Building identification within a connected object ecosystem

Tahar Nabil

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY

TELECOM ParisTech

2018-ENST-0001

EDITE - ED 130

Doctorat ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

TELECOM ParisTech

Spécialité « Signal et Images »

présentée et soutenue publiquement le 9 janvier 2018 par

Tahar NABIL

Identification de modèle thermique de bâtiment dans un environnement d'objets connectés

Directeurs de thèse : **François ROUEFF**
Éric MOULINES
Encadrants : **Alexandre GIRARD**
Jean-Marc JICQUEL

Jury

M. Stéphane DERRODE, Professeur, LIRIS, UMR 5205, Ecole Centrale de Lyon	Rapporteur
M. Stéphane PLOIX, Professeur, GSCOP, Grenoble INP	Rapporteur
M. Karim ABED-MERAIM, Professeur, PRISME, Université d'Orléans	Examinateur
Mme Émilie CHAUTRU, Enseignant-chercheur, Centre de Géosciences, Mines ParisTech	Examinatrice
M. Éric MOULINES, Professeur, CMAP, UMR 7641, Ecole Polytechnique	Directeur de thèse
M. François ROUEFF, Professeur, LTCI, Télécom ParisTech	Directeur de thèse
M. Alexandre GIRARD, Ingénieur-chercheur, Centre de Chatou, EDF R&D	Invité
M. Jean-Marc JICQUEL, Ingénieur-chercheur, Centre des Renardières, EDF R&D	Invité

T H È S E

TELECOM ParisTech
école de l'Institut Mines-Télécom - membre de ParisTech
46 rue Barrault 75013 Paris - (+33) 1 45 81 77 77 - www.telecom-paristech.fr

Remerciements

Le travail de thèse représente un effort à la fois individuel et collectif, et c'est un grand plaisir que de pouvoir profiter de ces quelques lignes pour remercier l'ensemble des personnes sans lesquelles ce travail n'aurait pu aboutir.

En premier lieu, j'ai bénéficié d'un encadrement nombreux, complémentaire et soutenu. C'est une grande richesse, et je remercie mes encadrants au sein d'EDF R&D, Jean-Marc Jicquel et Alexandre Girard. Jean-Marc, merci pour la confiance dont tu as fais preuve à mon égard, et que tu as transmise tout au long de la thèse. Merci également d'avoir su guider la thèse, en contre-balançant ma tendance à l'éparpillement... Alexandre, merci d'avoir enrichi nos comités mensuels par ta grande culture scientifique, impressionnante, et par tes suggestions toujours pertinentes. Du côté de l'encadrement académique à Télécom ParisTech, je remercie Éric Moulines, qui a permis de lancer la thèse sur des bases solides. Merci également d'avoir transmis le relai après quelques mois à François Roueff, suite à tes nouvelles activités, tout en maintenant ton intérêt pour le sujet. J'ai beaucoup appris grâce à toi, François, sur les séries temporelles et les mathématiques statistiques de manière générale, et je suis particulièrement ravi et reconnaissant que tu aies joué le jeu de devenir directeur de thèse en cours de route. Enfin, un chaleureux merci à Maurice Charbit de s'être intéressé au sujet de la thèse. Ce fut un réel plaisir de discuter avec toi, ta bonne humeur redonne de l'énergie et de la motivation même dans les périodes de creux ! Au final, la thèse s'est réalisée en harmonie entre EDF R&D et Télécom ParisTech, ce qui est si important mais peut-être pas si commun dans le cas d'une cotutelle, alors encore une fois, un grand merci à vous tous.

Je suis également honoré et reconnaissant que Karim Abed-Meraim ait présidé le jury de la soutenance, qu'Émilie Chautru ait accepté de faire partie de ce jury, et que Stéphane Derrode et Stéphane Ploix aient rapporté ce travail. Vous représentez les deux sensibilités de cette thèse, traitement statistique du signal et physique du bâtiment, et vos remarques à la fois exigeantes et encourageantes valorisent d'autant plus le travail fourni.

Cette thèse doit beaucoup aux conditions dans lesquelles elle a été réalisée, c'est-à-dire à la grande liberté accordée par EDF R&D à ses doctorants. Je remercie en particulier mon cher chef de projet, René S., qui a non seulement permis que la thèse voit le jour mais m'a en plus donné accès à des moyens d'essais qui enrichissent notablement ce travail. Ce qui m'amène naturellement à remercier sincèrement tous mes collègues du groupe E13 devenu E36 au sein du département TREE d'EDF R&D. Ce fut toujours un plaisir de faire le périple jusqu'aux Renardières, et c'est grâce à votre bonne humeur et l'agréable ambiance de travail qui y règne. Une pensée particulière pour Florence C., pour ta gentillesse et ta capacité si précieuse à résoudre tous les problèmes administratifs. Pêle-mêle, mais non moins sincèrement, je tiens également à remercier les collègues de BESTLab, Hervé, Mathilde et Matthieu, ainsi que Sandra pour la rédaction des demandes de brevet, et les doctorants passés et présents du K15, Alberto, Kévin, Maria, Nathanael et Paul qui ont été autant

de sources d'inspiration et de motivation. Et enfin, un grand merci aux membres du club ping-pong des Renardières, c'est-à-dire l'inarrêtable Bernard, ainsi qu'Anthony, Antoine, Jérôme, Kévin, Laurent, Nicolas et Xavier. Grâce à votre gentillesse, votre sportivité et vos blagues - qui, il faut bien l'admettre, ne sont pas toujours très fines... -, vous m'avez offert une respiration hebdomadaire absolument bienvenue. Et surtout, vous m'avez notamment appris que j'avais le droit de taper des coups droits, que demander de plus !

Malgré une présence plus aléatoire, je suis content d'avoir partagé quelques moments au sein du département IDS du LTCI à Télécom ParisTech, en particulier avec Adrien, Andrés, Claire, Eugène, Guillaume, Moussab, Romain. Votre haut niveau académique a toujours été un exemple à suivre pour moi !

Si ces trois années de thèse se sont si bien déroulées, c'est aussi parce que je dois beaucoup à mes amis, sur qui je peux toujours compter : Erwan, Hicham, Houssine, Marya et Romain, Farid, Omar, Rémi, Léa et Pierre, Noura et Antoine, Alexia et Raphaël, et les plus vieux, Nicolas et Vincent. Merci à vous tous, je suis tellement heureux de vous avoir ! *Special thoughts also for Anda, Erik, Rauan and Vlad who have inspired me into going for the PhD adventure.*

L'un de mes moteurs a toujours été de bien faire pour ma famille. Je pense en particulier à mes grands-parents, en Auvergne et au Maroc, à mes parents, qui m'ont tout donné et sans qui je n'aurais jamais eu ce parcours, et enfin mes frères, Thami et Hedi, à qui j'espère piquer un peu de leur créativité ! J'espère pouvoir vous rendre tout ce que vous m'apportez.

Pour finir, je remercie Sarah, qui m'a supporté dans cette thèse avec brio, (quasiment) sans stress et toujours avec amour, gaieté et patience. Tout est plus beau avec toi et grâce à toi !

Résumé

Cette thèse s'intéresse au développement de méthodes d'identification de modèles de comportement thermique dynamique d'un bâtiment, en saison de chauffe. Le problème est classiquement décrit de la façon suivante. Soit un bâtiment représenté par un modèle expliquant l'évolution temporelle de la température intérieure en fonction de sollicitations externes (température extérieure, gains solaires) et internes (puissance de chauffage, apports internes gratuits). Les paramètres du modèle étant inconnus, l'identification de bâtiment cherche à les estimer à partir de données d'observation de chaque entrée et sortie du système, mesurées à un pas de temps de l'ordre de quelques minutes. Le problème d'identification est donc un problème d'estimation statistique, qui dépend fortement de la disponibilité réelle des données.

Dans un premier temps, on se donne un modèle de bâtiment sous forme de circuit électrique équivalent, et on suppose que toutes les grandeurs sont correctement observées. Plutôt que de modéliser la régulation de la température intérieure, nous proposons d'estimer directement les paramètres du système en boucle ouverte, bien que les données d'exploitation soient en fait générées en boucle fermée. Cette approche est étayée par des arguments statistiques et empiriques.

Cependant, en pratique, l'accès aux données n'est pas garanti pour l'énergéticien, puisqu'il ne mesure *a priori* que la consommation totale du bâtiment. L'idée de la thèse est de se placer dans le cadre d'un "bâtiment intelligent", dont les objets connectés permettent d'améliorer l'observation de son comportement thermique, sans toutefois supprimer toute incertitude. Nous remettons donc en cause successivement la disponibilité effective de chaque entrée ou sortie du modèle, et proposons des méthodes d'identification adaptées, validées sur des données simulées et réelles, issues d'expériences à EDF R&D. En particulier, pour la température intérieure, une démarche non-intrusive, se passant de la mesure, est présentée. Pour la température extérieure, des capteurs connectés sont aujourd'hui largement répandus dans le commerce : la méthode proposée permet de corriger le biais induit par l'emplacement inconnu dudit capteur pour l'énergéticien. Puis, le flux solaire n'étant pas mesurable en raison du coût prohibitif du capteur associé, nous suggérons d'intégrer une seconde sonde connectée de température, volontairement exposée au rayonnement solaire. L'algorithme associé exploite l'idée que la différence de mesure entre ce capteur exposé au soleil et un capteur de référence, protégé du rayonnement, est liée à l'intensité du rayonnement. Enfin, en utilisant les informations On/Off des objets connectés du bâtiment, le dernier algorithme permet l'identification du bâtiment à partir de la courbe de charge totale du bâtiment, sans distinction entre les charges qui représentent ou non des apports thermiques.

Abstract

This thesis aims at developing methods for identifying dynamic thermal building models. The problem can be described as follows. The thermal behaviour of a building is given by a model, which explains the dynamics of the indoor air temperature in terms of weather conditions (outdoor air temperature, solar gains) and internal excitations (controlled heating power, free heating gains). Since the parameters of the chosen model are unknown, the purpose of building identification is to estimate them given some measurements of every input and output of the system. The identification is therefore an algorithmic task that strongly depends on the actual availability of the data.

In a first part, we assume that the building model takes the form of an equivalent electrical network and that the inputs and output are accurately measured. Although the system operates in closed loop, with a controlled indoor temperature, we suggest to estimate directly the parameters of the open-loop model. This approach allows us not to model the nonlinear regulation, whereas statistical and empirical arguments are provided in order to justify it.

However, an energy utility has a limited access to the data in practice, the only guaranteed measurement being the total energy consumption of the building. The main idea developed in the thesis is to consider a "smart building", whose connected objects make possible an enhanced observation of the building's thermal behaviour. In order to address the uncertainty inherent to this framework, we question successively the availability of every input and output of the model, before suggesting adapted algorithms validated on artificial and real data stemming from experiments at EDF R&D. We begin with a non-intrusive closed-loop identification procedure that does not require any measurement of the indoor temperature. Next, we use the already wide-spread connected weather stations for measuring the outdoor temperature. Yet, the lack of accuracy of such sensors may provide misleading measurements. Hence, we propose an estimation algorithm able to correct the possible temperature bias and account for the latent uncertainty. Then, we suppose that the solar radiation is not directly measured, because the corresponding sensor comes with a prohibitive cost. We recommend instead to replace it with another connected temperature sensor, intentionally left exposed to the sun. The estimation algorithm exploits the idea that larger values of the difference between the measurements provided by this sensor on one hand and a reference, sheltered, temperature sensor on the other hand, are caused by larger values of the solar flux. Finally, by taking advantage of On/Off information given by the connected objects of the building, the last algorithm enables the identification from the total consumption of the building, without having to disaggregate between heating and non-heating loads.

Contents

List of Figures

Nomenclature

Acronyms

ARX	Auto-Regressive model with eXogenous input
BESTLab	Building Envelope & Solar Technology laboratory
CLGSS	Conditionally Linear Gaussian State Space model
CLT	Central Limit Theorem
EM	Expectation Maximization
EWA	Exponentially Weighted Aggregation
HMC	Hybrid, or Hamiltonian, Monte Carlo
MCMC	Markov Chain Monte Carlo
MLE	Maximum Likelihood Estimator
PF	Particle Filter
PI	Proportional - Integral
PS	Particle Smoother
RBPF	Rao-Blackwellised Particle Filter
RMSE	Root Mean Square Error
SMC	Sequential Monte Carlo

Capital letters

Notation	Meaning	Unit
T	Number of discrete observations	-
A	State matrix of a given state-space model	-
B	Input matrix of a given state-space model	-
C	Output matrix of a given state-space model	-
C_r	Air mass capacity	J/°C
C_s	Envelope capacity	J/°C
G_0	Extraterrestrial solar constant	W/m^2
$G_{h,clr}$	Global horizontal solar flux under clear-sky conditions	W/m^2
I_0	Instants during which no non-heating load is On	-
I_1	Instants during which at least one non-heating load is On	-
M_1	Number of particles in the forward particle filter	-

M_2	Number of trajectories drawn in the backward particle smoother	-
N	Cloud cover index	-
Q	Measured total load curve	W
Q_h	Power delivered by the heating system	W
Q_h^{max}	Maximum power delivered by the heating system	W
Q_r	Heating flux, including the heating system, free internal and metabolic gains	W
Q_s	Solar gains	W
Q_u	Useful load curve, the sum of Q_h and Q_{fg}	W
Q_{fg}	Free heating gains due to the appliances (lighting, etc.)	W
Q_m	Metabolic heating gains	W
Q_{nh}	Portion of the total load curve corresponding to non-heating uses	W
R_f	Thermal resistance between the indoor and outdoor air (lighter components)	°C/W
R_i	Thermal resistance between the indoor air and the envelope (heavy materials)	°C/W
R_o	Thermal resistance between the envelope and the outdoor air (heavy materials)	°C/W
T_i	Noisy measurement of the indoor temperature	°C
T_o	Noisy measurement of the outdoor temperature	°C
T_r	Setpoint temperature	°C
U	Observed inputs of a given state-space model	-
UA	Heat transfer coefficient	W/°C
V	Gaussian white noise (state equation)	-
W	Gaussian white noise (observation equation)	-
X	Hidden states of a given state-space model	-
Y	Observed outputs of a given state-space model	-

