressures on the façade of the building under control. The very practical dilemma faced by the facilities manager of the controlled building is whether to incur the (possibly high) cost of BMS recalibration and balance this against future energy savings by returning their building to optimal operation, occupant satisfaction, etc. The effect of external influences on a building are far from a theoretical concern given the pace of urbanisation.

Given the reality of near-continual changes in building operation and environment, we believe that there is a compelling case for research on continuous updating of the building control system to ensure that the building is running at near-optimality throughout its whole lifetime. Fortunately, the basis for this continual refinement of the building model already exists in the iterative re-identification in closed-loop ('identification-for-control') paradigm (Gevers, 2005). Here the system remains under nominal control but is subjected to additional input stimuli designed to identify the current model. If MPC is to deliver its full potential in buildings, we believe model updating needs to be actively explored.

A highly-relevant subsidiary issue is allowing model structure to vary. Typically, hand-crafted predictive building models are of fixed functional form but if the building's characteristics change, it is likely that the form of the 'best-fitting' model will also change. This too is an area for future research. Note that allowing the functional form of the model to be updated is fundamentally different from periodic re-estimation of the parameters of a fixed model (West et al., 2014).

The Role of Occupants

Ultimately, all buildings – with occasional exceptions, such as data centres – are intended for people, and any control strategy should not adopt an excessively technological approach that reduces occupants to the role of nuisances who complain about the conditions. To date, comparatively little attention has been paid to occupants' perceptions of comfort in the MPC-related literature, the prevailing control-engineering-centric assumption being that maintaining an air temperature within bounds equates to occupant comfort.

In one of the few MPC-related studies to consider the subject, West et al. (2014) elicited occupant comfort on a seven-point scale using an online feedback tool installed on occupants' computers, and minimised a weighted objective within an MPC framework comprising: the percentage of people dissatisfied (PPD) with their comfort, energy costs, and carbon emissions. Interestingly, these authors observed that the mere mention of comfort sensitised occupants to the issue even before their MPC trial began, as measured by numbers of complaints. In addition, Ferreira et al. (2012) and Cigler et al. (2012) have directly optimised predicted mean vote (PMV), a measure of occupant thermal comfort.

Although these studies do consider comfort as being wider than merely temperature by accounting for other factors, such as the use of PMV, they are still basing the consideration of comfort on values that are physically measurable (e.g. air temperature, radiant temperature, humidity). In fact, there is growing evidence that comfort is more complex. Chappells and Shove (2005) have persuasively argued that ". . . comfort is a highly negotiable socio-cultural construct". Indeed, it has been suggested that occupants are more satisfied when they have the perception of control, the definition of which is complex and incorporates: the control options, the efficacy of the control action, individuals' expectations and preferences (Hellwig, 2015). The demands we set on our indoor environment are also acknowledged to vary if the building is ventilated and cooled by passive means (where control may still be necessary for automatic windows and night cooling) (de Dear & Brager, 1998). New guidelines in the UK base the internal setpoint on a relationship to the external running mean temperature (CIBSE, 2013). Since 2005, the government of Japan has run a "Cool Biz" campaign encouraging the wearing of more casual clothing at work in order to make higher indoor air temperatures more tolerable. Although some reports indicate annoyance at the higher temperatures (Tanabe, Iwahashi, Tsushima, & Nishihara, 2013), that the scheme has been running for nearly a decade indicates some success, particularly as the annoyance appears to occur mainly at temperatures greater than 27 C. Both the "Cool Biz" campaign and the fundamental re-examination of comfort highlight the hitherto little-considered role of policy on the requirements of a control system. A further issue is the necessity to define the controlled states, e.g. setpoint temperatures. Defining the model through an empirical process, however, opens up the possibility of incorporating other variables, such as some form of feedback from occupants (Gunay, O'Brien, & Beausoleil-Morrison, 2013). For instance information from adaptive actions by occupants (such as window opening, or changing a thermostat setting) can be used to improve the prediction of the control model. This raises the prospect of adapting models as our understanding of comfort develops in the future.

