Sequential Monte Carlo for on-line parameter estimation of a lumped building energy model

Simon Rouchier[a] **, Maria José Jiménez[b], Sergio Castaño[b]

[a]Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, LOCIE, 73000 Chambéry, France
[b]CIEMAT, Department of Energy, Energy Efficiency in Buildings Unit, Av. Complutense 22, E28040 Madrid, Spain

Rouchier S., Jiménez M.J., Castaño S. (2019) Sequential Monte Carlo for on-line parameter estimation of a lumped building energy model, *Energy and Buildings*, vol. 187, p. 86-94

Abstract

The characterisation of parameters of building energy models based on in-situ sensor information is generally performed after the measurement period, using all data in a single batch. Alternatively, on-line parameter estimation proposes updating a model every time a new data point is available: this establishes a direct link between external events, such as the weather, and the identifiability of parameters. The present study uses the Sequential Monte Carlo method to train a lumped building energy model (RC model), and thus estimate a Heat Loss Coefficient, and other parameters, sequentially. Results show the direct impact of solicitations (solar irradiance and indoor heat input) on this estimation, in real time. The method is validated by comparing its results with the Metropolis-Hastings algorithm for off-line parameter estimation.

Keywords

Bayesian calibration; on-line; Sequential Monte Carlo; building energy simulation

1 Introduction

The characterisation of parameters of simplified building energy models using in-situ measurements is now a widespread research topic [1] and is commonly performed for two general types of applications. The first is parameter estimation, where the aim is to identify quantities that can be directly interpreted. This includes the characterisation of intrinsic building performance by energy signature models [2, 3], co-heating tests [4] or other dynamic methods [5, 6, 7, 8]. The second use of in-situ measurements is the calibration of a model for predictive purposes [9, 10, 11], for instance in the aim of model predictive control [12, 13, 14] or building energy management. The aim is to build a model that correctly reproduces the energy dynamics of buildings, regardless of the physical meaning of its parameters. Although grey-box models may be suitable for this purpose, black-box models with no consideration of physics may also be appropriate: the former are less reliant on data availability but the latter are more flexible.

State-space models, which include the simplified resistor-capacitor (RC) model structures, are a popular choice for both of these applications. When written as a set of Stochastic Differential Equations, they allow

**Corresponding author. Email: simon.rouchier@univ-smb.fr

accounting for modelling approximations [15, 16, 17, 18] and offer a more reproducible parameter estimation than deterministic models that overlook modelling errors [19].

Parameter estimation is typically performed *off-line*: measurements of indoor and outdoor conditions are first carried in a test building, and data is processed after the experiment in a single batch. An interesting challenge is to carry parameter estimation *on-line*, during the observation period: starting from an initial guess for parameter values, these estimates are updated sequentially, every time a new observation becomes available.

There are several motivations for this: first, it would allow using the measurement period for computations, thus reducing the total time of the procedure [20]. It would also be a way to make use of the emerging wireless energy monitoring technology: smart meters, wireless sensor networks, etc. With frequent data collection and remote transfer, either off-line and on-line analysis can be performed non-intrusively during the monitoring period. The off-line alternative however requires restarting calculations from the beginning of the measurement period, which can become problematic if a frequent update on parameter estimates is expected. A second, more important advantage of on-line estimation lies in the amount of information gained from the experiment. Parameter estimates are to be updated after every new observation: this will allow directly observing which phenomena "bring information" to the parameters, by correlating the reduction in their estimation uncertainty with observed events. Such a thorough diagnosis can be useful for fault detection as well, particularly when the necessary length of the measurement period is not known a priori.

Bayesian inference offers the possibility of on-line estimation with Sequential Monte-Carlo (SMC) methods [21]. Originally developed for the sequential estimation of states [22], SMC was later adapted to state and parameter estimation [23, 24]. Building physics applications are scarce and very recent [20], but may become more common due to the motivations listed above.

The present paper applies SMC for the on-line estimation of the heat loss coefficient (HLC) of a test cell. Starting from a highly uncertain prior knowledge of HLC, the target is to dynamically observe what leads its estimation to narrow down to a more precise value. The identifiability of HLC regarding available data is then discussed. Sec. 2 presents the test cell and Sec. 3 the RC model chosen to simulate it. Sec. 4 shortly describes the Metropolis-Hastings and SMC algorithms for off-line and on-line Bayesian parameter estimation. Results are then showed and discussed on Sec. 5.

Thermophysical parameters and time series		State-space representation	
R_1, R_2, R_3	Heat resistance	$\mathbf{x}$	Vector of states
C_1, C_2	Heat capacitance	$\mathbf{P}$	State variance
k_1, k_2	Solar aperture coefficient	$\mathbf{y}$	Observations
T_i	Indoor temperature	$\mathbf{w}$	Process noise
T_e	Envelope temperature	$\mathbf{v}$	Measurement noise
T_a	Ambient (outdoor) temperature	$\mathbf{A}$	State matrix (continuous)
q	Indoor heating power	$\mathbf{B}$	Input matrix (continuous)
I_{sol}	Solar irradiance	$\mathbf{C}$	Observation matrix (continuous)
HLC	Heat loss coefficient	$\mathbf{F}_\theta$	State matrix (discrete)
Filtering and estimation		$\mathbf{G}_\theta$	Input matrix (discrete)
ε	Prediction error mean	$\mathbf{H}_\theta$	Observation matrix (discrete)
$\mathbf{\Sigma}$	Prediction error covariance	$\mathbf{Q}$	Model error covariance matrix
$\mathbf{K}$	Kalman gain	$\mathbf{R}$	Measurement error covariance matrix
L_y	Likelihood function	Subscripts and superscripts	
θ	Parameter	c	Continuous
$p(\theta)$	Prior function	d	Discrete
$p(\theta\|\mathbf{y})$	Posterior function	t	Time coordinate
g	Proposal distribution	(j)	Particle index
ω	Particle weight		

Table 1: Nomenclature

2 Case study

The present study uses measurements that were carried in the Round Robin Test Box (RRTB), within the framework of the IEA EBC Annex 58 [25]. This experimental test cell, shown by Fig. 1 has a cubic form, with exterior dimensions of 120x120x120 cm^3. The floor, roof and wall components of the box are all identical and have a thickness of 12 cm, resulting in an inner volume of 96x96x96cm^3. One wall contains an operable wooden window with overall dimensions of 71x71 cm^2 and a glazed part of 52x52 cm^2. The double glazing has a U-value of 1.1 W/m^2K and g-value of 0.63. Numerical simulation [25] has estimated the overall HLC of the box to a target value of 4.08 W/K, assuming constant standard surface heat transfer coefficients. This value presents an uncertainty in the range 3.49-4.14 W/K due to variations produced by the presence of a thin air or glue layer between the different material layers, or approximations in the estimation of surface heat transfer coefficients depending on wind and surface temperature. This range will serve as reference to check the validity of the results below. The total solar aperture of the box was estimated at 0.162 m^2.