Greek letters

Notation	Meaning	Unit
α	Linear state in a conditionally linear Gaussian state-space model	-
β	Parameters for the dynamics of the outdoor temperature bias	-
χ_{nh}	Indicator function for the non-heating events	-
ΔT_o	Noisy observation of the outdoor temperature bias	°C
κ	Constant in Perrin de Brichambault's model	-
ρ	Learning rate in the aggregation strategy	-
σ_t	High variance increment for the dynamics of the outdoor temperature bias	-
τ	Time constant	s
τ_1	Slow time constant	h
τ_2	Fast time constant	h
θ	Parameters of the building model	-
θ_z	Azimuth angle	°

ϑ_b	Model outdoor temperature bias	°C
ϑ_i	Model temperature of the indoor air	°C
ϑ_o	Model temperature of the outdoor air	°C
ϑ_s	Model temperature of the envelope of the building	°C
ξ	Nonlinear state in a conditionally linear Gaussian state-space model	-

Lowercase letters

Notation	Meaning	Unit
a	Parameter in the cloud cover correction model	-
b	Parameter in the cloud cover correction model	-
g	Solar transmittance	-
z_f	Inverse of the thermal resistance R_f	°C/W
z_i	Inverse of the thermal resistance R_i	°C/W
z_K	Gain of the PI controller	W/°C
z_o	Inverse of the thermal resistance R_o	°C/W
z_r	Inverse of the air mass capacity	°C/J
z_τ	Inverse of the time constant of the PI controller	s^{-1}
z_s	Inverse of the enveloppe capacity	°C/J

Operators

$\bar{\cdot}$	Empirical average
$\delta_{\cdot}$	Kronecker function
$\mathbb{E}$	Expectation
$\mathcal{N}(m, \Sigma)$	Gaussian distribution with mean m and covariance Σ
$\mathcal{N}(x; m, \Sigma)$	Gaussian density with mean m and covariance Σ, evaluated in x
$\mathcal{Q}$	Conditional expectation of the complete log-likelihood
∇	Gradient
$\widehat{\cdot}$	Statistical estimator
$f(\,\cdot\,)$	State function in a general state space representation
$h(\,\cdot\,)$	Observation function in a general state space representation
$p(\,\cdot\,)$	Generic probability density

Introduction en français

Identification de bâtiment et données manquantes

Contexte de l'identification de bâtiment

Le secteur du bâtiment représente actuellement un enjeu de taille, au cœur aussi bien des politiques énergétiques mondiales que de nombreux efforts de recherche. Deux facteurs en particulier expliquent cette tendance. Le premier est le coût environnemental élevé de ce secteur : il représentait 37% de la consommation finale d'énergie au sein de l'Union Européenne (UE) en 2004 [PLOP08], 44% en France en 2012 [ADE13]. La Figure 1(a) montre ainsi que cette part a toujours été élevée en France depuis 1970, et est même à son paroxysme en 2015. Ces consommations s'accompagnent en outre d'émissions de gaz à effet de serre : 19.6% du total des émissions de l'UE sont dûs aux bâtiments résidentiels uniquement, selon Eurostats. On observe une proportion similaire en France, comme indiqué en Figure 1(b). Par ailleurs, la grande majorité de ces émissions (77% en France en 2009, [ADE13]) sont générées par le chauffage. Par conséquent, ces éléments combinés aux faibles taux de renouvellement du parc de bâtiments (typiquement 1% par an en France) mettent en lumière la nécessité de mettre en œuvre des actions d'efficacité énergétique des bâtiments, afin de promouvoir une transition vers une demande énergétique plus soutenable.

Le second facteur quant à lui est lié à l'émergence des nouveaux réseaux, dits "intelligents", de distribution de l'électricité, les *Smart Grids* [FMXY12]. Un concept clé de ces réseaux est la notion de flexibilité, avec le défi de débloquer et exploiter la flexibilité de la demande afin de garantir l'équilibre avec une offre toujours plus intermittente, en raison de l'essor des énergies renouvelables. Les bâtiments apparaissent alors comme des acteurs centraux des *Smart Grids*. En effet, alors qu'il reste difficile de stocker l'énergie électrique, la forte consommation du secteur du bâtiment représente un gisement significatif de flexibilité de la demande, capable de maintenir l'équilibre du réseau. Des actions comme l'effacement de consommation électrique ou la gestion active de la demande essaient d'exploiter au mieux ce gisement. L'effacement vise par exemple à réduire le pic journalier de consommation afin de lisser la demande. Lorsque ce pic est trop important, notamment en hiver, la production d'électricité est assurée par des centrales thermiques, qui sont parmi celles contribuant le plus aux émissions de dioxyde de carbone. Ainsi, les technologies associées aux *Smart Grids* permettraient également de réduire le coût environnemental des bâtiments. Ces solutions reposent notamment sur des méthodes avancées de gestion et de contrôle du système de chauffage, afin de tirer profit des capacités de stockage thermique des bâtiments.

En outre, le nouvel écosystème digital, transformé par l'irruption des smartphones, du *cloud* et de l'internet des objets (*Internet of Things*, IoT), représente une opportunité pour relever ces défis. Par exemple, des études de cas exposées dans [VNB+16] ou [GAO+16] suggèrent qu'un retour d'information en temps réel via le smartphone pourrait abaisser la température de la pièce, et donc la consommation en énergie, sans affecter significativement le confort des occupants. Citons également le développement des thermostats connectés, qui pourraient aussi économiser l'énergie

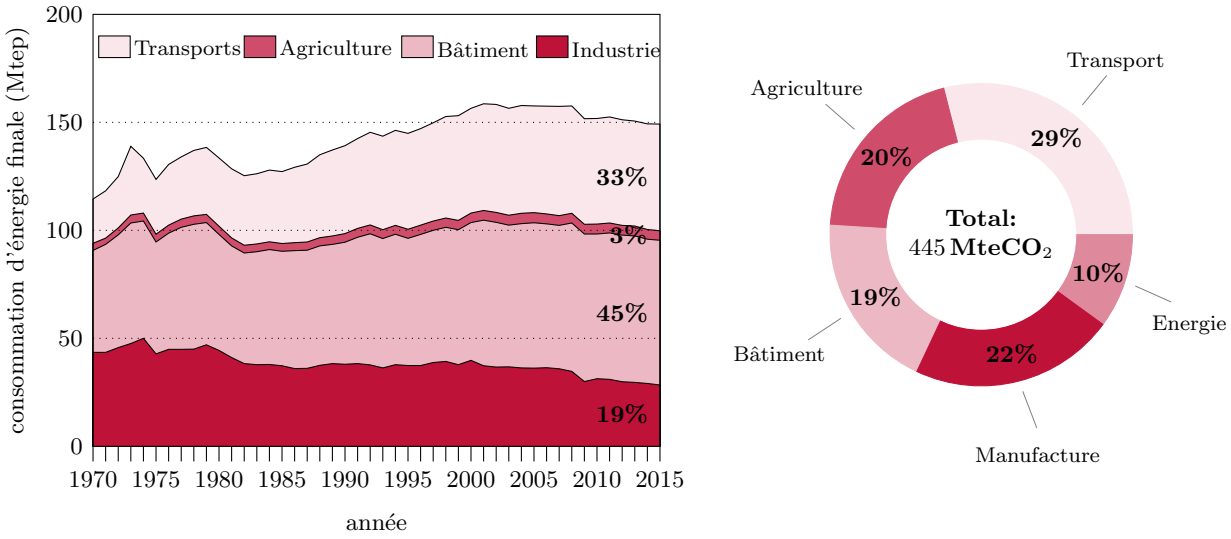

Figure 1: (gauche) Consommation finale annuelle par secteur en tonne équivalent pétrole (tep) en France, 1970 – 2015. Source: Service de l'observation et des statistiques, 2015. (droite) Emissions de gaz à effet de serre en tonne équivalent CO_2 ($teCO_2$) par secteur en France métropolitaine en 2015. Source: Citepa/Format SECTEN - april 2017.

en ajustant la température de consigne dès que la pièce contrôlée n'est plus occupée, sans non-plus altérer le confort thermique [PWM16]. Les fournisseurs d'énergie sont donc pressés de créer de nouveaux services énergétiques, voire de nouveaux objets, afin de répondre à ces enjeux tout en faisant face à la compétition, autrefois inattendue, des acteurs des technologies de l'information et de la communication tels les géants du Web.

Ce contexte marque donc un changement de paradigme, comme signalé dans [SBPW16]. Les approches antérieures en énergétique du bâtiment s'intéraissaient avant tout à la phase de conception, c'est-à-dire que les thermiques détaillés étaient simulés préalablement à l'occupation réelle du bâtiment, en s'appuyant par exemple sur des scénarios standardisés d'occupation. Le nouveau paradigme implique au contraire que l'utilisateur doit être au cœur des services, *via* des intéractions : en échange des informations qu'il/elle fournit, par exemple à travers la mesure de certaines grandeurs d'intérêt, comme la température intérieure, l'utilisateur a accès à un service l'aidant à contrôler et à réduire sa consommation énergétique.

Conformément à ce contexte, des méthodes avancées d'automatique, telle que la commande prédictive (voir e.g. [HGP14]), sont appliquées au système de chauffage, ventilation et climatisation. Leur impact sur la consommation du bâtiment est potentiellement élevé : des expériences détaillées dans [PŠFC11] et [ŠOCP11] atteignent des économies d'énergie respectives de 17-24% et 15-28%. De même, l'effacement de charge est réalisé grâce à des méthodes de pilotage dynamique de l'énergie du bâtiment par commande optimale, sous des contraintes de confort thermique [Mal12]. Ces différentes actions dépendent en premier lieu d'un bon **modèle de bâtiment** [PCV+13], capable d'évaluer leur potentiel et limites, ainsi que les performances énergétiques dudit bâtiment. Ainsi, dans [Mal12], la méthode d'effacement dépend d'une représentation fiable de la dynamique de la température intérieure, afin de garantir que la stratégie de pilotage respecte bien les contraintes de confort thermique. Le modèle de bâtiment constitue donc un *prérequis* à la création de ces services. Cette thèse s'intéresse ainsi plus particulièrement à la démarche d'estimation de tels modèles.

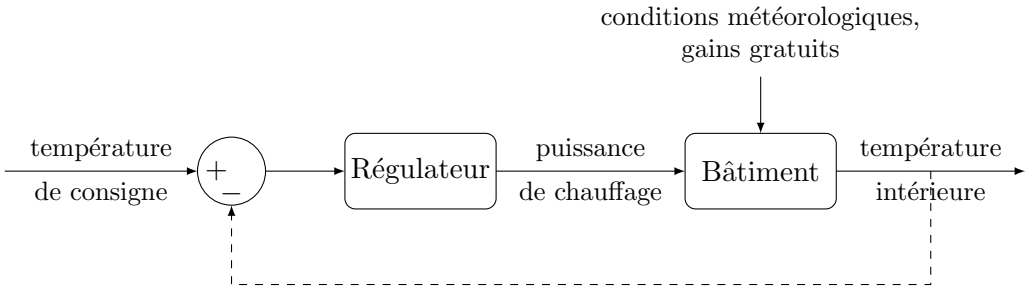

Figure 2: Schéma bloc d'un système de bâtiment avec régulation de la température intérieure. Sollicitations externes : gains internes gratuits liés aux autres appareils électriques et à l'occupation, conditions météorologiques incluant e.g. la température extérieure et le flux solaire.

Par modèle de bâtiment, on entend plus précisément un modèle représentatif de l'état thermique du bâtiment et capable d'expliquer la dynamique de la température intérieure. Cette dynamique s'explique en général par trois classes de facteurs, à savoir les facteurs liés à la structure même du bâtiment (composition, géométrie, etc.), ceux liés aux conditions météorologiques (température extérieure, flux solaire, vitesse du vent, etc.) et enfin les intéractions des occupants avec le bâtiment [Zay11, AAC+15]. Ces intéractions peuvent être contrôlables (température de consigne, système de chauffage) ou non-contrôlables et considérées alors comme des perturbations (gains métaboliques liés à la présence, gains internes gratuits causés par exemple par les ordinateurs, l'éclairage, etc. [PRMS08]). Nous étudions plus spécifiquement le bâtiment en saison de chauffe, typiquement en hiver. Le bâtiment évolue alors en boucle fermée, la température intérieure étant asservie afin de suivre un certain signal de consigne. Voir la Figure 2 pour une représentation du modèle de bâtiment d'un point de vue "système".

Plusieurs approches de modélisation existent afin d'expliquer la dynamique de la température intérieure à partir de ces facteurs. Il y a essentiellement deux paradigmes, avec d'un côté les modèles physiques détaillés reposant sur une connaissance experte, et de l'autre les modèles statistiques, construits uniquement à partir des données de fonctionnement du bâtiment. Nous optons dans la suite de la thèse pour la seconde catégorie. Ce choix est déterminé par le contexte précédemment décrit et par les usages potentiels du modèle de bâtiment. En effet, la commande prédictive nécessite des modèles avec un faible nombre de paramètres et pouvant simuler la dynamique du bâtiment pour un faible coût calculatoire. Ainsi, l'approche statistique est particulièrement adaptée, au contraire des modèles physiques d'ordre élevé. Les différences entre ces deux approches sont décrites dans le Chapitre 2 de ce document.