As to quantifying the impact of occupants on buildings, the development of stochastic models of behaviour suitable for incorporation into energy simulators has attracted a great deal of attention in recent years (Page, Robinson, Morel, & Scartezzini, 2008; Gunay et al., 2013). The implications of stochastic occupant behaviours on energy saving with MPC have been examined by Oldewurtel, Sturzenegger, and Morari (2013) who concluded from a simulation study of two different HVAC buildings that long-term occupancy predictions do not produce any significant energy savings, and that almost all the available savings due to occupant absence can be realised by reactive cessation of lighting and ventilation as soon as an occupant absence is detected. Gunay, Bursill, Huchuk, O'Brien, and Beausoleil-Morrison (2014), however, identified significant errors in the accuracy of a (simple RC) predictive model of zone temperature for a single-occupancy office due to the stochastic effects of presence, lighting and blind operation, with consequent increased energy use and thermal discomfort relative to a baseline (non-MPC) reactive controller. On the other hand, Tanner and Henze (2014) observed, again from a simulation (of a single building over a single month), that stochastic window-opening behaviours of occupants reduced the energy-saving potential of MPC by a factor of two in a mixed-mode building.

Clearly, more research on the effects of different stochastic occupant behaviours is needed to establish the generality (or otherwise) of the limited, often quite specific, simulation results described above.

Discussion: Open Questions in MPC Buildings Research

If MPC is to be seen as a viable and trustworthy technology, longer-term trials of successful MPC operation are needed, not just the tens of days that have been commonly reported to date, so that i) energy savings can be quantified for different building types, and ii) the long-term consistency of these savings can be assessed. In fact, along with management-level 'buy-in' to the process, Henze (2013) has suggested establishing a database of show-case projects. In this regard, the KTH Open Testbed (Pattarello, Wei, Ebadat, Wahlberg, & Johansson, 2013) and the US Department of Energy FlexLab facility at the Lawrence Berkeley National Laboratory, Berkeley, CA are an interesting approach to providing a common research platform. Greater experience in operating a number of full-scale buildings using MPC is urgently needed[5].

Any risks of disruption caused by MPC malfunction can be mitigated by switching to a back-up, known-reliable BMS in the event of problems, as implemented by Sturzenegger et al. (2013). In fact, disconnection of the controller is an established approach in MPC to handling a failure to find a feasible solution (Camacho & Bordons, 2004, p. 198).

Notwithstanding the large volume of research currently being carried-out on MPC in buildings, a number of fundamental technical challenges remain to be addressed.

The production of the predictive model of the building remains foremost, and is currently both expensive and needs highly-skilled personnel. Whether current research experience can be translated into commercial deployment is uncertain, and in a comment on a highly-successful trial of MPC on a real building, Sturzenegger et al. (2013, p.1035) observe that "... the efforts undertaken in this research project would be prohibitive in an industrial application of MPC". Indeed, Sturzenegger et al. (2013, p.1035) go on to remark, rather pessimistically, that ". . . the derivation of a good MPC applicable model is not expected to be easily standardized for commercial application, which makes . . . [modelling] . . . the most critical, if not the currently prohibitive factor".

We have noted above the estimates by attendees at a workshop on MPC in buildings held in Montréal in June 2011 that 70% of project costs are typically consumed by model creation and calibration (Henze, 2013). (The whiteboard notes of the same roundtable discussion[6] also express the sentiment that "tunable model complexity" was regarded as the "Holy Grail" of MPC in buildings.) Clearly more cost-effective model generation methods are urgently required if MPC is to be economically widely-adopted. Possibly the area of data-driven control (Hou & Wang, 2013) may present a way forward. Given the large amount building-specific information required to formulate accurate white-box and grey-box models (Sturzenegger et al., 2013), black-box models seem to be a highly-feasible approach. Possibly a different sort of hybrid method where appropriate physics-based models of individual items of plant (e.g. a single fancoil unit) are embedded in a black-box framework may be usefully pursued, as suggested by Ljung (2010), although such an approach may lack the robustness to accommodate degradation due to wear-and-tear of the white- box-modelled components. As a counter-argument, Váňa et al. (2014, p.792) contend that grey-box modelling is to be preferred as black-box modelling "usually spoils the real system structure" (although quite what these authors mean by this is unclear). Regardless, economically-feasible approaches to model generation are urgently required.