Figure 1: Round Robin Test Box

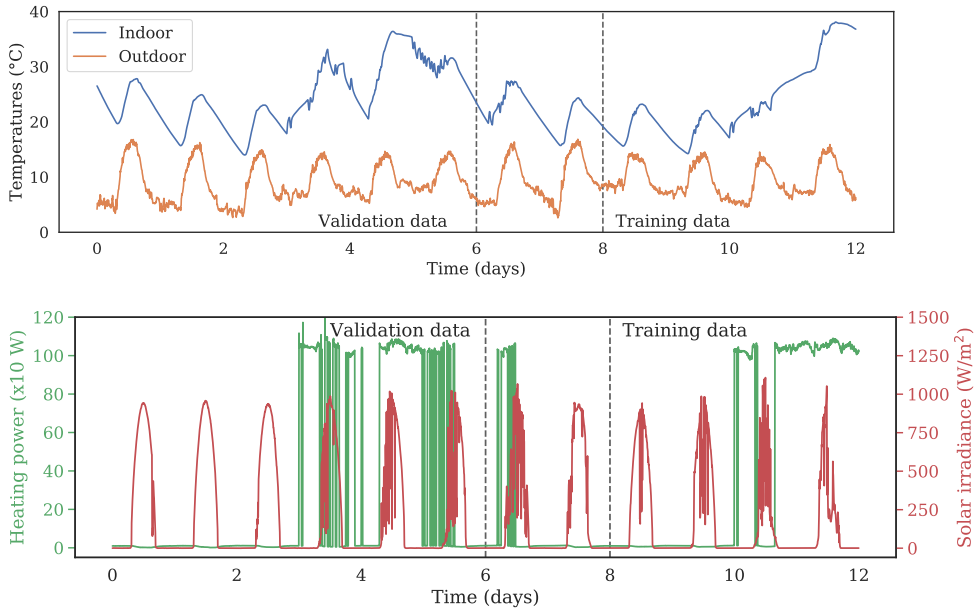

Figure 2: Measurements of indoor and outdoor temperature, heating power and solar irradiance

The test box was installed outdoors, in the LECE laboratory at Plataforma Solar de Almeria, in the South East of Spain. Experiments were carried during a 43-days period in the winter of 2013-2014. Measurements used in this study are: indoor temperature (which is the average of two type T thermocouples placed inside the box), outdoor air temperature, heating power, and global horizontal solar irradiance. The box is also equiped with sensors that were not used here: internal and external surface temperatures of each side, heat flow meters, diffuse horizontal solar irradiance, wind speed and direction, relative humidity, and horizontal and vertical long wave radiation from the sky. All sensor types are listed in [25]. All measurements were received with a sampling time of 1 min, but were then resampled to a time step of 5 min in order to reduce calculation time without compromising precision.

A period of 12 days was chosen for the present investigation, starting from the 6$^{\text{th}}$ of December 2013 at 00:00. Measurements are shown in Fig. 2. The first 6 days will be used as validation data for the trained model, and the last 4 days as training data. The measured indoor temperature during the validation period will be compared to the output of the models calibrated with the training data. This particular partition of the original 12-days dataset is motivated by the following reasons:

- Both the training and the validation dataset comprise a period of free-floating indoor temperature, and a period of controlled indoor heat input. In terms of model calibration, these boundary conditions are not very informative at first, then become more informative: we expect to witness their effects on the evolution of the estimation of the heat loss coefficient.
- The training and validation sets are separated by a short "buffer period" in order to make them relatively independent from each other. By this precaution, we want to avoid a correlation between both datasets, that would not guarantee that the trained model is generalizable.

3 Modelling

3.1 Model

In order to estimate its heat loss coefficient (HLC) and other properties, the test box is represented by a lumped Resistor-Capacitance model. It is a 3R2C model described by:

$$\begin{bmatrix}\dot{T}_i(t)\\ \dot{T}_e(t)\end{bmatrix} = \underbrace{\begin{bmatrix}-\frac{1}{R_1C_1}-\frac{1}{R_3C_1} & \frac{1}{R_1C_1}\\ \frac{1}{R_1C_2} & -\frac{1}{R_1C_2}-\frac{1}{R_2C_2}\end{bmatrix}}_{\mathbf{A}}\begin{bmatrix}T_i(t)\\ T_e(t)\end{bmatrix} + \underbrace{\begin{bmatrix}\frac{1}{R_3C_1} & \frac{1}{C_1} & \frac{k_1}{C_1}\\ \frac{1}{R_2C_2} & 0 & \frac{k_2}{C_2}\end{bmatrix}}_{\mathbf{B}}\begin{bmatrix}T_a(t)\\ q(t)\\ I_{sol}(t)\end{bmatrix} + \mathbf{w}(t) \quad (1)$$

$$\mathbf{y}(t) = \underbrace{\begin{bmatrix}1 & 0\end{bmatrix}}_{\mathbf{C}}\begin{bmatrix}T_i(t)\\ T_e(t)\end{bmatrix} + \mathbf{v}(t) \quad (2)$$

where T_i, T_a and T_e are the indoor, ambient (outdoor) and envelope temperatures. The envelope temperature is associated with the thermal mass of the opaque surfaces, and does not represent a specific coordinate within the envelope. The model has two states T_e (unobserved) and T_i (observed, shown in Fig. 2(a)); q (W) is the indoor heating power; I_{sol} (W/m^2) is the solar irradiance on a southern vertical plane. A schematic view of the 3R2C model is shown in Fig. 3. The choice of this model structure is motivated by simplicity: in a previous study [19], a 2R2C model was judged sufficient to describe the dynamics of a very simple mono-zone building. The 3R2C model is an extension of this model, applied to the RRTB where the influence of the window may be significant.

In the continuous state equation (Eq. 1), $\mathbf{w}(t)$ denotes a Wiener process that represents modelling errors with an incremential covariance $\mathbf{Q}_c$ [15], and $\mathbf{v}(t)$ is the measurement error of the indoor temperature, normally distributed white noise with zero mean and variance $\mathbf{R}_c$. The coefficients of the $\mathbf{Q}_c$ matrix and $\mathbf{R}_c$ are considered unknown and will be estimated along with the other parameters of the model.