La construction des modèles statistiques repose sur des méthodes d'*identification* de systèmes, et donc sur des expérimentations, comme définit par L. Ljung dans son ouvrage de référence [Lju87] consacré au sujet : *"Les signaux entrée/sortie du système sont enregistrés et soumis à de l'analyse de données afin d'inférer un modèle"*. Les principales étapes de la démarche d'identification sont le choix d'une classe de modèles, la planification d'une campagne de mesure des facteurs explicatifs les plus pertinents, le choix d'un critère pour la sélection du modèle, la construction (estimation) d'un modèle à partir des données mesurées, et enfin la validation [Zay11]. Se référer e.g. à [Lju87] pour une étude approfondie de cette démarche. Dans la section suivante, nous soulignons le fait que les observations des signaux entrée/sortie sont déterminantes pour l'identification du modèle thermique d'un bâtiment.

De l'importance des données

Etant fixée une certaine classe de modèles de bâtiment, l'identification est un problème statistique, et, par suite, algorithmique, qui dépend fortement de la disponibilité des données entrée/sortie du modèle. Plutôt que de se concentrer uniquement sur l'aspect algorithmique de la démarche, auquel une grande partie de l'état de l'art est déjà consacrée, cette thèse étudie plus particulièrement les contraintes pratiques qui surviennent dans l'implémentation de tels algorithmes. Ces contraintes sont dues au second aspect de l'identification, c'est-à-dire la disponibilité réelle des données.

Cette question de la disponibilité et de la qualité des données est souvent négligée dans la littérature relative à l'identification de bâtiment. Dans beaucoup de publications, les données sont obtenues soit par une instrumentation spécifique de laboratoire (e.g. [BM11, Zay11, Mej11]), soit par un mélange de bases de données météorologiques historiques moyennées et de la réponse simulée du bâtiment à ces sollicitations moyennes (e.g. [MCPF10, HGP12]). Ces deux approches supposent ainsi que chaque entrée et sortie du modèle thermique est disponible à une précision suffisante : en ce sens, elles sont déterministes. Nous pouvons toutefois illustrer qu'elles ne sont pas adaptées à un contexte industriel. La seconde approche ne convient de toute évidence pas à une mise en œuvre pratique, puisque la simulation de la dynamique s'obtient alors nécessairement par un modèle physique détaillé, alors que ces modèles sont précisément écartés de notre démarche en raison de la lourdeur de modélisation qu'ils représentent : c'est un travail peu adapté aux besoins industriels, car spécifique à chaque bâtiment, et qui requiert l'intervention d'un expert.

Prenons ensuite l'exemple du flux solaire pour motiver notre propos et montrer que la première approche, reposant sur une instrumentation complète du bâtiment, n'est pas non plus réalisable en pratique. Une mesure de cette quantité est nécessaire notamment pour estimer correctement les constantes de temps qui caractérisent l'inertie du bâtiment. Cette entrée représente de plus une source substancielle d'apports thermiques gratuits que les récents bâtiments basse consommation tendent à maximiser. Cependant, l'instrumentation associée a un coût prohibitif, elle est donc difficilement déployable à grande échelle pour un grand nombre de bâtiments. En effet, le flux solaire est mesuré par un pyranomètre dont le coût capteur excède typiquement 100€, hors installation : il semble peu envisageable d'équiper chaque utilisateur d'une telle instrumentation. Par conséquent, la première approche qui suppose une observation parfaite du bâtiment semble de fait illusoire en pratique.

Remarquons également que les deux approches ne permettent pas de prendre en compte les apports thermiques gratuits dûs à l'occupation du bâtiment, ces derniers étant fortement intermittents par essence et difficile à mesurer avec précision.

Plutôt que de supposer que chaque entrée et sortie du modèle est un signal déterministe, cette thèse explore donc une troisième option, en adoptant le point de vue du fournisseur d'énergie. Celui-ci dispose *a priori* de la mesure d'une unique grandeur, grâce au déploiement des compteurs intelligents, à savoir la consommation totale d'énergie à un pas de temps infra-horaire (typiquement 10 minutes), appelée courbe de charge. Il y a cependant des nuances selon le type de bâtiment considéré :

- une maison ou un appartement dont le vecteur d'énergie est uniquement l'électricité : alors la mesure correspond à la consommation agrégée de l'ensemble des usages;
- une maison ou un appartement dont le vecteur d'énergie sont l'électricité et le gaz : alors il y a généralement une mesure agrégée du chauffage et de l'eau chaude sanitaire d'un côté (gaz), et une mesure agrégée des autres usages de l'autre (électricité);

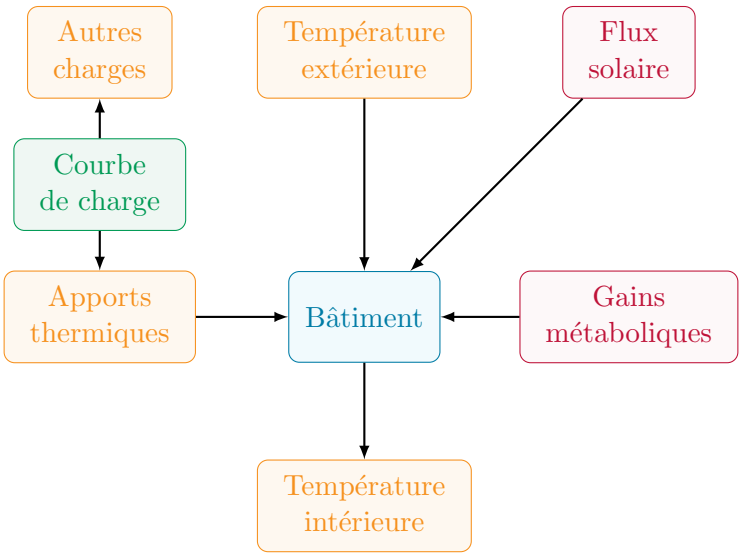

Figure 3: Disponibilité effective des entrées et sorties principales d'un modèle thermique de bâtiment, pour le fournisseur d'énergie. En rouge, les grandeurs non-mesurées. En orange, celles mesurées avec une incertitude non-négligeable. En vert, les grandeurs mesurées avec précision.

- pour le secteur tertiaire (bâtiments de bureaux, ...) ou industriel, la mesure est généralement l'agrégation de toutes les consommations.

Quoiqu'il en soit, les données mesurées ont besoin d'être séparées entre les apports thermiques et les autres charges, les premiers étant intégrés au modèle tandis que les secondes en sont écartées. Une approche possible pour y parvenir serait d'implémenter un algorithme de désagrégation de type NILM (*Non Intrusive Load Monitoring*, i.e. surveillance non-intrusive de la charge des appareils, voir [ZGIR12, EG09]). Cependant, ces méthodes reposent typiquement sur une mesure à haute-fréquence de la courbe de charge totale, en supplément du compteur déjà installé. Elles s'accompagnent donc d'un coût d'implémentation significatif.

Nous choisissons une méthode alternative, appliquée à toutes les entrées du modèle : on suppose que les données mesurées sont *incomplètes* ou *imprécises*, et on cherche à intégrer l'incertitude résultante dans la démarche d'identification. Les autres grandeurs, notamment les températures et les gains solaires, ne sont pas disponibles *a priori* pour le fournisseur d'énergie, comme indiqué en Figure 3. Ainsi, l'identification n'est plus menée dans un contexte d'observation parfaite du comportement du bâtiment, mais au contraire dans un environnement d'observation incomplet et dégradé : les données sont soit manquantes, soit mesurées avec un certain degré d'imprécision. Cela exige donc un cadre stochastique, plus que déterministe comme pour les autres approches. De par ce parti-pris, cette thèse est donc en accord avec certaines contributions précédentes, telle [FABG14] qui prend en considération l'incertiude sur la température intérieure et les apports internes gratuits, ou encore [BHZ14] qui intègre les données incertaines dans une approche de modélisation par des processus Gaussiens.

Pour améliorer l'observation du bâtiment, notre hypothèse première est alors que le fournisseur d'énergie a accès aux données fournies par les **objets connectés** du bâtiment, et qu'il utilise ces données pour l'identification. Cette hypothèse est détaillée dans la section suivante.

Bien que les objets connectés complètent utilement l'observation du bâtiment, nous ne sommes néanmoins toujours pas dans le contexte d'observation parfaite du bâtiment, comme pour les deux

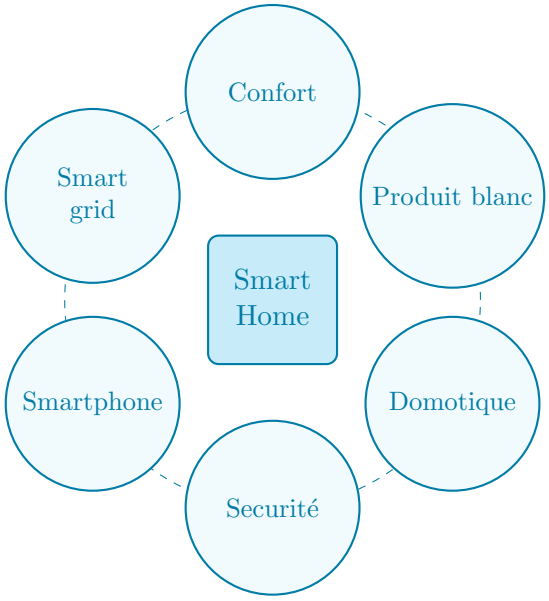

Figure 4: La maison intelligente (*smart home*) : un écosystème connecté.

méthodes mentionnées précédemment. L'idée de la thèse est donc de modéliser les informations apportées par les objets connectés, et de les intégrer dans l'identification. Supposons par exemple que nous disposons d'un côté d'une mesure agrégée de la consommation d'une maison, et des informations fournies par une machine à laver connectée de l'autre. Comment utiliser un tel appareil connectée pour retirer les consommations impropres (usages non-thermiques) de la courbe de charge totale ? Ce travail cherche à répondre à ce genre de questions.

L'internet des objets, source de données et d'incertitude

L'idée de recourir aux objets connectés se justifie par leur prépondérance croissante dans la société. Si on considère en particulier le cas d'un logement, le thermostat connecté commence à être largement répandu : 4.9M d'unités ont été vendues en 2015 d'après IoT Analytics, dont 70% en Amérique du Nord, soit une croissance de 123% entre le quatrième trimestre 2014 et le quatrième trimestre 2015. Ce marché est encore en naissance, mais la multiplicité des entreprises proposant des thermostats connectés ou encore des stations météorologiques connectées (en France, Sowee, Netatmo, Qivivo, Hector, Atlantic, Schneider Electrics, Ween,...) illustre le potentiel de développement de ces produits.

A propos du logement individuel, les objets connectés forment un écosystème définissant la *Smart Home*, ou maison connectée. Ces objets peuvent être classés en six catégories, comme représenté Figure 4. La classe *smart grid* regroupe les compteurs communicants donnant accès à la consommation globale du bâtiment à un pas de temps de quelques minutes. Des compteurs dédiés peuvent aussi mesurer la production locale d'énergie renouvelable, ou la charge d'un véhicule électrique. Ensuite, la classe des objets liés au *confort thermique* englobe le thermostat et la station météorologique conectés. Ces objets contrôlent la température et la qualité de l'air intérieurs. Comparé aux anciens thermostats programmables, le thermostat connecté vise à accroitre la fiabilité du gain d'économies d'énergie, avec de nouvelles fonctionnalités telle l'interface smartphone pour l'utilisateur, le retour sur l'utilisation de l'énergie, la détection de présence, la détection de défauts, la gestion active de la demande ou encore le contrôle multizone [PWM16]. Les *produits blancs*, i.e.

le gros électroménager, forment la troisième catégorie. Ces appareils (lave linge, etc.) commencent à être connectés, avec en application notamment le contrôle distant ou le suivi de consommation. Le ballon d'eau chaude, potentiellement connecté (voir [Bee16] pour une étude) est également inclus dans cette catégorie. La quatrième classe concerne tout ce qui a trait à la *domotique* : contacts sans fils pour la détection d'ouverture des portes et fenêtres, éclairage connecté, etc. Puis, les appareils liés à la *sécurité* regroupent les caméras d'intérieur ou même les serrures connectées. Certaines caméras connectées intègrent par exemple des fonctionnalités de reconnaissance faciale, envoyant des notifications d'identification en conséquence. Enfin, l'objet connecté le plus répandu est tout simplement le *smartphone*, qui peut servir notamment d'interface entre l'utilisateur et ses autres objets connectés.

Suivant cette brève description de l'écosystème formé par les objets connectés, l'hypothèse clé formulée dans cette thèse est que l'énergéticien peut utiliser les informations fournies par ces objets, afin d'améliorer les conditions d'observation du comportement thermique du bâtiment et ainsi identifier un modèle plus précis. Plus précisément, en commençant par les apports thermiques, le compteur communicant fournit au moins une borne supérieure à cette contribution. De plus, les informations des autres objets connectés peuvent permettre de diminuer cette borne, comme détaillé au Chapitre 7. Comment obtenir des informations sur les autres grandeurs du modèle ? Si l'énergéticien équipe également ses clients d'une station météorologique connectée, il aura alors accès à *une* mesure - qui n'est pas *la* mesure - de la température extérieure. De même pour la température intérieure. A noter qu'une station météorologique connectée ne mesure pas le flux solaire incident, de sorte que cette quantité est complètement cachée *a priori*. Finalement, les objets liés à la domotique, voire même la sécurité, peuvent fournir une information sur le nombre d'occupants dans le bâtiment à un instant donné, renseignant ainsi sur les gains métaboliques. Cela reste cependant très prospectif à l'heure actuelle (voir le Chapitre 7).