A currently unresolved, and important, issue in MPC for buildings is the form of the predictive models, in particular, the role of non-linear models. There seems widespread, even near-universal, agreement that building dynamics are non-linear. For example, CIBSE (2009, p.2-13) states "… the operation of a complete HVAC system is highly non-linear…"; similar comments would apply to naturally-ventilated buildings relying principally on convection. Little work has been reported on non-linear overall system models, and we are aware of no study that has sought to address the issue of linear versus non-linear modelling. As noted elsewhere, much of the published work on buildings MPC has employed a Hammerstein system model in which a static nonlinearity has been followed by a linear dynamical model formulated in terms of heat fluxes. Such an approach does not eliminate the non-linearity inherent in the system, but, rather, isolates it as a separate entity to be estimated independently. We are aware of no work which has explored the influence of the static nonlinearity and the accuracy of its estimation on the MPC performance of a Hammerstein building model. Traditionally, control engineers appear to have shied away from fully non-linear models due, we suspect, to the non-convex optimisations that result although increasingly-available computational power has diminished this objection in recent years; the lengthy sampling intervals (10-60 minutes) in buildings MPC further dilutes this objection. There remains, of course, the problem that a non-convex optimisation cannot be guaranteed to find a global minimum. Non-linear models are, nonetheless, established in the wider MPC literature (Camacho & Bordons, 2004), and have been successfully used by Bengea et al. (2014) on a centralised supervisory system controlling a mid-sized building. Ljung (2010) has discussed the research challenges in the identification of non-linear models for control, in general. We consider that the necessity, or otherwise, of using non-linear plant models for buildings should be investigated, and compared to Hammerstein models.

Although the economic generation of the predictive model is critical for the uptake of MPC in buildings, an equally important aspect is the updating of the model to reflect the inevitable changes in the building with time. We have discussed above the potential for 'identification for control', and this area has become established in the chemical engineering field, and elsewhere. Applying it to buildings, however, requires additional consideration. For example, what perturbations on internal conditions are acceptable during re-identification? Also, should re-identification take place periodically or in response to evidence that the predictive model is no longer appropriate – the latter, automated approach is starting to emerge in the chemical engineering literature (Conner & Seborg, 2005). Effective model updating is clearly important for ensuring whole life-cycle optimisation of building performance.

Our clear view is that MPC optimisation should minimise non-renewable primary energy, and impose bounds on internal temperature and CO_2 concentrations as constraints; those bounds may be based on adaptive principles (CIBSE, 2013). Moreover, we need to account for rare events when no feasible solution is possible due to a combination of extreme weather and physical limitations on plant capacity. In that situation, 'soft' constraints (Camacho & Bordons, 2004, p. 199) should be employed so that a solution is always found; when the constraints cannot be met, the constraint violation is thus reduced to its absolute minimum. Previous work on temperature control appears to be a straightforward transfer of existing control approaches to MPC and, as we have argued above, will not necessarily minimise energy usage. Further, the minimisation of non-renewable primary energy should allow for the straightforward integration of renewable energy sources into the building's energy management system, and avoid the problem cited by Carbon Trust (2012) where a conventional rule-based BMS failed to differentiate between energy supplied for heating by a gas boiler and that supplied by a ground-source heat pump.

To produce a robust implementation, MPC also needs to be able to handle missing data (i.e. data dropout) both in the model construction phase (Váňa et al., 2010) and during operation. We differentiate between temporary data loss due to communications 'glitches', and persistent loss of data. The effect of data glitches can probably be handled satisfactorily by Kalman filtering the sensor outputs (Sturzenegger et al., 2013). As for longer-term data loss, sensors can malfunction or even be stolen. In the case where their MPC optimisation failed to yield a solution, Sturzenegger et al. (2013) used the strategy of applying the second, normally discarded, element of the control sequence from the previous optimisation. Potentially, MPC should also serve as a diagnostic tool to detect malfunctions/degradation in the HVAC plant or sensors (Ljung, 2010). If the optimisation is consistently unable to produce a feasible solution, this should trigger an investigation in advance of occupant complaints about conditions.

Related to the above issue on model updating is the question of how do we obtain the first model to start controlling a building with MPC? This applies both to new build and to refurbishment projects: in either case, the model-predictive controller needs to be 'bootstrapped' somehow. By implication, any initial model will be replaced by a re-identified model and hence should be 'cheap' to generate. Potentially, one solution is to use an approximate but quick-to-compute initial model such as a grey-box model (Žáčeková et al., 2013) which will be rapidly replaced by a more accurate updated model identified in closed loop (identification-for-control). (The work by Žáčeková et al. (2013) is noteworthy as it is the first report, of which we are aware, of closed-loop re-identification in building MPC.) Landau (1999, p.1115) notes that "most practical results have shown the major improvement in performance occurs after the first identification in closed loop" suggesting such a procedure may converge very rapidly to a good controller. Clearly, the newly-commissioned control system would not initially run optimally but should soon approach optimal performance after a few model updates.