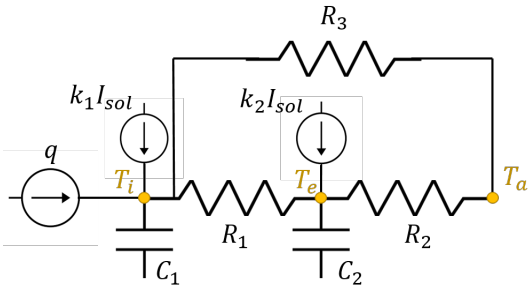

Figure 3: 3R2C model

The 3R2C model has 8 parameters that enable a physical interpretation of the RRTB: R_1 and R_2 (K/W) are two thermal resistances representing heat transfer through the opaque walls; R_3 is a resistance directly linking the outdoor and indoor temperatures, representing heat transfer through the window; C_1 and C_2 (J/K) are thermal capacities and k_1 and k_2 (m^2) are two solar aperture coefficients, one for each state of the model. The last parameter is the initial envelope temperature $T_e(0)$: since T_e is an unobserved state, its initial value is unknown. In the following, we denote as θ the vector of these parameters.

The stochastic model described by Eq. 1 must be discretized in order to specify its evolution between discrete time coordinates. Let us denote the sample interval length Δt and assume that the inputs $\mathbf{u}(t) = [T_a(t)\ q(t)\ I_{sol}(t)]$ are constant during this interval. Eq. 1 and 2 can be discretized into the following discrete linear system:

$$\mathbf{x}_t = \mathbf{F}_\theta\, \mathbf{x}_{t-1} + \mathbf{G}_\theta\, \mathbf{u}_t + \mathbf{w}_t \tag{3}$$

$$\mathbf{y}_t = \mathbf{H}_\theta\, \mathbf{x}_t + \mathbf{v}_t \tag{4}$$

where $\mathbf{x}_t$ denotes the vector of states at the time coordinate t, and $\mathbf{y}_t$ denotes the observations. The $\mathbf{F}_\theta$ and $\mathbf{G}_\theta$ matrices of the discrete equation result from the matrices of the continuous equation 1 using the usual state-space discretization method. Their coefficients are functions of θ and of the time step size Δt. Similarly, the process noise in discrete time $\mathbf{w}_t \sim \mathcal{N}(0, \mathbf{Q}_d)$ has a covariance matrix $\mathbf{Q}_d$ that can be calculated from the covariance matrix of the process noise in continuous time $\mathbf{w}(t) \sim \mathcal{N}(0, \mathbf{Q}_c)$. The observation error $\mathbf{v}_t$ has a covariance $\mathbf{R}_d$, which depends on $\mathbf{R}_c$ and the time step size.

The discretization equations are given here, and are available with more detail in [15, 1]:

$$\mathbf{F}_\theta = \exp(\mathbf{A}\,\Delta t) \tag{5}$$

$$\mathbf{G}_\theta = \mathbf{A}^{-1}\,(\mathbf{F}_\theta - \mathbf{I})\,\mathbf{B} \tag{6}$$

$$\mathbf{H}_\theta = \mathbf{C} \tag{7}$$

$$\mathbf{Q} = \int_0^{\Delta t} \exp(\mathbf{A}\,\Delta t)\ \mathbf{Q}_d \exp\left(\mathbf{A}^T\,\Delta t\right) \mathrm{d}t \tag{8}$$

$$\mathbf{R}_d = \frac{1}{\Delta t}\mathbf{R}_c \tag{9}$$

where $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ are the matrices of the continuous linear system 1 and 2.

3.2 Kalman filter

Let us first suppose that the parameters θ of the system are known, and a sequence of output observations $\mathbf{y}_{1:T}$ (in our case the indoor temperature) and input variables $\mathbf{u}_{1:T}$ (outdoor temperature, heating power and solar irradiance) has been obtained.

Given a state transition probability $p(\mathbf{x}_t|\theta, \mathbf{x}_{t-1}, \mathbf{u}_t)$ (Eq. 3) and an observation probability $p(\mathbf{y}_t|\mathbf{x}_t)$ (Eq. 4), a Kalman filter produces $p(\mathbf{x}_t|\mathbf{y}_{1:T}, \theta)$, the probability distribution function of each state $\mathbf{x}_t$ given measurements and parameter values, and the marginal likelihood function $L_y(\theta) = p(\mathbf{y}_{1:T}|\theta)$. Its algorithm has been described by many authors including [15, 1] and is shortly recalled here.

In the following, definitions adapted from [26] are used: $\mathbf{x}_{t|s}$ is the expected state at time t given observations up to time s. $\mathbf{P}_{t|s}$ is the variance of the state $\mathbf{x}_t$, i.e. the mean-squared error.

$$\mathbf{x}_{t|s} = \mathrm{E}(\mathbf{x}_t|\mathbf{y}_{1:s}, \theta) \tag{10}$$

$$\mathbf{P}_{t|s} = \mathrm{Var}(\mathbf{x}_t|\mathbf{y}_{1:s,\theta}) = \mathrm{E}\left[(\mathbf{x}_t - \mathbf{x}_{t|s})(\mathbf{x}_t - \mathbf{x}_{t|s})^T|\mathbf{y}_{1:s}, \theta\right] \tag{11}$$

The Kalman filter algorithm is described here and illustrated by Fig. 4:

- Set the initial states $\mathbf{x}_{0|0}$ and their covariance $\mathbf{P}_{0|0}$
- for $t = 1...T$:
  1. **Prediction step**: given the previous state $\mathbf{x}_{t|t}$ and its covariance $\mathbf{P}_{t|t}$, the model estimates the one-step ahead prediction.

$$\mathbf{x}_{t+1|t} = \mathbf{F}_\theta\, \mathbf{x}_{t|t} + \mathbf{G}_\theta\, \mathbf{u}_{t+1} \tag{12}$$

$$\mathbf{P}_{t+1|t} = \mathbf{F}_\theta\, \mathbf{x}_{t|t}\, \mathbf{F}_\theta^T + \mathbf{Q} \tag{13}$$

  2. **Innovations** (prediction error) ε_{t+1} and their covariances $\mathbf{\Sigma}_{t+1}$ are then calculated, along with the Kalman gain $\mathbf{K}_{t+1}$, by comparing **measurements** $\mathbf{y}_{t+1}$ (see Fig. 4) with the one-step ahead prediction $\mathbf{x}_{t+1|t}$:

$$\varepsilon_{t+1} = \mathbf{y}_{t+1} - \mathbf{H}_\theta\, \mathbf{x}_{t+1|t} \tag{14}$$

$$\mathbf{\Sigma}_{t+1} = \mathbf{H}_\theta\, \mathbf{P}_{t+1|t}\, \mathbf{H}_\theta^T + \mathbf{R} \tag{15}$$

$$\mathbf{K}_{t+1} = \mathbf{P}_{t+1|t}\, \mathbf{H}_\theta^T\, \mathbf{\Sigma}_{t+1}^{-1} \tag{16}$$

  3. **Updating step**: the new states at time $t+1$ are updated, as a compromise between the one-step ahead prediction and the measurement.

$$\mathbf{x}_{t+1|t+1} = \mathbf{x}_{t+1|t} + \mathbf{K}_{t+1}\, \varepsilon_{t+1} \tag{17}$$

$$\mathbf{P}_{t+1|t+1} = (\mathbf{I} - \mathbf{K}_{t+1}\, \mathbf{H}_\theta)\, \mathbf{P}_{t+1|t} \tag{18}$$

- The total (negative) log-likelihood can be calculated up to a normalizing constant:

$$-\ln L_y(\theta) = \frac{1}{2}\sum_{t=1}^{T} \ln|\mathbf{\Sigma}_t(\theta)| + \frac{1}{2}\sum_{t=1}^{T} \varepsilon_t(\theta)^T\, \mathbf{\Sigma}_t(\theta)^{-1}\, \varepsilon_t(\theta) \tag{19}$$

This standard Kalman filter algorithm works for linear systems only. Non-linear systems require another filter, such as the Extended Kalman Filter (used by [16]), the Unscented Kalman Filter [27], or the particle filter.