Du point de vue du fournisseur d'énergie, il semble donc raisonnable de supposer que la seule grandeur mesurée avec précision est la courbe de charge totale, tandis que le flux solaire et les gains métaboliques sont eux au contraire inconnus, et que les apports thermiques et les températures intérieure et extérieure sont mesurées avec un certain degrée d'incertitude. Cet état des lieux correspond à la représentation simplifiée entrées/sortie du modèle de bâtiment donnée en Figure 3. Les contraintes pratiques altèrent donc considérablement les conditions d'application des algorithmes classiques d'identification de bâtiment. Trouver des solutions à ces problèmes soulève par conséquent de nouvelles questions pour l'identification de bâtiment. Quelles sont les informations appportées par les objets connectés pour limiter les contraintes pratiques ? Comment modéliser mathématiquement ces informations ? Comment peut-on adapter la démarche de l'identification afin qu'elle intègre un certain niveau d'incertitude portant sur les entrées du système ? Peut-on estimer précisément un modèle de bâtiment à partir d'un ensemble réduit et peu onéreux de capteurs, malgré l'incertitude intrinsèque de ce dispositif ? Quels sont les capteurs clés d'une bonne estimation ? L'objet de cette thèse est de répondre à de telles interrogations. Les sections suivantes décrivent succinctement les contributions principales de la thèse pour résoudre la problématique exposée dans cette introduction, en suivant l'ordre des chapitres dans le reste du document. Nous présentons donc d'abord la modélisation retenue tout au long de ce travail pour représenter le comportement dynamique d'un bâtiment, puis étudions l'identification de ce modèle d'une part lorsque toutes les grandeurs sont correctement observées, d'autre part lorsque chaque entrée ou sortie du modèle est supposée incertaine.

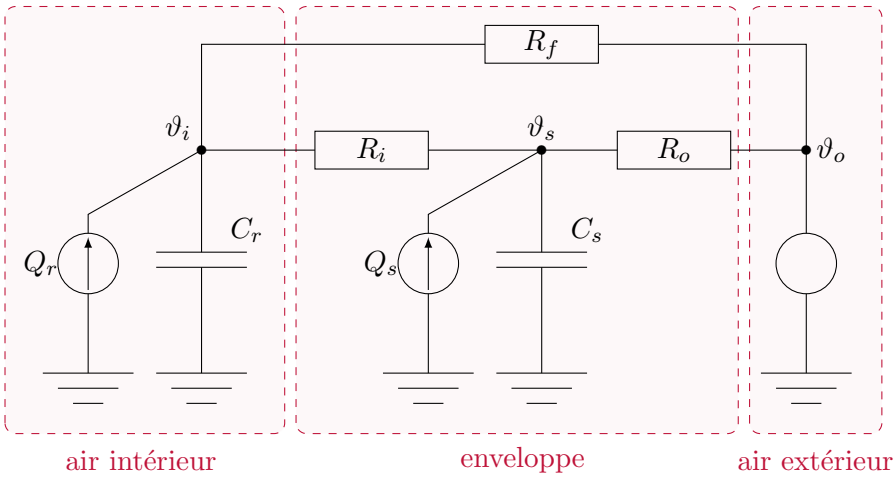

Figure 5: Modèle de bâtiment : le circuit électrique équivalent R3C2.

Modélisation du bâtiment

Le modèle qui sert de support à cette thèse est un circuit électrique équivalent à trois résistances et deux capacités, dénoté R3C2. Il est représenté en Figure 5. Les circuits électriques ont en effet souvent été utilisés pour représenter de manière simplifier la dynamique thermique des bâtiments, voir par exemple [AAC+15, Ber13, FVLA02, LMDP14]. Le modèle R3C2 en particulier a été présenté par EDF en 1991 [MR91] et repris récemment dans la thèse de C. Zayane, [Zay11]. Il fait intervenir trois températures : la température intérieure ϑ_i, la température extérieure ϑ_o et une température de structure ϑ_s représentative de l'état thermique de l'enveloppe du bâtiment. Il y a donc une hypothèse sous-jacente de modèle monozone thermique, c'est-à-dire que ϑ_i est représentative de l'état thermique de l'ensemble des zones du bâtiment. C'est une hypothèse classique en identification de bâtiment. L'air intérieur est chauffé directement par le flux Q_r qui rassemble le système de chauffage Q_h et les gains internes gratuits Q_{fg}, tandis que le chauffage par gains solaires Q_s est injecté indirectement au nœud de température de structure du bâtiment.

L'avantage de la représentation R3C2 est qu'elle cumule interprétabilité physique et faible coût calculatoire, ce qui justifie son choix dans la thèse, conformément aux éléments de contexte précédemment évoqués. En effet, le modèle est d'ordre 2, ordre minimal suffisant pour reproduire la dynamique thermique [MCPF10]. Il comporte cinq paramètres regroupés dans le vecteur $\theta := \begin{pmatrix} 1/C_r & 1/C_s & 1/R_f & 1/R_o & 1/R_i \end{pmatrix}$, à partir desquels deux constantes de temps peuvent être calculées, une lente τ_1 et une rapide τ_2, ainsi que deux paramètres statiques, le coefficient de déperditions thermiques UA et la transmittance solaire g. Ces calculs sont présentés dans le Chapitre 2. On parle en ce sens de modèle "*boîte grise*", hybride entre les modèles purement physiques ("boîte blanche") ou purement statistiques ("boîte noire"). A noter cependant que le propos de la thèse n'est pas d'affirmer que le modèle R3C2 est le "meilleur" modèle de bâtiment. Il s'agit avant tout de montrer comment adapter l'estimation d'un modèle classique de bâtiment lorsque les données sont manquantes ou partielles. En particulier, les outils présentés dans la suite de la thèse sont valables pour tout type de modèle électrique équivalent.

Dans le Chapitre 2, on montre également que le circuit R3C2 se met en équation sous la forme

d'une représentation d'état discrète linéaire Gaussienne de la forme

$$X_t = A(\theta)X_{t-1} + B(\theta)U_t + V_t, \tag{1}$$
$$Y_t = CX_t + W_t, \tag{2}$$

où $X := \begin{pmatrix}\vartheta_i & \vartheta_s\end{pmatrix}^\intercal$ est le vecteur des *états cachés*, non-observés, $Y := T_i$ est une mesure *bruitée* de la température intérieure de modèle ϑ_i, $U := \begin{pmatrix}\vartheta_o & Q_r & Q_s\end{pmatrix}^\intercal$ le vecteur des entrées mesurées du modèle, et V et W sont des bruits blanc Gaussiens indépendants. t désigne l'indice temporel discret. Les matrices A et B sont paramétrées par θ, mais pas la matrice d'observation C (simple retour d'état).

Etant donnée la classe de modèle fixée par le circuit R3C2, l'enjeu de l'identification de bâtiment est donc le suivant : être capable d'estimer les paramètres θ du modèle à partir de données d'observation des grandeurs d'entrée U_t et de sortie Y_t du modèle, et malgré la non-connaissance de l'état X_t du système. Un jeu de données typique est représenté en Figure 6. Le Chapitre 3, présenté ci-après, traite cette problématique lorsque toutes ces grandeurs sont effectivement observées, tandis que les chapitres 4 à 7 proposent des solutions lorsque les données sont seulement partiellement observées.

Identification boucle ouverte d'un bâtiment observé en boucle fermée

Lorsque toutes les grandeurs d'entrée/sortie du système sont correctement observées, une difficulté pratique subsiste malgré tout : la question se pose de savoir si le bâtiment doit être identifié en *boucle ouverte* ou en *boucle fermée*, ce que la littérature sur l'identification de bâtiment ne tranche pas véritablement. L'identification boucle fermée correspond à l'exploitation réelle du bâtiment en saison de chauffe, comme illustré Figure 2. Cependant, la boucle fermée implique que le contenu informatif du signal de commande (la puissance de chauffe) est moindre comparé à un signal choisi librement en boucle ouverte : cette dernière situation constitue par conséquent de manière générale l'approche préférentielle pour l'identification de modèle. Elle reste toutefois difficile à mettre en œuvre en pratique dans le cas particulier d'un bâtiment habité ou occupé, puisqu'elle est intrusive (le fournisseur détermine la séquence de chauffage à la place de l'utilisateur) et génère un inconfort thermique (la température intérieure ne suit plus une certaine consigne).

Nous montrons ainsi dans le Chapitre 3 que cette problématique est peu traitée dans l'état de l'art, avec soit des approches reposant sur des expérimentations en boucle ouverte (e.g. [BM11]), soit le recours à un modèle physique détaillé du bâtiment pour obtenir des données simulées en boucle ouverte (e.g. [HGP12]), soit des approches en boucle fermée (e.g. [Zay11]) qui peinent à intégrer la saturation de la commande de chauffage.

Dans le Chapitre 3, on se place donc dans le cas où les données sont générées en boucle fermée, cas le plus probable en pratique. La contribution principale de ce chapitre est de justifier l'estimation des paramètres du modèle de bâtiment comme si les données étaient en fait générées en boucle ouverte, bien que ce ne soit pas le cas en réalité. Cette approche est appelée *approche directe* en identification. L'argumentation repose sur les propriétés de l'estimateur du maximum de vraisemblance (MLE, *Maximum Likelihood Estimator*). Si les données sont générées en boucle ouverte, alors les propriétés statistiques asymptotiques (consistance, normalité) du MLE d'un système linéaire tel que le circuit R3C2 sont déjà établies dans la littérature, voir notamment [HD88, Cai87]. Ce n'est plus le cas lorsque les données sont générées en boucle fermée, en raison de la nouvelle structure de dépendance des variables, illustrée en Figure 7 : le processus d'entrée U_t n'est plus exogène, mais dépend de l'observation passée de $\{Y_{t-1}\}$.

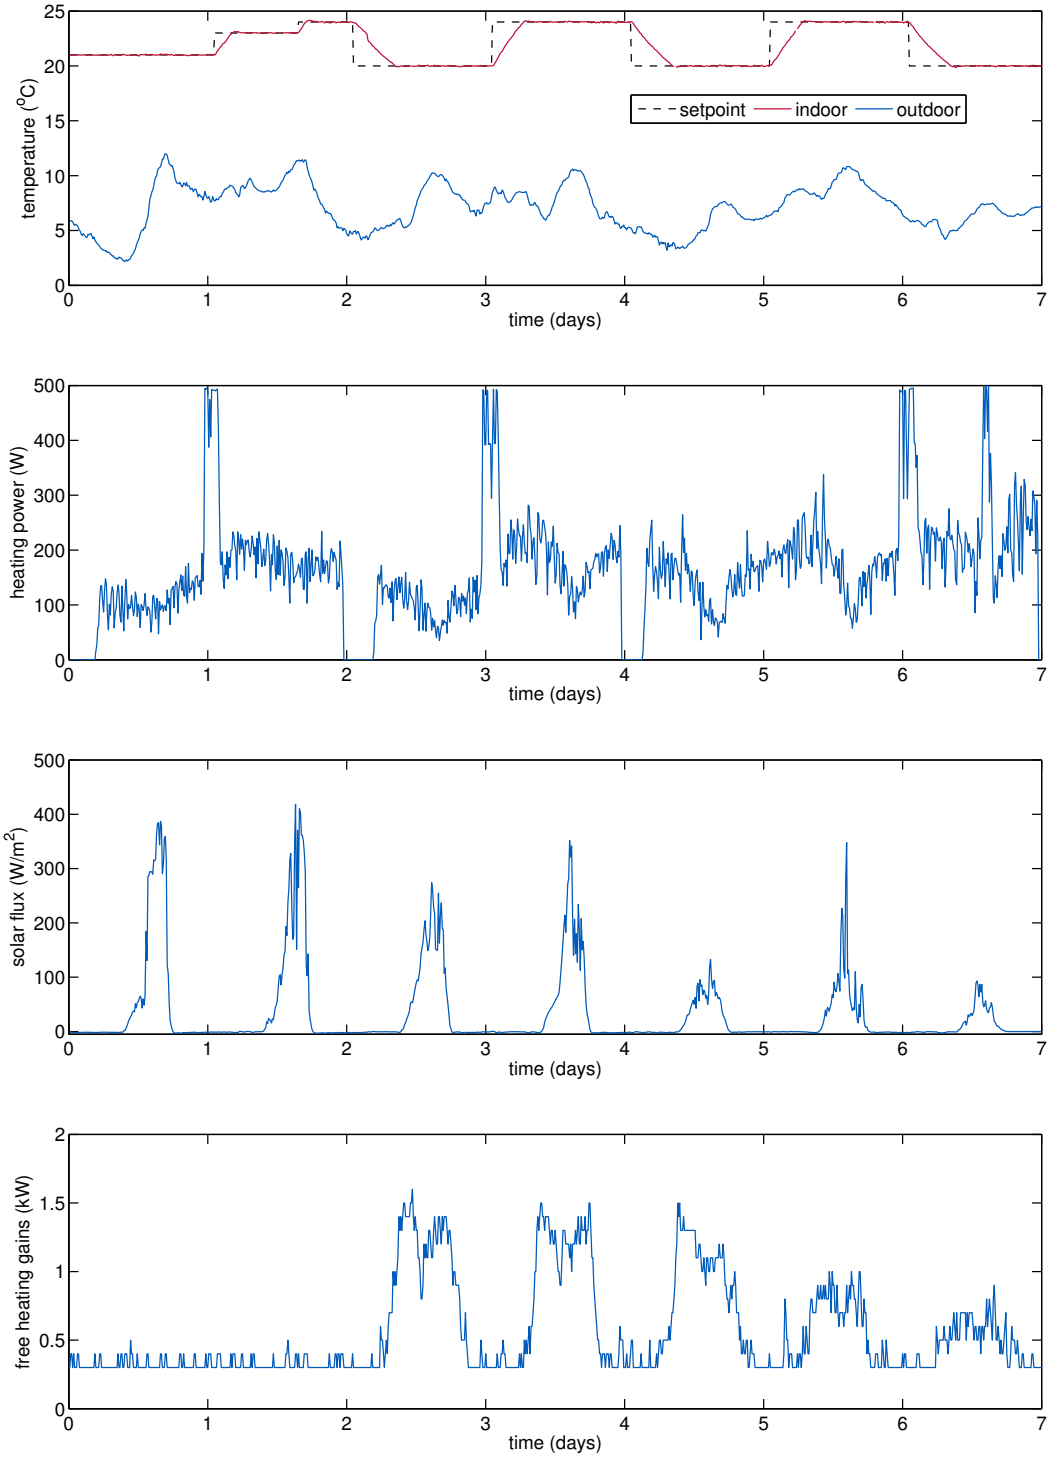

Figure 6: Exemple de données pour l'identificiation. De haut en bas : température intérieure T_i, température extérieure ϑ_0, consommation du système de chauffage Q_h, gains solaires at apports internes gratuits Q_{fg}. Données mesurées sur le site d'EDF R&D Les Renardières, laboratoire BESTLab.