Although the principal focus of this paper has been on control, it is an obvious truism that control can only ever be as good as the sensor information upon which it acts. Unfortunately, very little work has been done on sensor placement and the accuracy with which temperature can be measured in a zone. Heat transfer in occupied spaces is usually very complicated; meaningful temperature sensing in the large, open-plan spaces which are typical of modern office buildings is even more problematic. In addition, convective flow in high spaces, such as atria, means that the notion of a single zone temperature does not exist in practice. We also know of a BMS installer who handed over a building to an occupier who immediately fitted a ceiling-mounted data projector directly beneath the single temperature sensor in a room. Since it then 'measured' serious overheating, the BMS applied maximum cooling, which made the space unacceptably cold. There has been little published work on sensor placement for temperature control; we are aware of only the work by Riederer, Marchio, and Visier (2002), which considered only a single, medium-sized room. The critical subject of sensor placement related to control requires more research if it is not to become the achille's heel of advanced control systems. Two areas of investigation suggest themselves:

Firstly, using a large, redundant set of sensors and selecting the most 'informative' during periodic updating of the predictive model. Methodologies for selecting the most useful subset of inputs are well-established in the machine learning literature where they are known as feature, or variable, selection – see, for example, Guyon and Elisseeff (2003). The emergence of wireless technologies has made installing large numbers of sensors economically feasible although, conversely, the widespread use of metallised surfaces in modern, energy-efficient buildings presents a challenge for reliable radio-frequency propagation (Subrt & Pechac, 2012).

Secondly, temperature has traditionally been taken as synonymous with comfort although the difficulties of accurately measuring the temperature of complex spaces makes this association tenuous. A potentially useful, but radical, avenue is to simply abandon temperature as a control objective and treat the output of temperature sensors as quantities that are *correlated* with occupant comfort as opposed to being a direct measure of comfort. Eliciting occupant perceptions of comfort would then become key to successful building operation; essentially, this appears similar to the current practice whereby the building operator 'tweaks' the set points in response to occupant complaints although more automated ways of achieving this are beginning to emerge in the commercial sector[7].

Finally, it is also worth highlighting what appears to be a fundamental difference in emphasis between US and European research on MPC. US work emphasises the management of thermal storage elements, their effective integration into MPC for load shifting and peak 'shaving', and monetary cost as an objective to be minimised – see, for example, J. Ma et al. (2012). European work tends to emphasise energy saving. This difference probably reflects the economic motivators of the US research, driven by the dynamic electricity tariffs now common in the US. European work, on the other hand, emphasises energy saving probably driven by policy goals laid-out in Directive 2010/31/EU of the European Union on the energy performance of buildings (European Union, 2010).

Conclusions

We have critically reviewed previous work on, and the prospects for, model-predictive control (MPC) in non-domestic buildings. Since it has optimisation at its heart, MPC seems a strong candidate to supersede current, reactive rule-based control in buildings and to deliver significant energy savings. In order to make this an economically-viable method of delivering optimal building control, however, there are several areas which require more research. Most critically, more cost-effective methods of model generation need to be developed. The structural form of the model best-placed to meet this requirement still needs significant research – whether non-linear models provide a solution generalizable across the building stock, or whether a Hammerstein solution is more effective needs significant attention.

In order to enable whole-life cycle optimisation of building energy use, the model underpinning MPC needs to track the changes in building performance over time, and therefore investigation into model re-estimation is required to ensure continuous, efficient operation. To meet climate change objectives, the MPC optimisation needs to minimise non-renewable primary energy using the internal conditions as constraints. If we employ model updating, the first model generated when we begin controlling the building does not need to be perfect. It is then possible to use an simple initial model to begin control knowing it will be replaced by an improved model in the near-future.

The predictive model, and hence the control, relies on the quality of the input data, therefore the role of sensor reliability and location needs to be understood. More radically, the use of MPC opens-up possibilities for the control to be based on occupant comfort, rather than on the temperature in the space. This would require further research into ways of eliciting occupant feedback as input to the predictive model.