4 Bayesian parameter estimation

The goal of the on-line parameter estimation exercise is to assess the value of all static parameters of the model, at each time coordinate of the measurement period: the expected output is a sequence of posterior distributions $\{p(\theta|\mathbf{y}_{1:t}), t \in 1...T\}$, where T is the number of data points in the measurement period. This sequential estimation is performed by the SMC algorithm. In order to validate the results of this method, the estimation has also been carried in a "traditional" off-line fashion with the Metropolis-Hastings algorithm. Both methods are described below.

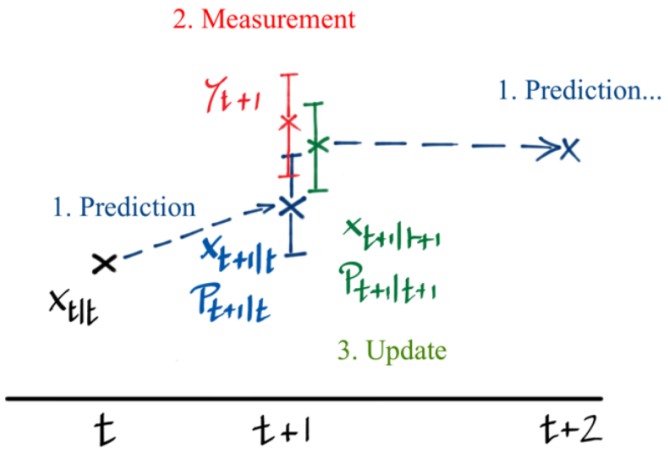

Figure 4: Schematic view of one iteration of the Kalman filter

4.1 Off-line parameter estimation: Metropolis Hastings

The target of off-line Bayesian parameter estimation is the estimation of the posterior distribution $p\left(\theta|\mathbf{y}_{1:T}\right)$, which is the probability of the parameters θ *given* a batch of data $\mathbf{y}_{1:T}$. Bayes equation reads:

$$p\left(\theta|\mathbf{y}_{1:T}\right) \propto L_y(\theta)\, p(\theta) \tag{20}$$

where $p(\theta)$ is the prior over the parameter values and $L_y(\theta)$ is the likelihood function, calculated by Eq. 19 in the case of a linear state-space model. The prior allows accounting for expert knowledge; an informative prior contributes to regularizing the inverse problem of parameter estimation.

Algorithm 1 Metropolis Hastings algorithm

for $n = 1...N$ **do**
  Draw a new value from a proposal distribution g
  $\theta' \leftarrow g\left(\theta'|\theta_{n-1}\right)$
  Compute the marginal likelihood using a Kalman filter (for linear models):
  $\left(p\left(\mathbf{x}_{1:T}|\theta, \mathbf{y}_{1:T}\right), L_y(\theta')\right) \leftarrow \textsc{KalmanFilter}(\theta')$
  Accept or reject the proposal:
  $\alpha \sim \mathrm{U}(0,1)$
  if $\alpha \leq \dfrac{L_y(\theta')\, p(\theta')\, g(\theta_{n-1}|\theta')}{L_y(\theta_{n-1})\, p(\theta_{n-1})\, g(\theta'|\theta_{n-1})}$ **then**
    $\theta_n \leftarrow \theta'$
  else
    $\theta_n \leftarrow \theta_{n-1}$
  end if
end for

The Marginal Metropolis Hastings (MMH) algorithm is part of the family of Markov Chain Monte Carlo (MCMC) methods. It calculates a finite sequence of samples $\{\theta_n, n \in 1...N\}$ approximating the posterior distribution. Algorithm 1 employs a Kalman filter to compute the states $p\left(\mathbf{x}_{1:T}|\theta, \mathbf{y}_{1:T}\right)$ and likelihood $L_y(\theta)$ associated to each proposal for θ. If the state-space model (Eq. 3) is non-linear, this filter can be replaced by a particle filter: this approach is known as Particle Markov Chain Monte Carlo (PMCMC) [28].

The choice of the proposal disribution g, and a good initialisation, are critical for the performance of the algorithm. A burn-in phase at the beginning of the Markov chain must be discarded as it does not

reflect the posterior distribution. [29] construct the proposal distribution by using the gradient and Hessian of the posterior, calculated by differentiation of the Kalman filter equations. Alternatively, the Adaptive Metropolis Hastings algorithm is used by [19].

4.2 On-line estimation: Sequential Monte Carlo

We now consider the procedure for on-line parameter estimation. The target is to construct a sequence of posterior distributions $\{p(\theta|\mathbf{y}_{1:t}), t \in 1...T\}$, one for each observation point, that will allow us to visualize the information gained during the experiment in real time.

The SMC algorithm for parameter estimation is an adaptation of particle filtering for state variables. The foundation of this method is the Importance Sampling paradigm as described by [30]: simulating samples under an instrumental distribution and then approximating the target distributions by weighting these samples using appropriately defined importance weights. The reader is referred to [30] and [24] for a deeper explanation of SMC and its application to parameter estimation. The method used here is inspired from the Iterated Batch Importance Sampling algorithm [31]. It is described in Fig. 5 and Algorithm 2.