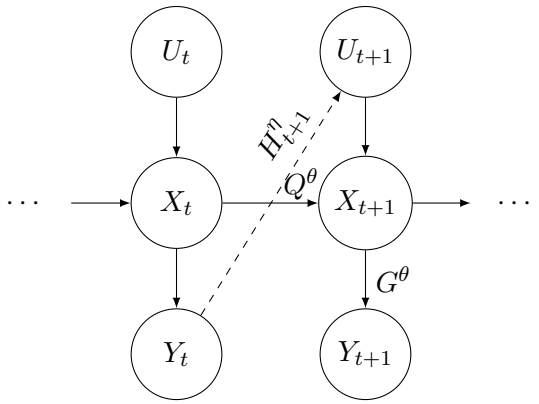

Figure 7: Graphe de dépendance d'un modèle de Markov caché avec données générées en boucle fermée. Le processus $\{Y_t\}$ est l'observation, $\{X_t\}$ est l'état caché, et le processus d'entrée $\{U_t\}$ est généré à partir de la mesure passée de $\{Y_{t-1}\}$. Le cas boucle ouverte correspond à la suppression du noyau $H^\eta_{t+1}(Y_t, U_{t+1})$.

Pour ces raisons, il faut redémontrer les propriétés asymptotiques du MLE dans le cadre de la boucle fermée. Cependant, établir de tels résultats va au-delà du sujet de la thèse. Sous l'hypothèse que la boucle de régulation ne dépend pas des paramètres θ du système en boucle ouverte, le Chapitre 3 montre toutefois que si la consistance et la normalité asymptotique du MLE en boucle fermée sont admises, alors

1. maximiser la vraisemblance des observations générées en boucle fermée revient à maximiser la vraisemblance obtenue avec les mêmes données mais sous l'hypothèse boucle ouverte;

2. le calcul de la covariance asymptotique du MLE en boucle fermée utilise la même statistique que celle du MLE sous hypothèse boucle ouverte (à savoir la Hessienne de la log-vraisemblance sous hypothèse boucle ouverte).

Ainsi, malgré l'exploitation réelle du bâtiment en boucle fermée, il semble raisonnable d'estimer par maximum de vraisemblance les paramètres θ du circuit R3C2 *comme s'il était en boucle ouverte*, et d'utiliser le théorème central limite pour calculer des intervalles de confiance. Une telle procédure d'estimation est décrite dans le Chapitre 3, en utilisant l'algorithme espérance-maximisation (EM, *Expectation-Maximization*). Les résultats numériques, aussi bien sur données simulées que réelles, confirment la validité de l'approche. Différentes conditions expérimentales sont testées afin de quantifier la qualité des estimations en fonction du contenu informatif des données. Il apparait qu'une estimation précise des paramètres statiques et dynamiques du modèle thermique du bâtiment est possible à partir de 10 à 14 jours d'expérimentation, pourvu que la température de consigne présente (au moins) deux niveaux séparés de quelques degrés.

Ce chapitre montre ainsi de manière empirique qu'il n'est pas nécessaire de modéliser la régulation pour l'identification d'un bâtiment observée en conditions réelles d'exploitation, lorsque les signaux d'entrée/sortie sont bien tous mesurés.

Prise en compte des contraintes pratiques

Le reste de la thèse est ensuite consacré aux contraintes qui empêchent la mise en œuvre pratique de l'algorithme d'identification du Chapitre 3. Nous remettons ainsi successivement en cause la disponibilité effective de chaque entrée et sortie du modèle, et proposons des algorithmes adaptés.

Température intérieure cachée

Supposons dans un premier temps que la température intérieure n'est pas mesurée, tandis que toutes les entrées du modèle le sont. C'est une hypothèse contraignante, puisque cette grandeur est la sortie du modèle thermique de bâtiment. En particulier, l'approche boucle ouverte du Chapitre 3 ne s'applique plus, et il faut déterminer un modèle pour la régulation de la température à partir d'une consigne. Une solution proposée dans la thèse de C. Zayane [Zay11] est donc d'adopter un modèle boucle fermée dont la sortie est la commande de chauffage, et les entrées sont les conditions météorologiques, les apports thermiques internes gratuits, i.e. autres que le système de chauffage, ainsi que la température de consigne. Cependant une limite du travail [Zay11] est que la régulation prend la forme d'un correcteur Proportionnel Intégral (PI) linéaire. Or la commande est en réalité non-linéaire, car saturée : elle est forcément positive (pas de refroidissement) et bornée (limitation par la puissance nominale du système de chauffage). La méthode semble donc perfectible, notamment dans l'optique de la création de services énergétiques tels que la prédiction de consommation ou le lissage de la courbe de charge.

Pour mieux prendre en compte cette contrainte, et la non-linéarité qu'elle génère, nous proposons dans le Chapitre 4 d'appliquer une méthode d'inversion Bayésienne, avec un algorithme de type Markov Chain Monte Carlo, l'échantillonneur de Monte Carlo à dynamique Hamiltonienne (HMC, [Nea11]). Cette méthode a déjà été étudiée pour un circuit R1C1 dans [Laj11], et le Chapitre 4 étendons donc l'étude dans le cas du R3C2.

Pour cela, la régulation Q_h est représentée par un correcteur PI dont la sortie est saturée par une fonction sigmoïdale, approximation $\mathcal{C}^\infty$ de la fonction de saturation classique :

$$Q_h = \Phi\left[K\left(T_r - \vartheta_i + \frac{1}{\tau}\int(T_r - \vartheta_i)\right)\right], \tag{3}$$

où K et τ sont respectivement le gain et la constante de temps du PI, T_r est la température de consigne, ϑ_i la température intérieure et $\Phi : x \in \mathbb{R} \mapsto \frac{Q_h^{max}}{1+\exp(-\lambda(x-Q_h^{max}/2))} \in (0, Q_h^{max})$, pour une puissance limitante Q_h^{max}, avec $\lambda := 4/Q_h^{max}$. Le modèle du bâtiment est donc une représentation d'état non linéaire avec sept paramètres, les paramètres θ du circuit R3C2 étant augmentés de K et τ. Ses états cachés sont $X := \begin{pmatrix}\vartheta_i & \vartheta_s & d\end{pmatrix}^\intercal$, où $d := \frac{1}{\tau}\int(T_r - \vartheta_i)$, tandis que les entrées et la sortie observés sont respectivement $U := \begin{pmatrix}\vartheta_o & Q_{fg} & Q_s & T_r\end{pmatrix}^\intercal$ et $Y := Q_h$. Les notations sont celles utilisées lors de la description du modèle R3C2. Le Chapitre 4 détaille la méthode Bayésienne d'apprentissage du modèle, en utilisant l'algorithme HMC. L'avantage de cet algorithme est qu'il utilise des informations de gradient pour explorer plus efficacement la distribution *a posteriori* des paramètres. De manière originale, nous proposons d'estimer conjointement les paramètres et les états cachés du système, notamment la température intérieure ϑ_i, par l'algorithme HMC. Ainsi, la densité *a posteriori* considérée est $p(\theta, x_{0:\mathrm{T}}|y_{1:\mathrm{T}})$, où T désigne le nombre d'observations. De plus, des priors sur la trajectoire des états sont également ajoutés. En particulier, le prior sur la température intérieure pénalise fortement les écarts à la température de consigne, que la température intérieure est supposée suivre. Ce prior très informatif favorise la convergence de l'algorithme.

Les résultats numériques sur des données simulées sont encourageants, avec une bonne estimation des constantes physiques du modèle R3C2. La régulation est plutôt bien estimée également, notamment la constante de temps τ. Ainsi, la prédiction de la température intérieure sur des données de test est précise, avec par exemple une erreur quadratique moyenne de 0.35°C pour la trajectoire représentée en Figure 8.

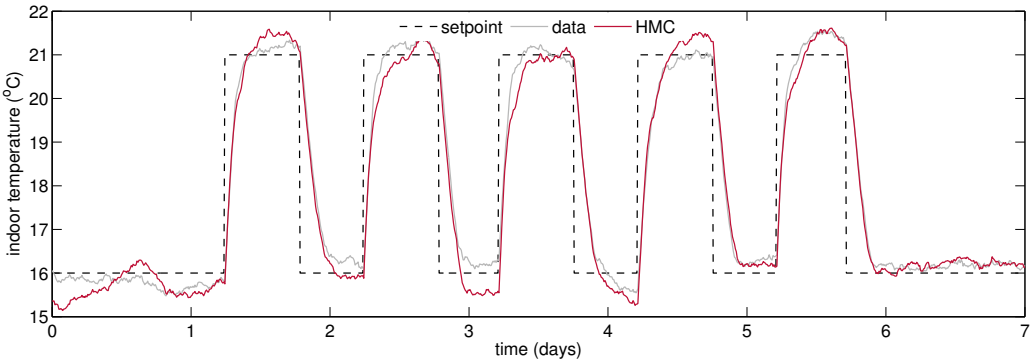

Figure 8: Exemple de prédiction de la température intérieure sur données test.

Réduction du biais sur la température extérieure

Après avoir envisagé dans le Chapitre 4 que la sortie du modèle de bâtiment n'est pas mesurée, les chapitres 5 à 7 examinent le cas des trois entrées du modèle R3C2, à savoir la température extérieure, les apports solaires et les gains thermiques internes. La démarche est similaire pour ces trois grandeurs : contrairement au cas de la température intérieure, nous disposons ici d'une version, certes *incertaine*, de la grandeur d'intérêt. Cela signifie en particulier que nous pouvons appliquer le résultat du Chapitre 3. Ainsi, bien que les données soient toujours générées en boucle fermée, nous proposons dans la suite des algorithmes d'estimation du modèle R3C2 en *boucle ouverte*.

Le Chapitre 5 traite plus spécifiquement le cas de la température extérieure. Dans le contexte des bâtiments connectés, la mesure de cette grandeur peut s'obtenir de deux façons différentes. La première possibilité est de disposer d'un capteur de température, par exemple dans le cadre d'une station météorologique connectée, capteur situé au niveau du bâtiment en question. En l'absence de mesure *in-situ*, on peut envisager de prendre la mesure du plus proche capteur dans une base de donnée accessible en open-source (voir par exemple la Weathermap de Netatmo). Dans les deux cas, le fournisseur d'énergie a accès à *une* mesure de la température extérieure, cette mesure n'étant pas forcément *la* mesure représentative du micro-climat local au niveau du bâtiment étudié. Cela se comprend aisément dans le cas de la seconde solution proposée : le champ de température extérieure n'est pas uniforme, même sur des distances de quelques kilomètres seulement. Cela reste valable dans le premier cas également. En effet, bien que le capteur connecté soit situé au niveau même du bâtiment, le fournisseur d'énergie ne connait pas la position exacte du capteur, qui est déterminée par l'utilisateur final. Or, il est nécessaire de placer le capteur à l'ombre, pour ne pas mélanger l'influence du rayonnement solaire. Ainsi, un capteur exposé au soleil présentera un fort biais de mesure. La Figure 9 montre que ce biais peut être élevé, jusqu'à 6°C, même au mois de novembre.

Du point de vue du fournisseur d'énergie, la mesure disponible de la température extérieure est donc possiblement trompeuse. Pour prendre en compte cette incertitude, nous introduisons une représentation dynamique de la température extérieure, au pas de temps 10 minutes, sous la forme d'une représentation d'état :

$$\vartheta_{o,t} = \vartheta_{o,t-1} + \varepsilon_t, \tag{4}$$
$$\vartheta_{b,t} = \vartheta_{b,t-1} + \epsilon_t, \tag{5}$$
$$T_{o,t} = \vartheta_{o,t} + I_t \vartheta_{b,t} + \eta_t, \tag{6}$$

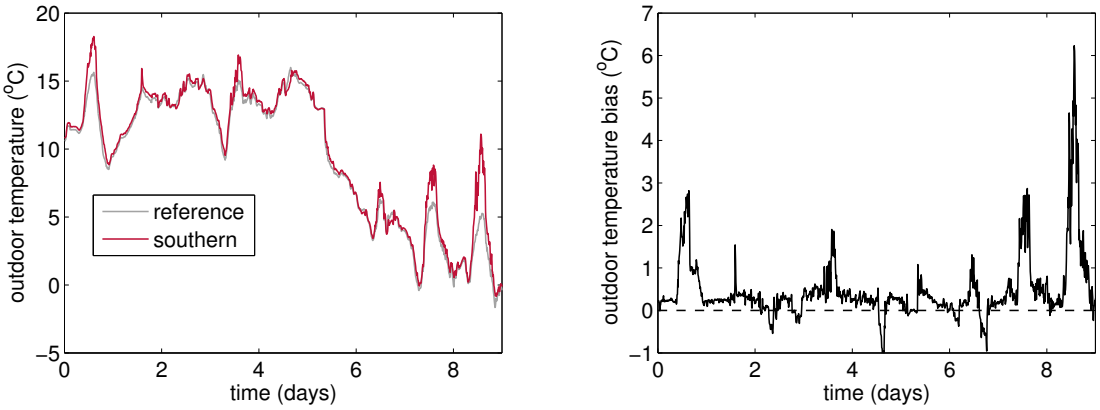

(a) (b)

Figure 9: Exemple d'écart de mesure entre deux capteurs de température extérieure. (a) Capteur de référence protégé du rayonnement (gris clair) et capteur sur une façade orientée Sud, sans cache (rouge). (b) Ecart entre ces deux mesures. Laboratoire EDF R&D BESTLab, novembre 2015.

où, T_o est la mesure disponible de la température extérieure, ϑ_o est la température extérieure de référence du modèle R3C2, ϑ_b est le processus représentatif du biais et I_t est une indicatrice valant 0 ou 1. ε_t, ϵ_t et η_t sont ici des bruits blanc Gaussiens indépendants. Ainsi, l'observation T_o est une version bruitée de la température de référence ϑ_o, éventuellement corrigée par le biais ϑ_b. Au pas de temps 10 minutes, on suppose que les états cachés ϑ_o et ϑ_b sont des marches aléatoires indépendantes. Deux hypothèses sont examinées dans le Chapitre 5 concernant I_t:

1. de façon déterministe, on fixe $I_t = 0$ la nuit, et $I_t = 1$ le jour. On observe en effet, au moins dans le cas d'un capteur placé au niveau du bâtiment mais non-protégé du rayonnement solaire, que l'écart de mesure est très réduit durant la nuit (voir par exemple la Figure 5.1).