Finally, experience in operating full scale buildings over long periods of time is urgently required in order to demonstrate that MPC is indeed a viable technology for optimising building performance.

References

Afram, A., & Janabi-Sharifi, F. (2014). Theory and applications of HVAC control systems - A review of model predictive control (MPC). *Building and Environment, 72*, 343-355. doi: 10.1016/j.buildenv.2013.11.016

Álvarez, J., Redondo, J., Camponogara, E., Normey-Rico, J., Berenguel, M., & Ortigosa, P. (2013). Optimizing building comfort temperature regulation via model predictive control. *Energy and Buildings*, *57*, 361-372. doi: 10.1016/j.enbuild.2012.10.044

Aswani, A., Master, N., Taneja, J., Culler, D., & Tomlin, C. (2012). Reducing transient and steady state electricity consumption in HVAC using learning- based model-predictive control. *Proceedings of the IEEE* , *100* (1), 240-253.

Behl, M., Nghiem, T., & Mangharam, R. (2014). IMpACT: Inverse model accuracy and control performance toolbox for buildings. In *IEEE International Conference on Automation Science and Engineering* (p. 1109-1114). Taipei. doi: 10.1109/CoASE.2014.6899464

Bengea, S. C., Kelman, A. D., Borrelli, F., Taylor, R., & Narayanan, S. (2014). Implementation of model predictive control for an HVAC system in a mid-size commercial building. *HVAC & R Research*, *20*(1), 121-135. doi: 10.1080/10789669.2013.834781

Billings, S. A. (2013). *Nonlinear System Identification: NARMAX Methods in the Time, Frequency, and Spatio-temporal Domains*. Hoboken, NJ: John Wiley.

Bordass, B., Leaman, A., & Ruyssevelt, P. (2001). Assessing building performance in use 5: Conclusions and implications. *Building Research & Information*, *29* (2), 144-157. doi: 10.1080/09613210010008054

Box, G. E. P., Jenkins, G. M., & Reinsel, G. C. (1994). *Time Series Analysis: Forecasting and Control* (3rd ed.). Englewood Cliffs, N.J.: Prentice Hall.Bunn, R. (2014). *Soft Landings Framework* (Tech. Rep. No. BG 54/2014). BSRIA.

Camacho, E. F., & Bordons, C. (2004). *Model Predictive Control* (2nd ed.). London: Springer.

Camponogara, E., Jia, D., Krogh, B., & Talukdar, S. (2002). Distributed model predictive control. *IEEE Control Systems Magazine*, *22* (1), 44-52. doi: 10.1109/37.980246

Carbon Trust. (2012). *Closing the Gap: Lessons Learned on Realising the Potential of Low-Carbon Building Design* (Tech. Rep.). London: Carbon Trust.

Chappells, H., & Shove, E. (2005). Debating the future of comfort: Environmental sustainability, energy consumption and the indoor environment. *Building Research and Information*, *33* (1), 32-40. doi: 10.1080/0961321042000322762

CIBSE. (2008). *TM46: Energy Benchmarks*. London: Chartered Institute of Building Services Engineers.

CIBSE. (2009). *CIBSE Guide H: Building Control Systems*. London: Chartered Institute of Building Services Engineers.

CIBSE. (2013). *TM52 The Limits of Thermal Comfort: Avoiding Overheating in European Buildings*. London: Chartered Institute of Building Services Engineers.

Cigler, J., & Prívara, S. (2010). Subspace identification and model predictive control for buildings. In *11th International Conference on Control, Automation, Robotics and Vision (ICARCV2010)* (p. 750-755). Singapore. doi: 10.1109/ICARCV.2010.5707821

Cigler, J., Prívara , S., Váňa, Z., Žáčeková, E., & Ferkl, L. (2012). Optimization of predicted mean vote index within model predictive control framework: Computationally tractable solution. *Energy and Buildings*, *52*, 39-49. doi: 10.1016/j.enbuild.2012.05.022

Cigler, J., Široký, J., Korda, M., & Jones, C. N. (2013). On the selection of the most appropriate MPC problem formulation for buildings. In *11th REHVA World Congress (CLIMA 2013)*. Prague, Czech Republic.