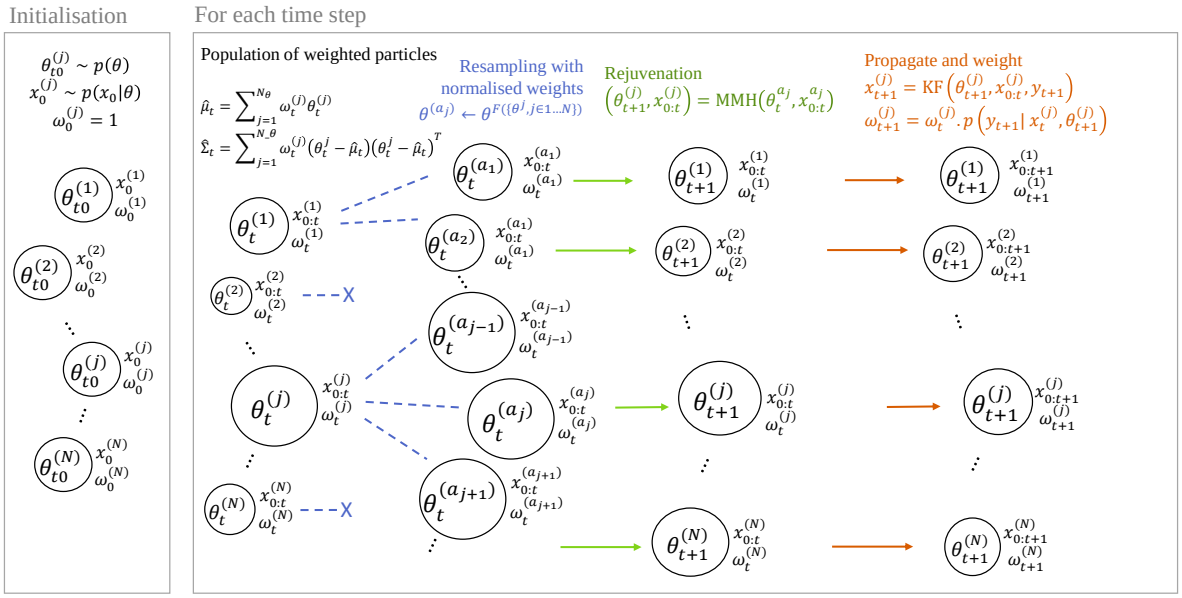

Figure 5: Principle of the SMC algorithm

The algorithm starts with the generation of a population of N_θ particles drawn from a prior distribution $p(\theta)$. Each parameter is assigned an initial state $\mathbf{x}_0$ and weight. At each time step t, a Kalman filter computes the states $\mathbf{x}_t^{(j)}$ and likelihood $L_t^{(j)}$ associated to each particle $\theta_t^{(j)}$. If the state-space model (Eq. 3) is non-linear, this filter can be replaced by a particle filter: this approach is known as the SMC^2 algorithm [32]. By this operation, the population of particles is updated so that at each time t they are a properly weighted sample from $p(\theta|y_{1:t})$ [32]. After several time steps, there is a risk that only a few of the initial N_θ particles are significantly more likely than the others and concentrate the majority of the total weight: a resampling step is then performed in order to generate a new population of particles from the most influencial ones, and a MCMC rejuvenation step then restore the diversity of particles [33].

Resampling does not occur every time a new observation becomes available, but only when required: this is measured by the effective number of particles that significantly contribute to the total weight of all

Algorithm 2 Sequential Monte Carlo algorithm

Initialisation: generate a population of N_θ particles, their states and weights
for all $j \in \{1...N_\theta\}$ **do**
  $\theta_0^{(j)} \sim p(\theta)$
  $\mathbf{x}_0^{(j)} \sim p(X_0)$
  $\omega_0^{(j)} = 1$
end for
for $t = 1...T$ **do**
  for all $j \in \{1...N_\theta\}$ **do**
    Resampling
    $\{a_j, j \in 1...N_\theta\} \leftarrow \textsc{Multinomial}\left(\omega_{t-1}^{(j)}, j \in 1...N_\theta\right)$
    Rejuvenation by a single MMH step with proposal distribution $\mathcal{N}\left(\hat{\mu}_{t-1}, \hat{\Sigma}_{t-1}\right)$
    $\left(\theta_t^{(j)}, \mathbf{x}_{0:t-1}^{(j)}, L_{t-1}^{(j)}\right) \leftarrow \text{MMH}\left(\theta_{t-1}^{(a_j)}, \mathbf{x}_{0:t-1}^{(a_j)}, \mathbf{y}_{0:t-1}\right)$
    Propagate and weight
    $\left(\mathbf{x}_t^{(j)}, L_t^{(j)}\right) \leftarrow \textsc{KalmanFilter}\left(\mathbf{x}_{t-1}^{(j)}, \theta_t^{(j)}, \mathbf{y}_t\right)$
    where $L_t^{(j)} = p\left(\mathbf{y}_t | \mathbf{x}_{t-1}^{(j)}, \theta_t^{(j)}\right)$ is the incremential likelihood.
    $\omega_t^{(j)} = \omega_{t-1}^{(j)} \cdot L_t^{(j)}$
  end for
  Normalise weights
  $\omega_t^{(j)n} = \omega_t^{(j)} / \sum_{j=1}^{N_\theta} \omega_{t-1}^{(j)}$
  Calculate weighted mean and covariance of parameters
  $\hat{\mu}_t = \sum_{j=1}^{N_\theta} \omega_t^{(j)n} \theta_t^{(j)}$
  $\hat{\Sigma}_t = \sum_{j=1}^{N_\theta} \omega_t^{(j)n} \left(\theta_t^{(j)} - \hat{\mu_t}\right) \left(\theta_t^{(j)} - \hat{\mu_t}\right)^T$
end for

particles [33]. This operation decreases the number of unique particles, hence the subsequent rejuvenation step that restores diversity. The choice of $\mathcal{N}\left(\hat{\mu}_t, \hat{\Sigma}_t\right)$ as the proposal distribution for the MCMC rejuvenation step was proposed by [31] and ensures a reasonable acceptance ratio while leaving $p\left(\theta|y_{1:t}\right)$ invariant. The rejuvenation step makes the algorithm quite computationally expensive, since the total likelihood of all particles $p\left(y_{1:t}|\theta\right)$ must be recalculated every time resampling occurs. This problem is mitigated by the fact that particles can be resampled independently, making this effort parallelisable.

4.3 Algorithm settings and performance

Both the MMH and the SMC algorithms require a large number of model evaluations. Choosing the appropriate settings (number of iterations, convergence criteria) for each is crucial to ensure convergence in a reasonable time, before comparing their evaluation results.

- Each MMH run is a chain of 20,000 iterations, including a 10,000 burn-in and a thinning factor of 10. The posterior is thus approximated by a chain of 1,000 uncorrelated samples. Convergence was ensured by checking the autocorrelation function and the stationarity of this chain.
- The SMC algorithm used 3,000 θ-particles. Other authors [20] have used 2,000 in a similar case. There is no procedure to ensure the convergence of SMC at each time step, since each particle is only propagated once.

The present work has not investigated ways to reduce the computational cost. It is difficult to compare the performance of both algorithms in terms of number of function evaluations until convergence for two reasons. First, the performance of SMC mostly relies on the number of required resampling-rejuvenation steps, which is not predictable. Second, the outcome of both algorithms is different, as SMC returns a posterior distribution at each time step: requiring the same amount of information from a sequence MMH runs would be prohibitive.

5 Results

The RRTB was monitored for 12 days, 4 of which were used to train a 3R2C model. The model was trained separately by off-line and on-line Bayesian inference, using the MMH and SMC algorithms. In order to compare both methods at different points in time, the MMH algorithm was run several times by using 1 day, 2 days, 3 days or 4 days of training data, respectively.