2. I_t est un processus non-stationnaire, indépendant, à deux états, avec à chaque instant t les probabilités *a priori* $\pi_j(t) := P(I_t = j)$, $j = 0, 1$. Le système alterne donc de façon stochastique entre les deux configurations $I_t = 0$ ou $I_t = 1$. L'enjeu est alors de calculer les probabilités filtrées $\pi_j(t|t) := P(I_t = j|Y_1, \ldots, Y_t)$, $j = 0, 1$. Il s'agit d'une relaxation de la première hypothèse.

Dans les deux cas, le modèle dynamique de la température extérieure est ajouté au modèle R3C2 du bâtiment, formant une représentation d'état linéaire ou conditionnellement linéaire, selon l'hypothèse sur I_t, ayant pour sorties et entrées observées $Y := \begin{pmatrix} T_i & T_o \end{pmatrix}^\intercal$ et $U := \begin{pmatrix} Q_r & Q_s \end{pmatrix}^\intercal$, respectivement, et pour états cachés $X := \begin{pmatrix} \vartheta_i & \vartheta_s & \vartheta_o & \vartheta_b \end{pmatrix}^\intercal$. L'avantage de la modélisation choisie est qu'elle n'introduit pas de paramètres supplémentaires aux 5 paramètres du circuit R3C2, tout en permettant une estimation de la température extérieure de référence ϑ_o via le filtrage-lissage de l'état. Les deux modèles sont appris par l'algorithme EM [SS91].

Les résultats numériques, sur des données simulées, mettent en avant les bénéfices des deux approches suggérées dans cette thèse, par rapport à une approche naïve, qui reproduirait l'estimation décrite au Chapitre 3 (observation parfaite du bâtiment) en considérant que la mesure disponible est une entrée exogène du modèle de bâtiment. En particulier, quelque soit l'hypothèse fait sur I_t, le biais d'estimation du coefficient de déperditions statiques UA est bien réduit par rapport à la méthode naïve, qui le surestime. Cette surestimation est par ailleurs cohérente avec l'interprétation physique

de UA : à consommation de chauffage constante, le modèle est identifié avec une température extérieure plus chaude qu'en réalité, ce qui veut dire que le bâtiment estimé est moins performant thermiquement qu'en réalité, donc UA augmente. De manière générale, l'hypothèse stochastique sur I_t permet la meilleure identification des paramètres R3C2. Enfin, la trajectoire de la température extérieure de référence ϑ_o est bien estimée par les deux méthodes, avec les valeurs extrêmes du biais qui sont éliminées.

Apprentissage de la dynamique du flux solaire

Le Chapitre 5 montre notamment que la température extérieure de référence peut être déduite de la donnée d'un capteur possiblement mal orienté, ainsi que de la mesure du flux solaire et de l'observation du comportement thermique du bâtiment. Dans le Chapitre 6, nous proposons de renverser le raisonnement pour le cas des *gains solaires*. En se donnant *deux capteurs* de température extérieure, un *volontairement exposé* au soleil, au niveau des surfaces vitrées, et un second *volontairement abrité* et protégé du rayonnement solaire, on cherche alors à estimer le modèle thermique du bâtiment sans mesure directe du flux solaire, i.e. sans pyranomètre. L'idée sous-jacente est que la différence entre les deux températures extérieures est essentiellement due au flux solaire : plus le rayonnement solaire incident est fort, plus cette différence est grande. L'intérêt de la démarche est donc de remplacer le coûteux pyranomètre par un second capteur de température extérieure, au coût réduit.

L'approche développée dans le Chapitre 6 est similaire à celle du Chapitre 5. Il s'agit par conséquent de construire un modèle du flux solaire à partir de l'observation du biais ΔT_o de température extérieure (i.e. la différence entre les deux capteurs de température extérieure), ce modèle étant ensuite intégré dans la représentation d'état du circuit R3C2. Les principales étapes de la modélisation sont les suivantes :

1. Représentation physique du flux solaire global horizontal G_{cld} avec correction par la couverture nuageuse :

$$G_{h,cld} = G_{h,clr}(1 - aN^b), \tag{7}$$

   où $G_{h,clr}$ est le flux solaire global horizontal par ciel clair, qui ne dépend que des coordonnées géographiques et temporelles du bâtiment, $N \in [0, 1]$ est la couverture nuageuse, ou nébulosité. Dans cette thèse, on choisit le modèle de Perrin de Brichambault [PdBV82, RW99] pour calculer $G_{h,clr}$.

2. Reconstruction de la nébulosité à partir du biais ΔT_o de température extérieure :

$$N = \phi(\Delta T_o), \text{ avec } \phi : x \in \mathbb{R} \mapsto 1 - \frac{1}{1 + \exp(-\alpha(x - x_0))} \in (0, 1). \tag{8}$$

   Ce modèle est paramétré par x_1 et x_2, où $\alpha = 4/(x_2 - x_1)$ et $x_0 = (x_1 + x_2)/2$. Il assimile ΔT_o à un indice de clarté, dans le sens où les grandes valeurs $\Delta T_o > x_2$ correspondent aux ciels les plus clairs (N proche de 0), et les faibles valeurs $\Delta T_o < x_1$ aux ciels les plus sombres (N proche de 1).

3. Estimation des paramètres a, b, x_1 et x_2. Ces paramètres peuvent être fixés *a priori* (valeurs par défaut valables pour de larges zones géographiques) ou appris si des mesures *in situ* du rayonnement solaire sont disponibles.

4. Représentation d'état de ΔT_o :

$$\eta_t = \eta_{t-1} + (q_\eta + \sigma_t) v_{\eta,t}, \tag{9}$$
$$\Delta T_{o,t} = \eta_t + w_{\eta,t}, \tag{10}$$

où v_η est un bruit blanc Gaussien unitaire indépendant du bruit blanc Gaussien d'observation w_η de variance r_η^2, $\Delta T_{o,t}$ est l'observation bruitée de l'état caché η_t. La dynamique de l'état η est une marche aléatoire dont la variance des incréments vaut q_η^2 au minimum, sauf à quelques instants inconnus, où cette variance vaut $(q_\eta + \sigma_t)^2$, pour $\sigma_t \geq 0$. On suppose en effet que le processus σ vaut exactement 0 pour une majorité d'instants, tandis que quelques valeurs positives modélisent les sauts dans la dynamique de η. Cette hypothèse permet de représenter plus finement l'instabilité climatique au pas de temps 10 minutes (voir la Figure 9(b)), avec ponctuellement des phases de plus grande variabilité.

5. Estimation de $\beta := \{q_\eta, r_\eta, \sigma_1, \ldots, \sigma_\mathrm{T}\}$ en pénalisant la log-vraisemblance $\log L_\beta$ du modèle d'état (9):

$$\widehat{\beta}(\lambda) := \arg\min_\beta \left\{ -2 \log L_\beta(\Delta T_{o,1}, \ldots, \Delta T_{o,\mathrm{T}}) + \lambda \sum_{t=1}^{\mathrm{T}} |\sigma_t| \right\}. \tag{11}$$

En choisissant soigneusement le coefficient λ, la pénalité $\lambda \sum_{t=1}^{\mathrm{T}} |\sigma_t|$ assure une structure parcimonieuse à la série des σ_t.

Une fois connus les coefficients a, b, x_1, x_2 et β, le modèle global est formé en ajoutant le modèle (9) aux équations du circuit R3C2. Cela résulte en un modèle d'état conditionellement linéaire et Gaussien, puisque conditionnellement à la dynamique du flux solaire, non-linéaire, la dynamique du bâtiment est linéaire. Les entrées et sorties observées du modèle sont respectivement $U := \begin{pmatrix} \vartheta_o & Q_r & Q_{clr} \end{pmatrix}^\intercal$ et $Y = \begin{pmatrix} T_i & \Delta T_o \end{pmatrix}^\intercal$, les états cachés sont $X = \begin{pmatrix} \vartheta_i & \vartheta_s & \eta \end{pmatrix}^\intercal$. Il reste alors à estimer les 5 paramètres du circuit R3C2. Bien que non-linéaire, la structure de linéarité conditionnelle est exploitée pour appliquer l'algorithme EM et résoudre l'étape E le plus efficacement possible en utilisant un lissage particulaire Rao-Blackwellisé [LS10b, Lin11]. Cet algorithme de la littérature fournit une approximation numérique des densités de filtrage et de lissage des états cachés, avec une réduction de la variance par rapport au lissage particulaire simple appliqué aux modèles non-linéaires; il est détaillé dans le Chapitre 6.2.

Les résultats numériques, sur des données simulées et réelles, montrent l'intérêt de la démarche présentée dans ce chapitre. En particulier, ils confirment qu'une bonne estimation du flux solaire reçu peut être déduite à partir de ΔT_o. Un exemple d'estimation est reproduit en Figure 10. A noter que l'estimation de la trajectoire lissée de η se fait dans le modèle global, et non pas uniquement à partir de la représentation purement "météorologique" (9), ce qui permet de bénéficier du contenu informatif lié à l'observation du comportement thermique du bâtiment. De même, l'estimation des paramètres physiques du circuit R3C2 est largement améliorée par cette méthode, en comparaison à des approches classiques de la littérature. En particulier, la variance des estimations est significativement réduite grâce au lissage particulaire Rao-Blackwellisé. L'ensemble de l'étude numérique conforte donc l'idée qu'une solution simple au coût prohibitif du pyranomètre consiste à le remplacer par un capteur de température extérieure, sous réserve de l'exposer volontairement au rayonnement solaire et de conserver le capteur de température extérieure de référence.

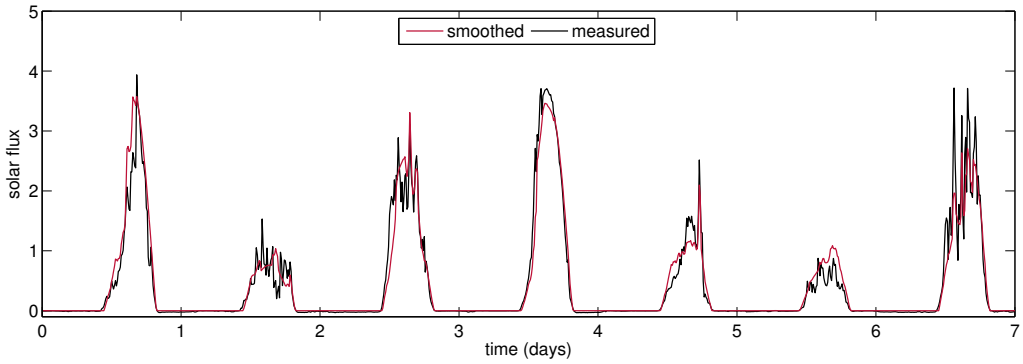

Figure 10: Exemple d'estimation du flux solaire: en noir la mesure réelle (non utilisée pour l'apprentissage du modèle R3C2), et en rouge la trajectoire obtenue à partir de l'état lissé η.

Incertitude sur la courbe de charge

Le dernier chapitre de la thèse est consacré à la prise en compte de l'incertitude sur les apports thermiques internes du bâtiment, du point de vue du fournisseur d'énergie. Celui-ci mesure la consommation totale du bâtiment, à un pas de temps infra-horaire, voire de quelques minutes. Cette donnée constitue ainsi une borne supérieure des apports thermiques internes, puisque la consommation totale correspond à des usages qui contribuent ou non au chauffage. Pour un meilleur apprentissage du modèle, l'énergéticien cherche donc à estimer les apports non-thermiques, afin de les retirer de la consommation totale et obtenir une courbe de charge *utile*. A noter que d'après les résultats du Chapitre 3, il n'est pas nécessaire de distinguer parmi les apports thermiques, notamment entre système de chauffage et apports gratuits, le modèle R3C2 pouvant être identifié directement en boucle ouverte, avec l'ensemble des contributions thermiques internes agrégées dans le terme noté Q_r. Par ailleurs, les gains métaboliques, liés à l'occupation effective du bâtiment, ne sont pas étudiés dans ce chapitre.

Sans information supplémentaire, la courbe de charge totale est une première approximation de la courbe de charge utile, et peut donc être utilisée telle quelle dans un algorithme d'identification tel que celui développé dans le Chapitre 3. Pour aller plus loin, nous nous plaçons dans le Chapitre 7 dans le cas d'un bâtiment connecté. Nous considérons plus précisément le cas d'une maison connectée. En s'appuyant sur une typologie des objets connectés d'après le niveau actuel de maturité des technologies correspondantes, nous formulons alors plusieurs hypothèses dont les deux principales sont les suivantes :

1. l'énergéticien mesure la courbe de charge totale à un pas de temps infra-horaire;
2. l'énergéticien a accès à une *chronologie* spécifiant les événements de type Marche/Arrêt des charges qui ne représentent pas un apport thermique.