Cole, R. J., Robinson, J., Brown, Z., O'Shea, M., (2008) Re-contextualising the notion of comfort. *Building Research and Information*. *36* (4), 323-336

Coley, D., & Penman, J. (1992). Second order system identification in the thermal response of real buildings. Paper II: Recursive formulation for on-line building energy management and control. *Building and Environment* , *27* (3), 269-277. doi: 10.1016/0360-1323(92)90028-N

Conner, J., & Seborg, D. (2005). Assessing the need for process re-identification. *Industrial & Engineering Chemistry Research*, *44* (8), 2767-2775. doi: 10.1021/ie049439g

Crawley, D., Pedersen, C., Lawrie, L., & Winkelmann, F. (2000). EnergyPlus: Energy simulation program. *ASHRAE Journal* , *42* (4), 49-56.

de Dear, R. & Brager, G. (1998). Towards an adaptive model of thermal comfort and preference. *ASHRAE Transactions 104* (1), 145-167.

European Union. (2010). *DIRECTIVE 2010/31/EU OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 19 May 2010 on the Energy Performance of Buildings. Official Journal of the European Union, L153/13. doi: 10.3000/17252555.L 2010.153.eng*

Ferkl, L., & Široký, J. (2010). Ceiling radiant cooling: Comparison of ARMAX and subspace identification modelling methods. *Building and Environment*, *45*(1), 205-212. doi: 10.1016/j.buildenv.2009.06.004

Ferreira, P., Ruano, A., Silva, S., & Conceição, E. (2012). Neural networks based predictive control for thermal comfort and energy savings in public buildings. *Energy and Buildings*, *55*(0), 238-251. doi: 10.1016/j.enbuild.2012.08.002

Florita, A. R., & Henze, G. P. (2009). Comparison of short-term weather forecasting models for model predictive control. *HVAC & R Research*, *15*(5), 835-853. doi: 10.1080/10789669.2009.10390868

Genceli, H., & Nikolaou, M. (1996). New approach to constrained predictive control with simultaneous model identification. *AIChE Journal*, *42*(10), 2857-2868. doi: 10.1002/aic.690421015

Gevers, M. (2005). Identification for control: From the early achievements to the revival of experiment design. *European Journal of Control* , *11* (4-5), 335-352. doi: 10.3166/ejc.11.335-352

Gunay, H. B., Bursill, J., Huchuk, B., O'Brien, W., & Beausoleil-Morrison, I. (2014). Shortest-prediction-horizon model-based predictive control for individual offices. *Building and Environment*, *82* , 408-419. doi: 10.1016/j.buildenv.2014.09.011

Gunay, H. B., O'Brien, W., & Beausoleil-Morrison, I. (2013). A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices. *Building and Environment*, *70* , 31-47. doi: 10.1016/j.buildenv.2013.07.020

Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, *3* , 1157-1182.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). New York: Springer-Verlag.

Hathway, A., Rockett, P., & Carpenter, P. (2013). The use of machine learning in model predictive control for the operation of non-domestic buildings. In *Futurebuild 2013.* Bath, UK.

Hazyuk, I., Ghiaus, C., & Penhouet, D. (2012). Optimal temperature control of intermittently heated buildings using model predictive control: Part II – Control algorithm. *Building and Environment*, *51* , 388-394. doi: 10.1016/j.buildenv.2011.11.008

Hellwig, R. T. (2015). Perceived control in indoor environments: A conceptual approach, *Building Research and Information*, 43(3), 302-315. doi=10.1080/09613218.2015.1004150..

Henze, G. P. (2013). Editorial – Model predictive control for buildings: A quantum leap? *Journal of Building Performance Simulation*, *6* (3), 157- 158. (Special Issue on Model Predictive Control for Buildings) doi: 10.1080/19401493.2013.778519

Hjalmarsson, H. (2009). System identification of complex and structured systems. *European Journal of Control* , *15* (3-4), 275-310. doi: 10.3166/ejc.15.275-310

Hou, Z.-S., & Wang, Z. (2013). From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, *235* (0), 3-35. doi: 10.1016/j.ins.2012.07.014

Huang, G. (2011). Model predictive control of VAV zone thermal systems concerning bilinearity and gain nonlinearity. *Control Engineering Practice*, *19*(7), 700-710. doi: 10.1016/j.conengprac.2011.03.005

Hussain, M. A. (1999). Review of the applications of neural networks in chemical process control – simulation and online implementation. *Artificial Intelligence in Engineering* , *13* (1), 55-68. doi: 10.1016/S0954-1810(98)00011-9