Both methods used the same parameter prior $p(\theta)$, which will be displayed along with the results. It is a Gaussian prior with a wide support for each of the individual parameters. Indeed, we found that using a uniform prior could compromise the convergence ability of each algorithm in the case of parameters with low identifiability.

Results are presented in the following steps:

- First, the on-line and off-line estimation results of the heat loss coefficient and solar aperture of the box are shown. Theoretical values of these characteristics [25] are available for comparison. We discuss the events that bring information to the estimates.
- Then, we show the estimation of each individual parameter of the 3R2C model: their practical identifiability is briefly discussed.
- Finally, we compare the model predictions over the test measurement period, in order to provide some validation for these results.

5.1 Estimation of HLC and solar aperture

We first consider two parameters that allow us to compare estimation results with reference values of the test box: the heat loss coefficient and the total solar aperture:

$$\mathrm{HLC} = \frac{R_1 + R_2 + R_3}{(R_1 + R_2)\ R_3} \tag{21}$$

$$k = k_1 + k_2 \tag{22}$$

One of our targets is to determine which specific part of the data drives the parameter estimation towards more confident values. Therefore, Fig. 6(b) and 6(c) show the estimation results of HLC and the solar aperture k by comparing them with measurement data in Fig. 6(a). The blue line and blue area show the average and 95% confidence interval of the posterior distributions obtained by SMC at each time coordinate. The box-and-whisker plots show the prior distribution at $t = 0$ in grey, and the four posterior distributions obtained by MMH using either 1, 2, 3 and 4 days of measurements, in red.

The 95% confidence interval of a parameter estimated by either SMC or MMH narrows down progressively, as data is sequentially added to the problem. A quick, stepwise decrease is an indicator of an event that "brings information" to the parameter. Both the HLC and the k properties have a similar behaviour in this matter: their confidence intervals are first narrowed down during the first day of measurements, as the solar irradiance rises. Then, a high information gain occurs as indoor heating is turned on, on the third day. It is general knowledge that the parameters of a building energy model are hardly identifiable without a heat source. This study however quantifies the effect of this input signal on the parameter uncertainty.

A fair match can be seen between results from both MMH and SMC algorithms: the distributions mostly overlap.

The SMC algorithm has two advantages compared to MMH in this situation: first, it was only run once to produce all sequential posterior distributions, whereas each off-line parameter estimation had to be started from the beginning of the sequence. The second advantage is the higher resolution in the temporal evolution of parameter estimates.

5.2 Individual parameters

The overall heat loss coefficient (see Fig. 6(b)) is an aggregated property, calculated from the value of three separate thermal resistances R_1, R_2 and R_3. Similarly, the total solar aperture k shown in Fig. 6(c) is the sum of two coefficients. Finding a consistent value for either does not guarantee that each parameter of the 3R2C model will be individually identifiable. For this reason, the results for all parameters of the 3R2C model are shown in Fig. 7, using the same symbols as in Fig. 6.

Here are some points that arise from these results:

- Most parameter estimates see their confidence intervals narrow down abruptly when heating is turned on ($t = 48$h).
- R_3 (Fig 7(c)) is an exception to this result. This resistance was added to the model in an attempt to represent the heat flow through the window, separately from the opaque envelope. There is close to no variation between the prior distribution and posterior distributions for this parameter: it can be considered as non identifiable in this problem.
- Conversely, Fig. 7(d) shows the estimation of the initial condition on the envelope temperature, which is an unobserved state of the 3R2C model. This parameter is estimated with a fairly high confidence as soon as the experiment starts.

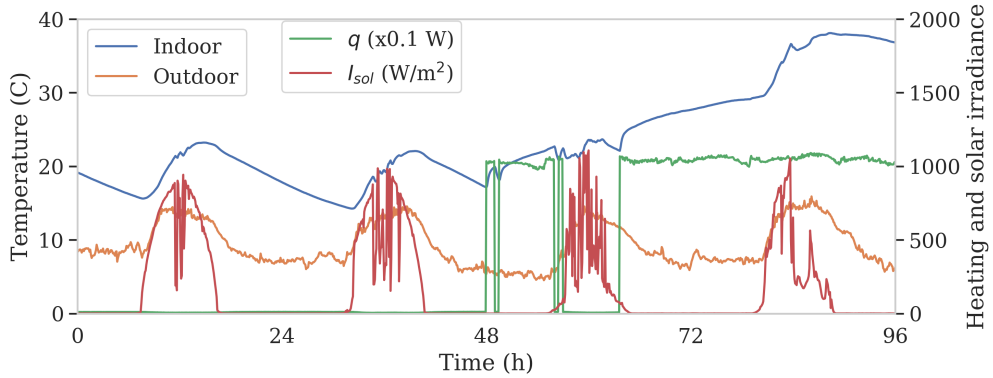

(a) Measurements in the RRTB

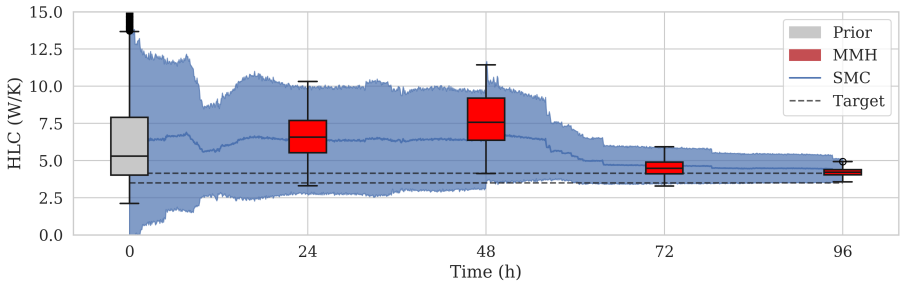

(b) Estimation of the HLC

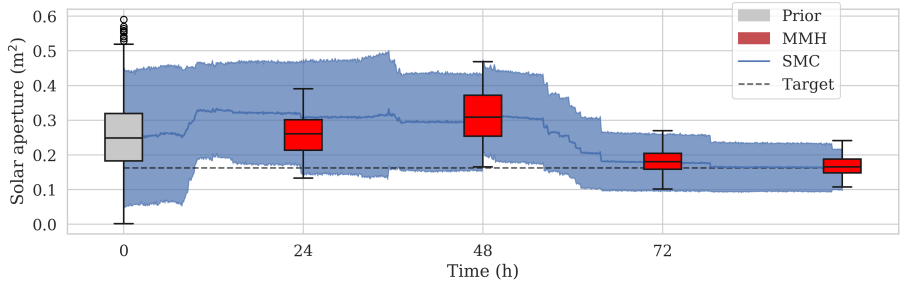

(c) Estimation of the solar aperture

Figure 6: (a) Measured indoor and outdoor temperature, heating power and solar irradiance on the RRTB; (b) Estimation of HLC by the SMC and MMH algorithms compared to the reference value; (c) Estimation of the total solar aperture