Il apparait en effet que contrairement aux apports thermiques, les autres charges correspondent surtout à des usages ponctuels, et ne sont pas continuellement présentes. C'est le cas par exemple du lave-vaisselle, de la machine à laver ou encore du ballon d'eau chaude. Ainsi, la seule information requise consiste à savoir à quels moments ce type de ce charge est activé puis désactivé.

D'après ces hypothèses, l'énergéticien sait à chaque instant si la consommation agrégée correspond uniquement à des apports thermiques, quels qu'ils soient, ou à un mélange d'apports thermiques et non-thermiques, en proportions inconnues. Il existe donc des intervalles de temps connus

pour lesquels *la courbe de charge mesurée est la courbe de charge utile*. L'algorithme d'estimation exploite cette connaissance.

Dans un premier temps, l'algorithme classique d'estimation du modèle R3C2 tel que présenté au Chapitre 3 s'applique aux instants où il n'y a pas de charges non-thermiques.

Cela fournit une première estimation des paramètres θ du circuit, à partir de l'observation exacte du système. Nous suggérons cependant de corriger cette estimation, en tirant profit de toutes les données. En particulier, les données mesurées en présence d'apports non-thermiques n'ont pas encore été exploitées. Le schéma correctif consiste alors à apprendre un modèle de régulation, puis à simuler une trajectoire de température intérieure à partir de la connaissance de la température de consigne, du modèle de bâtiment et de la régulation. Il suffit ensuite de mettre à jour l'estimation de θ en comparant d'un côté la simulation de la température avec la trajectoire observée, et de l'autre les puissances de chauffage simulée et observée (en se restreignant toutefois aux instants où la courbe de charge totale est exactement la courbe utile). La démarche globale est itérative, en alternant correction locale des paramètres θ du circuit R3C2 et estimation d'un modèle de régulation. Ces différentes étapes sont détaillées en Section 7.2.

L'un des éléments clés réside dans l'apprentissage de la régulation. Nous proposons une stratégie en deux étapes :

1. Pour chaque intervalle de temps durant lequel la courbe de charge mesurée est la courbe de charge utile, apprentissage d'un correcteur PI par moindres carrés. Cette étape fournit donc un ensemble d'estimateurs de la régulation.

2. Agrégation à poids exponentiels des modèles PI calculés lors de l'étape 1.

L'agrégation procède par analogie avec la théorie d'agrégation séquentielle d'experts [CBL06, Gai15]. Dans la théorie classique, les experts sont des "boîtes noires" proposant uniquement une séquence de prédictions, et l'agrégation vise à former une nouvelle prédiction par mélange des prédictions expertes. Ici, les experts sont de structure connue, ce sont les modèles PI appris lors de la première étape, et les prédictions sont les températures intérieures simulées à partir des PI experts et de l'estimation du R3C2. Cependant, plutôt que d'obtenir une nouvelle prédiction agrégée, l'objectif est de fournir en étape 2 un nouveau modèle PI dont les paramètres se rapprochent des paramètres du meilleur PI expert. L'idée sous-jacente est d'être capable de sélectionner à chaque instant le meilleur expert PI, sachant que certains experts peuvent être très mal appris. En effet, le PI est estimé à partir de la courbe de *charge utile*, alors qu'il faudrait l'apprendre à partir de la courbe de *charge chauffage*, i.e. retirer les apports gratuits, non-contrôlés. Chaque PI expert est donc plus ou moins bien appris, en fonction notamment de la qualité des données, i.e. du taux d'apports internes gratuits sur l'intervalle considéré. Le modèle agrégé est calculé comme mélange des modèles experts, où les poids de mélange sont ceux obtenus pour l'agrégation des prédictions. Les résultats théoriques de l'agrégation d'experts présentés dans [CBL06, Gai15] ne sont donc plus valables dans ce contexte. Toutefois, les résultats empiriques obtenus au Chapitre 7 valident la démarche.

Plus spécifiquement, l'étude numérique porte sur des données simulées et réelles. L'analyse d'une expérience sur un jeu de données simulé montre que la méthode proposée dans ce chapitre améliore significativement l'estimation des quatre paramètres physiques du modèle R3C2, tant par rapport à une approche classique de l'état de l'art que par rapport à l'estimation naïve sans le schéma correctif précédemment décrit. Les trajectoires simulées de la courbe de charge utile et de la température intérieure, reproduites en Figure 11 correspondent également aux observations. Enfin, l'analyse de l'agrégation illustre la capacité de la méthode à sélectionner les meilleurs modèles PI à chaque

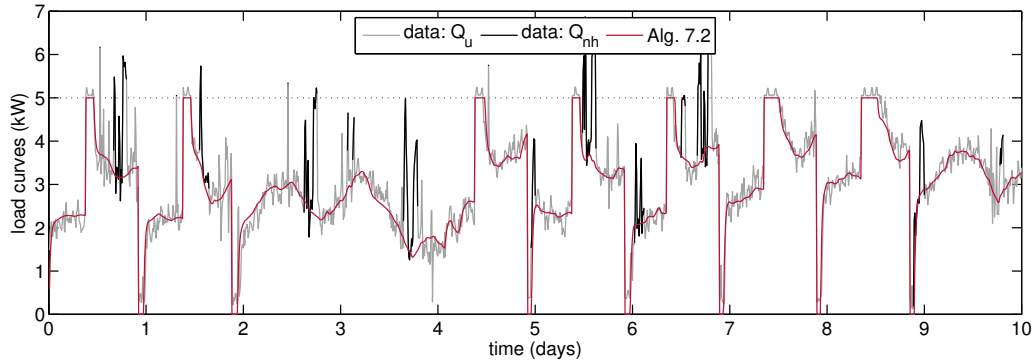

(a)

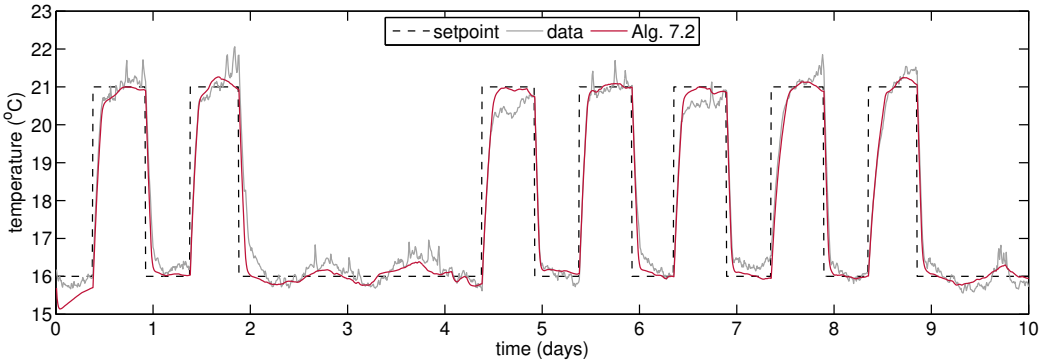

(b)

Figure 11: Exemple d'identification de bâtiment à partir de la courbe de charge totale. (a) En rouge, la courbe de charge utilie équivalente simulée par l'algorithme présenté au Chapitre 7, en gris la courbe de charge utile réelle et en noir les autres charges non-thermiques. (b) En rouge, estimation de la température intérieure, en nour gris la donnée observée. Données simulées.

instant. La répétition de plusieurs expériences avec données simulées confirme ces conclusions, et montre que la qualité d'estimation du PI impacte directement la qualité d'estimation du modèle final. Enfin, les données réelles permettent un test dans un cas limite, lorsque la courbe de charge contient une grande proportion d'apports internes gratuits et une faible part due à la commande : bien que dégradées, les estimations restent plausibles.

Ainsi, ce Chapitre 7, encore prospectif, propose d'anticiper le flux d'informations qui pourraient être disponibles prochainement grâce aux objets connectés du bâtiment, et de tirer profit de ces informations pour limiter l'instrumentation nécessaire à l'identification de bâtiment. Cette démarche est celle de l'ensemble de la thèse, et a été appliquée à chaque entrée et sortie du système, en donnant des résultats prometteurs. Parmi les perspectives à ce travail, présentées au Chapitre 8, la plus évidente consiste à étendre les résultats dans le cas où plusieurs entrées sont simultanément incertaines. Par exemple, en se plaçant dans la situation d'un bâtiment équipé d'une station météorologique connectée, avec deux capteurs de température, ainsi que d'un termostat connecté et d'un compteur communicant, il faudrait alors adapter les démarches des Chapitres 6 et 7 pour prendre en compte l'incertitude sur le flux solaire et les apports internes.

Note Certaines des contributions de cette thèse ont été publiées, dans des conférences avec actes, ou sous forme de brevet. Ces publications sont listées ci-dessous.

Liste des articles de conférence

[NMJ+17] Tahar Nabil, Éric Moulines, Jean-Marc Jicquel, Alexandre Girard, and Christian Lajaunie. Méthode de Monte Carlo à dynamique hamiltonienne pour estimation d'un modèle thermique de bâtiment. In *XXVIe Colloque GRETSI*, 2017.

[NMR+16] Tahar Nabil, Éric Moulines, Francois Roueff, Jean-Marc Jicquel, and Alexandre Girard. Maximum likelihood estimation of a low-order building model. In *24th European Signal Processing Conference (EUSIPCO)*, pages 702–707, 2016.

[NRJG17] Tahar Nabil, Francois Roueff, Jean-Marc Jicquel, and Alexandre Girard. Identification of a thermal building model by learning the dynamics of the solar flux. In *Machine Learning for Signal Processing (MLSP), 2017 IEEE 27th International Workshop on*, pages –. IEEE, 2017.

Liste des brevets

[NJGR16] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Système et procédé d'estimation du comportement thermique d'un bâtiment, pour un contrôle optimal de chauffage. Patent applied for, 2016. FR 1657902 / EP 17187160.

[NJGR17] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Procédé et dispositif de détermination indirecte d'un flux solaire incident. Patent applied for, 2017. FR 1758788.

[NJGR18] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Estimation d'un circuit électrique équivalent, à résistances et capacités thermiques, d'un bâtiment pour le contrôle optimal du chauffage du bâtiment. Patent applied for, 2018. FR 1850030.

Chapter 1

Introduction

1.1 The context of building identification

The building sector is currently at the heart of global energy policies as well as many research efforts. This is motivated by at least two reasons. The first one is the elevated environmental cost of this sector: it represented 37% of the final energy consumption in the European Union (EU) in 2004 [PLOP08], 44% in France in 2012 [ADE13]. Figure 1.1(a) shows that this proportion has always been elevated in France since 1970, and actually is at its highest in 2015. Moreover, these consumptions trigger also greenhouse gases emissions: 19.6% of the total EU emissions for the households only, according to Eurostats. A similar proportion is also observed in France, see Figure 1.1(b). Most of them (77% in France in 2009, [ADE13]) are actually due to space heating. As a consequence of this, and because the building renewal rate is small (for example, 1% per year in France), there is a need for building energy efficiency actions in order to foster the transition towards a more sustainable energy demand.

The second reason has to do with the raise of the new Smart Grids [FMXY12]. One of the core concepts of these grids is the notion of flexibility, with the challenge of unlocking and exploiting the flexibility of the demand in order to maintain a balance with an increasingly intermittent supply, due to the penetration of renewables. As such, the buildings are emerging as key interactive actors of the Smart Grids. Indeed, while it is still difficult to store the electrical energy, the elevated consumption of the building sector represents a significant potential of flexibility of the demand likely to keep the grid in balance, that actions like load shifting and demand response try to take advantage of. Load shifting aims for instance at reducing the daily electricity peak-demand in order to smoothen the demand. Whenever the peak is substantially high, the supply is ensured by thermal power plants, which contribute for most of the carbon dioxide emissions compared to other types of power stations. Hence, the Smart Grid technologies may also reduce the environmental cost of the building sector. In particular, advanced management and control techniques applied to the space heating system are at the center of such solutions, by exploiting the thermal storage capacities of the building.

Furthermore, the new digital ecosystem, with disruptions caused for instance by the smartphone, the cloud-based models and the Internet of Things (IoT) represents an opportunity to meet these challenges. As such, the software-based optimization of the energy consumption of the buildings appears as an efficient way of addressing these issues. For instance, field studies reported in [VNB+16] or [GAO+16] suggest that a smartphone-based real-time context-aware feedback could lower the room temperature, and thereby the energy consumption, without significantly affecting the thermal comfort. Another example is the spread of occupancy-responsive learning thermostats, which may

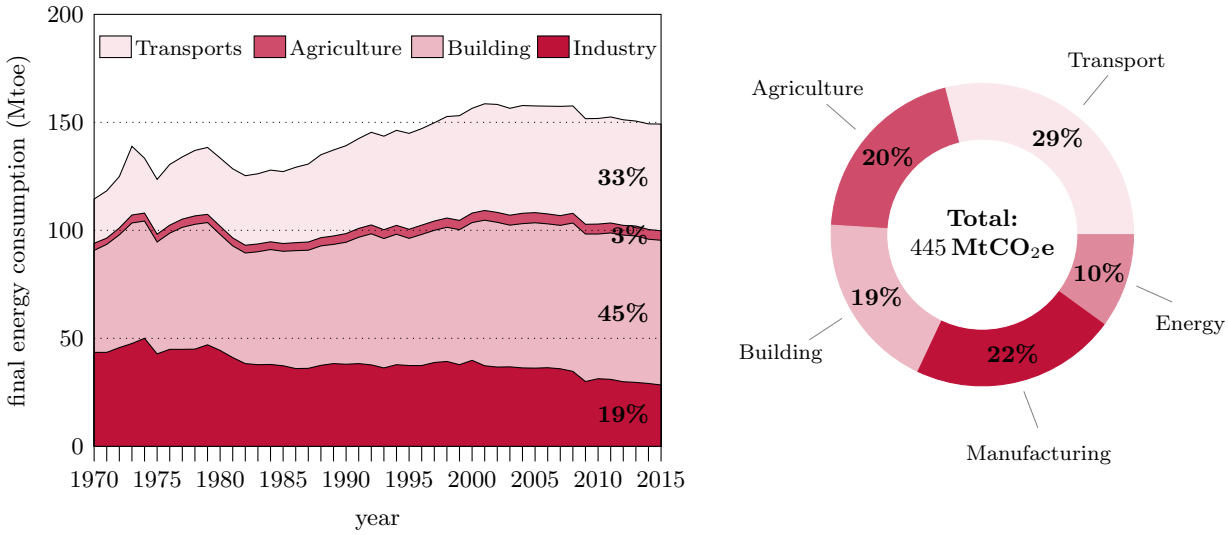

Figure 1.1: (left) Annual final energy consumption in tons of oil equivalent (toe) by sector in France, 1970 – 2015. Source: Service de l'observation et des statistiques, 2015. (right) Greenhouse gases emissions in tons equivalent carbon dioxide (CO_2e) by sector in metropolitan France in 2015. Source: Citepa/Format SECTEN - april 2017.

also save energy by automatically adjusting the setpoint temperature whenever the room under control is vacant, without altering the thermal comfort [PWM16]. As a result, the energy utilities are urged to create new energy services, information services or electrical devices, in order to tackle these challenges and face the formerly unexpected competition of information and communication technology actors such as the Web giants.