Lamoudi, M. Y., Alamir, M., & Béguery, P. (2012). Model predictive control for energy management in buildings. Part 2: Distributed model predictive control. In *4th IFAC Nonlinear Model Predictive Control Conference* (p. 226-231). Leeuwenhorst, Netherlands. doi: 10.3182/20120823-5-NL-3013.00036

Landau, I. (1999). From robust control to adaptive control. *Control Engineering Practice*, *7* (9), 1113-1124. doi: 10.1016/S0967-0661(99)00076-3

Levermore, G. J. (2000). *Building Energy Management Systems: Applications to Low-energy HVAC and Natural Ventilation Control* (2nd ed.). London: E and F N Spon.

Ljung, L. (1999). *System Identification: Theory for the User* (2nd ed.). Upper Saddle River, N.J.: Prentice Hall.

Ljung, L. (2010). Perspectives on system identification. *Annual Reviews in Control* , *34* (1), 1-12. doi: 10.1016/j.arcontrol.2009.12.001

Low Carbon Innovation Coordination Group. (2012). *Technology Innovation Needs Assessment - Non-domestic Building Summary Report* (Tech. Rep.). Low Carbon Innovation Coordination Group.

Ma, J., Qin, J., Salsbury, T., & Xu, P. (2012). Demand reduction in building energy systems based on economic model predictive control. *Chemical Engineering Science*, *67*(1), 92-100. doi: 10.1016/j.ces.2011.07.052

Ma, Y., Kelman, A., Daly, A., & Borrelli, F. (2012). Predictive control for energy efficient buildings with thermal storage: Modeling, simulation, and experiments. *IEEE Control Systems*, *32* (1), 44-64. doi: 10.1109/MCS.2011.2172532

Marafioti, G., Bitmead, R. R., & Hovd, M. (2014). Persistently exciting model predictive control. *International Journal of Adaptive Control and Signal Processing*, *28* (6), 536-552. doi: 10.1002/acs.2414

Myhren, J., & Holmberg, S. (2009). Design considerations for ventilation-radiators: Comparisons to traditional two-panel radiators. *Energy and Buildings*, *41* (1), 92-100. doi:10.1016/j.enbuild.2008.07.014

Nelles, O. (2001). *Nonlinear System Identification: From Classical Approaches to Neural Networks and Fuzzy Models*. Berlin: Springer.

Nishiguchi, J., Konda, T., & Dazai, R. (2010). Data-driven optimal control for building energy conservation. In *Sice Annual Conference* (p. 116-120). Taipei.

Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., . . . Morari, M. (2012). Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, *45*, 15-27. doi: 10.1016/j.enbuild.2011.09.022

Oldewurtel, F., Sturzenegger, D., & Morari, M. (2013). Importance of occupancy information for building climate control. *Applied Energy*, *101*, 521-532. doi:10.1016/j.apenergy.2012.06.014

Page, J., Robinson, D., Morel, N., & Scartezzini, J.-L. (2008). A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings*, *40* (2), 83-98. doi: 10.1016/j.enbuild.2007.01.018

Pattarello, G., Wei, L., Ebadat, A., Wahlberg, B., & Johansson, K. H. (2013). The KTH open testbed for smart HVAC control. In *5th ACM Workshop on Embedded Systems for Energy-efficient Buildings (BuildSys'13)* (p. 34:1-34:2). Rome, Italy. doi: 10.1145/2528282.2534154

Prívara, S., Cigler, J., Váňa, Z., Oldewurtel, F., Sagerschnig, C., & Žáčeková, E. (2013). Building modeling as a crucial part for building predictive control. *Energy and Buildings*, *56*(0), 8-22. doi: 10.1016/j.enbuild.2012.10.024

Prívara, S., Široký, J., Ferkl, L., & Cigler, J. (2011). Model predictive control of a building heating system: The first experience. *Energy and Buildings*, *43*(2-3), 564-572. doi: 10.1016/j.enbuild.2010.10.022

Rathouský J., & Havlena, V. (2013). MPC-based approximate dual controller by information matrix maximization. *International Journal of Adaptive Control and Signal Processing* , *27* (11), 974-999. doi: 10.1002/acs.2370

Riederer, P., Marchio, D., & Visier, J. C. (2002). Influence of sensor position in building thermal control: Criteria for zone models. *Energy and Buildings*, *34*(8), 785-798. doi: 10.1016/S0378-7788(02)00097-X

Rivera, D. E., Lee, H., Braun, M. W., & Mittelmann, H. D. (2003). "Plant friendly" system identification: A challenge for the process industries. In *13th IFAC Symposium on System Identification (SYSSID 2003)*. Rotterdam, Netherlands.