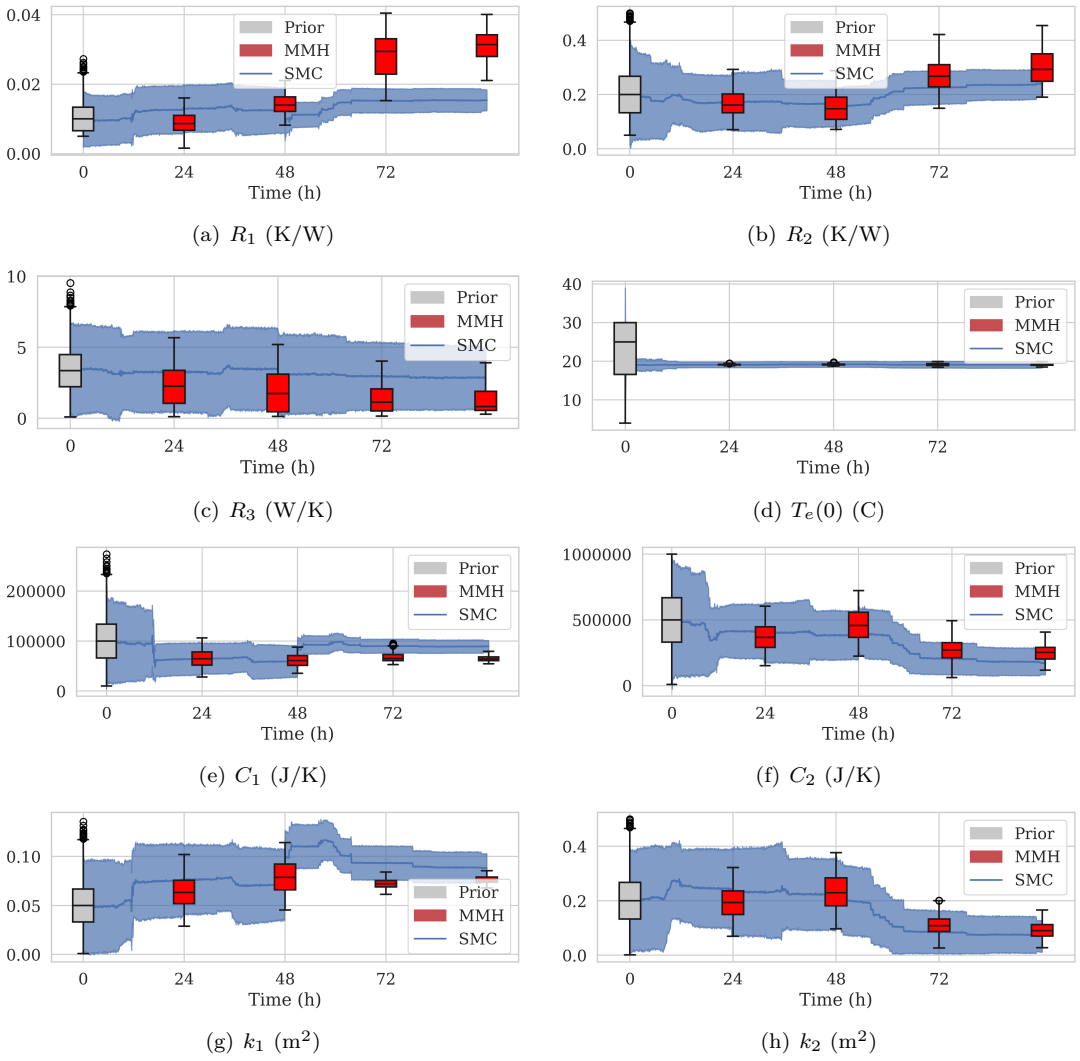

(a) R_1 (K/W)

(b) R_2 (K/W)

(c) R_3 (W/K)

(d) $T_e(0)$ (C)

(e) C_1 (J/K)

(f) C_2 (J/K)

(g) k_1 (m^2)

(h) k_2 (m^2)

Figure 7: Estimation of each individual parameter

5.3 Validation

In order to ensure the interpretability of the estimated parameter values, calibrated models should be validated. As shown in Fig. 2, the original dataset has 4 days of training data and 6 days of test data. The validation is shown in Fig. 8 by two plots.

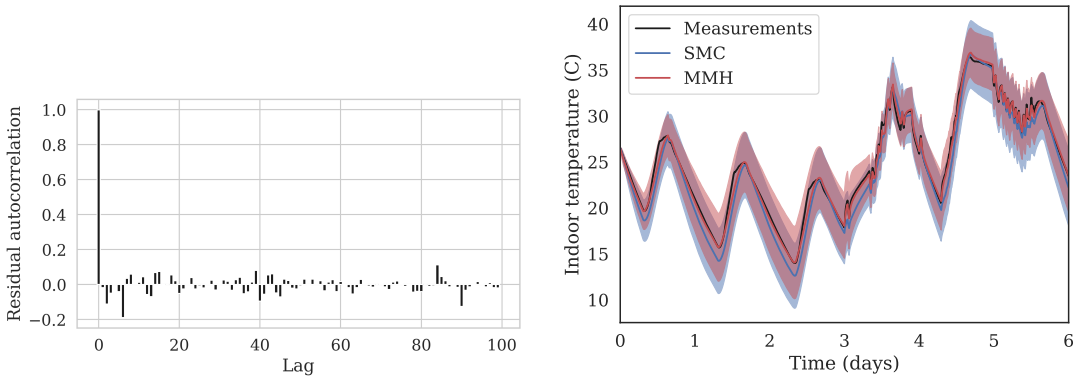

(a) (b)

Figure 8: (a) Autocorrelation function of one-step prediction residuals in the training dataset; (b) 6 days indoor temperature prediction over the test dataset with 95% confidence intervals

The auto-correlation function (ACF) of the one-step prediction residuals is shown in Fig. 8(a). Residuals were calculated with the 3R2C stochastic model calibrated by the SMC algorithm, using the mean values of parameter posterior distributions obtained at the end of the training set. The ACF indicates that the residuals are close to white noise: this implies that the 3R2C model order is sufficient to describe the dynamics of the system. The same observation can be made about the results of the MMH algorithm.

Fig. 8(b) shows the forecasts by models calibrated by both MMH and SMC, over the 6 days validation period. Forecasting uncertain states over an uncertain parameter space can be done by drawing a sample of θ from the posterior distribution, and averaging the state expectancies and variances over this sample. The initial temperature of the unobserved state at the beginning of the validation set is unknown. The mean for its initial distribution was set so that it has the same interpolation ratio between indoor and outdoor temperature, than the estimated parameter value $T_e(0)$. This has however little influence on the remaining validation period. More detail on this forecasting methodology is given in [19]. Fig. 8(b) shows a good overlap of the predictions by models calibrated with MMH and SMC. Measured indoor temperatures of the validation dataset are mostly comprised within the confidence intervals of the forecasts.

6 Conclusion

Sequential Monte Carlo is a method for on-line parameter estimation: initial probability distributions for each parameter of a model are updated sequentially, every time a new observation is received. The outcome of this method is a time series of posterior distributions, that describe the progression of acquired knowledge on the parameter values at each time step. In the present work, SMC is applied to the on-line characterisation of a simple test box (RRTB) and its results are compared to another identification technique, and validated.

- The temporal progression of posterior distributions shows which factors influence parameter estimation. Unsurprisingly, the most influencial event in the training dataset is the activation of indoor heating. This concerns all parameters of the calibrated model.
- Estimated values of the heat loss coefficient and solar aperture were in accordance with target values.

- SMC results are in accordance with parameters estimated with an off-line method (Metropolis Hastings).

7 Acknowledgements

The authors would like to thank the French National Research Agency (ANR) for funding this work through the BAYREB research project (ANR-15-CE22-0003).

References

[1] S. Rouchier, Solving inverse problems in building physics: an overview of guidelines for a careful and optimal use of data, Energy and Buildings 166 (2018) 178–195.

[2] M. F. Fels, PRISM: An introduction, Energy and Buildings 9 (1-2) (1986) 5–18.

[3] A. Rabl, A. Rialhe, Energy signature models for commercial buildings: test with measured data and interpretation, Energy and Buildings 19 (2) (1992) 143–154.

[4] G. Bauwens, S. Roels, Co-heating test: A state-of-the-art, Energy and Buildings 82 (2014) 163–172.

[5] O. Gutschker, Parameter identification with the software package LORD, Building and Environment 43 (2) (2008) 163–169.

[6] Y. Heo, R. Choudhary, G. A. Augenbroe, Calibration of building energy models for retrofit analysis under uncertainty, Energy and Buildings 47 (2012) 550–560.

[7] L. Castillo, R. Enríquez, M. J. Jiménez, M. R. Heras, Dynamic integrated method based on regression and averages, applied to estimate the thermal parameters of a room in an occupied office building in Madrid, Energy and Buildings 81 (2014) 337–362.

[8] R. Enríquez, M. J. Jiménez, M. R. Heras, Towards non-intrusive thermal load Monitoring of buildings: BES calibration, Applied Energy 191 (2017) 44–54.

[9] J. E. Braun, N. Chaturvedi, An Inverse Gray-Box Model for Transient Building Load Prediction, HVAC&R Research 8 (1) (2002) 73–99.

[10] B. Dong, C. Cao, S. E. Lee, Applying support vector machines to predict building energy consumption in tropical region, Energy and Buildings 37 (5) (2005) 545–553.

[11] X. Li, J. Wen, Building energy consumption on-line forecasting using physics based system identification, Energy and Buildings 82 (2014) 1–12.

[12] J. A. Clarke, J. Cockroft, S. Conner, J. W. Hand, N. J. Kelly, R. Moore, T. O'Brien, P. Strachan, Simulation-assisted control in building energy management systems, Energy and Buildings 34 (9) (2002) 933–940.

[13] I. Hazyuk, C. Ghiaus, D. Penhouet, Optimal temperature control of intermittently heated buildings using Model Predictive Control: Part I - Building modeling, Building and Environment 51 (2012) 379–387.

[14] Y. Lin, T. Middelkoop, P. Barooah, Issues in identification of control-oriented thermal models of zones in multi-zone buildings, in: 2012 IEEE 51st Annual Conference on Decision and Control (CDC), 2012, pp. 6932–6937. doi:10.1109/CDC.2012.6425958.

[15] H. Madsen, J. Holst, Estimation of continuous-time models for the heat dynamics of a building, Energy and Buildings 22 (1) (1995) 67–79.

[16] N. R. Kristensen, H. Madsen, S. B. Jørgensen, Parameter estimation in stochastic grey-box models, Automatica 40 (2) (2004) 225–237.

[17] M. J. Jiménez, B. Porcar, M. R. Heras, Estimation of building component UA and gA from outdoor tests in warm and moderate weather conditions, Solar Energy 82 (7) (2008) 573–587.

[18] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy and Buildings 43 (7) (2011) 1511–1522.

[19] S. Rouchier, M. Rabouille, P. Oberlé, Calibration of simplified building energy models for parameter estimation and forecasting: Stochastic versus deterministic modelling, Building and Environment 134 (2018) 181–190.

[20] L. Raillon, C. Ghiaus, Sequential Monte Carlo for states and parameters estimation in dynamic thermal models, San Francisco, USA, 2017.

[21] A. Doucet, S. Godsill, C. Andrieu, On sequential Monte Carlo sampling methods for Bayesian filtering, Statistics and Computing 10 (3) (2000) 197–208.

[22] J. E. Handschin, Monte Carlo techniques for prediction and filtering of non-linear stochastic processes, Automatica 6 (4) (1970) 555–563.

[23] J. S. Liu, R. Chen, Sequential Monte Carlo Methods for Dynamic Systems, Journal of the American Statistical Association 93 (443) (1998) 1032–1044.

[24] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, N. Chopin, On Particle Methods for Parameter Estimation in State-Space Models, Statistical Science 30 (3) (2015) 328–351.

[25] M. J. Jiménez, Reliable building energy performance characterisation based on full scale dynamic measurements. Report of subtask 3, part 1: Thermal performance characterization based on full scale testing - description of the common exercises and physical guidelines, in: International Energy Agency, EBC Annex 58., 2016.

[26] R. Shumway, D. Stoffer, Time series analysis and its applications, Springer, 2016.

[27] E. A. Wan, R. V. D. Merwe, The unscented Kalman filter for nonlinear estimation, in: Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No.00EX373), 2000, pp. 153–158.

[28] C. Andrieu, A. Doucet, R. Holenstein, Particle Markov chain Monte Carlo methods, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72 (3) (2010) 269–342.

[29] L. Raillon, C. Ghiaus, An efficient Bayesian experimental calibration of dynamic thermal models, Energy 152 (2018) 818–833.

[30] O. Cappé, S. Godsill, E. Moulines, An overview of existing methods and recent advances in sequential monte carlo, Proceedings of the IEEE 95 (5) (2007) 899–924.

[31] N. Chopin, A sequential particle filter method for static models, Biometrika 89 (3) (2002) 539–552.

[32] N. Chopin, P. E. Jacob, O. Papaspiliopoulos, SMC2: an efficient algorithm for sequential analysis of state space models, Journal of the Royal Statistical Society: Series B (Statistical Methodology) 75 (3) (2013) 397–426.

[33] L. M. Murray, Bayesian State-Space Modelling on High-Performance Hardware Using LibBi, arXiv:1306.3277 [stat]ArXiv: 1306.3277.