Söderström, T., & Stoica, P. (1989). *System Identification*. New York: Prentice Hall.

Sotomayor, O. A., Odloak, D., & Moro, L. F. (2009). Closed-loop model reidentification of processes under MPC with zone control. *Control Engineering Practice*, *17* (5), 551-563. doi: 10.1016/j.conengprac.2008.10.002

Sturzenegger, D., Gyalistras, D., Gwerder, M., Sagerschnig, C., Morari, M., & Smith, R. (2013). Model predictive control of a Swiss office building. In *Clima - RHEVA World Congress* (p. 3227-3236). Prague, Czech Republic.

Sturzenegger, D., Gyalistras, D., Morari, M., & Smith, R. S. (2012). Semi-automated modular modeling of buildings for model predictive control. In *4th ACM Workshop on Embedded Sensing Systems for Energy- Efficiency in Buildings (BuildSys '12)* (p. 99-106). Toronto, Canada. doi: 10.1145/2422531.2422550

Subrt, L., & Pechac, P. (2012). Intelligent walls as autonomous parts of smart indoor environments. *IET Communications*, *6* (8), 1004-1010. doi: 10.1049/iet- com.2010.0544

Tanabe, S., Iwahashi, Y., Tsushima, S., & Nishihara, N. (2013). Thermal comfort and productivity in offices under mandatory electricity savings after the Great East Japan earthquake. *Architectural Science Review* , *56* (1), 4-13.

Tanner, R. A., & Henze, G. P. (2014). Stochastic control optimization for a mixed mode building considering occupant window opening behaviour. *Journal of Building Performance Simulation*, *7* (6), 427-444. doi: 10.1080/19401493.2013.863384

Váňa, Z., Cigler, J., Široký, J., Žáčeková, E., & Ferkl, L. (2014). Model-based energy efficient control applied to an office building. *Journal of Process Control*, *24* (6), 790-797. doi: 10.1016/j.jprocont.2014.01.016

Váňa, Z., Kubeček, J., & Ferkl, L. (2010). Notes on finding black-box model of a large building. In *IEEE International Conference on Control Applications (CCA 2010)* (p. 1017-1022). Yokohama, Japan. doi: 10.1109/CCA.2010.5611157

Žáčeková, E., Váňa, Z., Hoogmartens, J., Verhelst, C., Sourbron, M., Ferkl, L., & Helsen, L. (2013). Identification for model based predictive control applied to an office building with a thermally activated building systems. In *11th REHVA World Congress (Clima 2013)* (p. 2430-2439). Prague, Czech Republic.

West, S. R., Ward, J. K., & Wall, J. (2014). Trial results from a model predictive control and optimisation system for commercial building HVAC. *Energy and Buildings*, *72* (0), 271-279. doi: 10.1016/j.enbuild.2013.12.037

Zhu, Y., & Butoyi, F. (2002). Case studies on closed-loop identification for MPC. *Control Engineering Practice*, *10* (4), 403-417. doi: 10.1016/S0967-0661(02)00007-2

---

1 Hereafter, we use the term ‘building’ to mean a non-domestic building unless otherwise stated.

2 . . . although it should be noted – as we discuss below – that a great many implementations of MPC in buildings to date have minimised only temperature deviation from a scheduled setpoint.

3 See http://www.opticontrol.ethz.ch/

4 We neither make nor imply any criticism of such consultancy operations. Indeed, they provide a valuable service to their clients by delivering energy savings, but do so within the constraints of current BMSs.

5 Ferkl (L. Ferkl, Private communication, 2014) reports that a building at the Czech Technical University, Prague has been continuously controlled by MPC for six years, and that an office building in Belgium has been

similarly controlled for two years; the experience on these extended trials, however, has yet to be published.

6 http://www.ibpsa.us/mpc2011/Post_Workshop_Roundtable_Discussion.jpg

7 See, for example, http://buildingrobotics.com/