As pointed out in [SBPW16], this leads to a change of paradigm. Previous approaches in building energy dealt mainly with the conception phase of the building, with detailed thermal models simulated beforehand with some standardized occupation scenarios. On the contrary, the new paradigm implies that the user should be at the center of the services, by way of interactions: in exchange for the information he or she provides, for instance through measuring some quantities of interest, the user has access to a service that helps lowering or monitoring its energy consumption.

In conformity with this context and in order to address these challenges, advanced control techniques, such as model predictive control (MPC, see e.g. [HGP14]) applied to the heating, ventilation and air conditioning (HVAC) system are being investigated. Their impact on the building energy consumption is indeed potentially high: some experiences reported in [PŠFC11] and [ŠOCP11] achieved 17-24% and 15-28% energy savings, respectively. Similarly, actions such as load shifting can be accomplished through the dynamic control of the building energy, under some thermal comfort constraints [Mal12]. These actions depend first of all on a good **building model** [PCV+13], able to evaluate their potential and limitations, as well as the energy performances of the building. For instance, the load shifting in [Mal12] depends on a reliable model of the dynamics of the indoor temperature in order to make sure that the thermal comfort constraint is always respected by the optimal control strategy. The building model is thus a *prerequisite* for these services. Hence, this thesis is interested more particularly in the estimation procedure of such models.

By a building model, we mean more precisely a model representative of the thermal behaviour of the building and able to explain the dynamics of the indoor temperature. These dynamics are

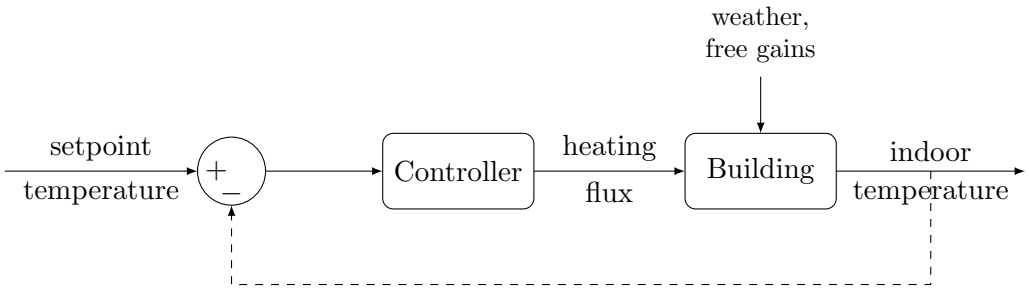

Figure 1.2: Block representation of a building system with regulated indoor temperature and other exogenous explaining factors (weather conditions including e.g. outdoor temperature and solar radiations, free internal gains including other electric appliances and occupancy profile).

usually explained by three classes of factors, namely the factors related to the intrinsic structure of the building (composition, geometry, etc.), the factors related to the weather conditions (outdoor temperature, solar radiation, wind speed, etc.) and the user's interactions with the building [Zay11, AAC+15]. These interactions are either controllable (setpoint temperature, heat flux) or non-controllable and seen as perturbations (metabolic heat gains due to the occupancy, free internal gains due to other appliances such as the computers, lighting, ventilation, etc. [PRMS08]). Considering more specifically the heating case, typically during winter periods, the building usually operates in closed loop, the indoor temperature being regulated such that is follows a certain setpoint signal. This is depicted in Figure 1.2.

Accordingly, several modelling approaches exist in order to explain the dynamics of the indoor temperature in terms of these factors. There are mainly two paradigms, with either the detailed physics-based or the statistically-based, data-driven, models. We opt in the sequel for the latter category. This choice is determined by the aforementioned context and the potential uses of the building model. Indeed, actions such as model predictive control require lightly parameterized models which can simulate the dynamics of the indoor temperature at a cheap computational cost [HGP12, PCV+13]. Hence, unlike the high-order detailed models, the data-driven approach is particularly relevant. The differences between these two types of models are further described in Chapter 2.

The construction of the data-driven model is thus based on *system identification* techniques, and therefore on experimentations: *"Input and output signals from the system are recorded and subjected to data analysis in order to infer a model"*, as defined by L. Ljung in his reference textbook devoted to system identification [Lju87]. The main steps in the system identification procedure are the choice of a class of models, the planification of a measurement campaign of the most significant explaining factors, the choice of a criterion for the model selection, the construction (estimation) of a model from the measured data, and the validation [Zay11]. An in-depth study of the identification procedure is provided in [Lju87]. In the next section, we emphasize the fact that the actual observation input and output signals are key to the building identification.

1.2 On the importance of the data

Given a certain class of building models, the identification is an algorithmic task that strongly depends on the availability of the data. Rather than focusing on the purely algorithmic aspect of the identification procedure, to which a significant part of the literature is already dedicated, this thesis

studies more particularly the practical constraints arising in the implementation of such algorithms. These constraints are due to the second side of system identification, namely data availability.

The issue of data availability and quality is often overlooked in the literature on building identification. In many publications, the data are either obtained from a specific accurate instrumentation in a laboratory (e.g. [BM11, Zay11, Mej11]) or by a mix of averaged historical weather records and simulated indoor climate (e.g. [MCPF10, HGP12]). Both approaches assume that every input and output of the thermal model is available with good enough precision: they are deterministic. We can illustrate how inadequate to an industrial context these two choices are. The latter approach is obviously not suited to practical cases because the simulated indoor climate is then necessarily obtained from a physical model, whereas such detailed models were precisely discarded in our methodology due to the cumbersome modelling effort that they represent: a tiresome task, not reproducible because specific to each building, that requires the intervention of an expert.

To motivate our point and show that the former appproach, based on a comprehensive instrumentation of the building, is not feasible in practice either, let us take the example of the solar flux. Having a measurement of the solar flux is necessary in order to have a good estimation of the time constants of the building, that characterize its inertia. This input represents moreover a substantial free heating source in low-energy buildings which tend to maximize this contribution. Yet, resorting to a specific instrumentation is costy, hence hardly reproducible at large scale for every building. Indeed, the solar flux is measured by a pyranometer, an expensive sensor (typically at least 100 €). As a result, the first approach, which assumes a perfect observation of the building is actually rather illusory in practice.

As a second example, we may remark that both approaches fail into taking into account the free heating gains due to the occupancy of the building which are highly stochastic by nature and difficult to measure with precision.

Consequently, instead of assuming that every input of the model is a deterministic signal, this thesis explores a third option and adopts the standpoint of an energy utility. From this standpoint, a unique data is available at first sight and thanks to the deployment of the smart meters, namely the overall energy consumption, called load curve. Note that there are differences according to the type of building:

- a house or a flat whose energy vector is purely electricity: then, the only measurement is the aggregated overall consumption;
- a house or a flat whose energy vector is electricity and gas: then, there is usually an aggregated measurement of the consumption of the heating system and hot water tanks (gas) and an aggregated measure for the other uses (electricity);
- for the tertiary sector (office buildings, industries, ...) the measurement is usually the aggregated overall consumption.

Yet, in any case, the collected data need to be disaggregated between heating and non-heating loads, the former being integrated in the thermal model whereas the latter are discarded. One way of dealing with this challenge would be to implement a Non-Intrusive Load Monitoring (NILM) algorithm (see [ZGIR12, EG09] for instance). However, such methods typically require a high-frequency measurement of the global load curve, in addition to the smart meter already installed. They come thus with a significant implementation cost.

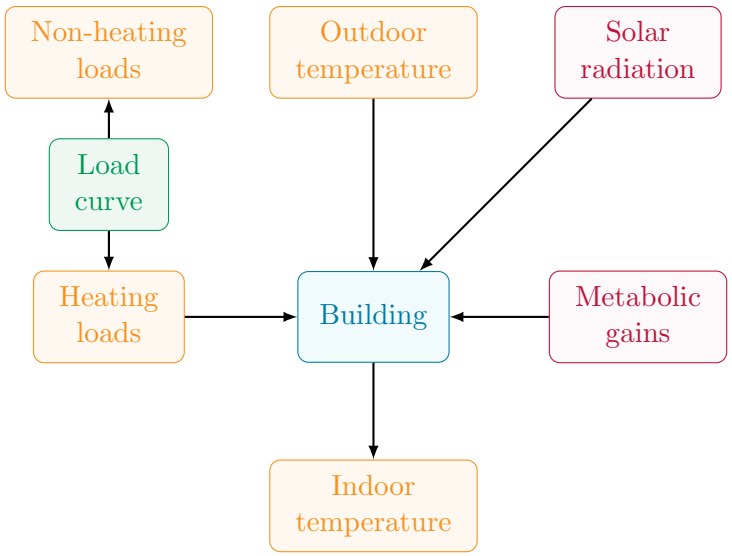

Figure 1.3: Availability of the main inputs and output of a thermal building model, from the standpoint of an energy utility. In red, quantities that are not measured, in orange, quantitites measured with some non-negligible uncertainty, in green, quantities that are accurately measured.

We opt here for a different method, that will be applicable to the other inputs: we suppose that the measured data is either *incomplete* or *inaccurate*, and we want to integrate the resulting uncertainty into the identification procedure. The other inputs, namely the temperatures and solar gains, are not available *a priori* to the energy utility (see Figure 1.3). Our primary hypothesis is then that the energy utility has access to the data supplied by the **connected objects** of the building, and that the utility will use these data for the building identification. This is further detailed in the next section. This is a change of paradigm, in the sense that the identification is no longer carried out in a context of perfect observation of the building, but rather in an incomplete, degraded context of observation: the data are either missing or collected with a certain degree of inaccuracy. This calls for a stochastic framework, rather than for the standard deterministic approach. As such, this thesis is in line with some previous contributions, such as [FABG14], which takes into consideration the uncertainty on the indoor temperature and the free internal gains, or [BHZ14] which integrates uncertain data in a Gaussian process modelling approach.

Although the connected objects help improving the observation of the building, we are nevertheless not in the perfect context of the two methods mentioned above. The idea is therefore to model the information the connected objects bring, and integrate them in the identification procedure. For instance, assume that we have a measurement of the aggregated consumption of a house on one hand, and a connected smart washing machine or hot water tank on the other hand. How can we use these smart appliances in order to discard their specific consumption from the aggregated consumption? Our aim is to answer this type of questions.

1.3 The Internet of Things, source of data and of uncertainty

The idea of resorting to the connected objects is justified by the fact that these objects are being increasingly ubiquitous in the society. In particular, if we consider a single house, the smart thermostat is starting to be wide-spread: 4.9M devices were sold in 2015 according to IoT Analytics (70% of them in North America), representing a 123% growth from Q4/2014 to Q4/2015. The

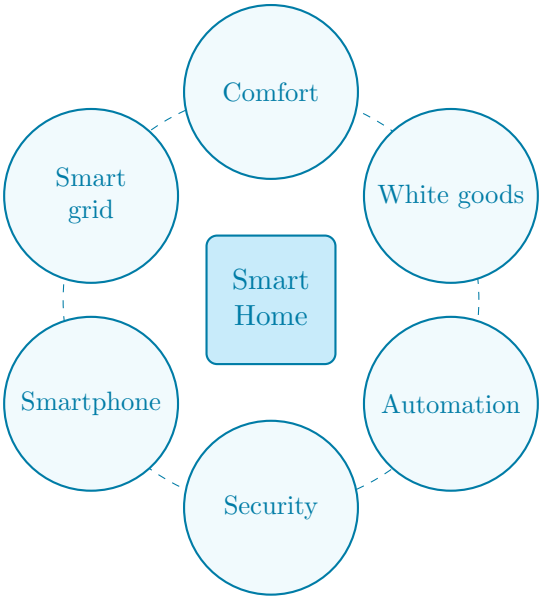

Figure 1.4: The smart home: a connected ecosystem.

market is still emerging, but the multiplicity of companies proposing smart thermostats or connected weather stations (in France, Sowee, Netatmo, Qivivo, Hector, Atlantic, Schneider Electrics, Ween,...) shows the potential of such products.

Regarding the specific case of individual housing, the connected objects form an ecosystem defining the so-called Smart Home. As represented in Figure 1.4, these objects can be grouped into six categories. The *smart grid* category encompasses the smart meters, which give access to the overall energy consumption of the building every few minutes. Dedicated meters might also measure specifically the local production of renewable energy or the consumption of an electric vehicle charging station. Then, the class of objects related to the *thermal comfort* covers the smart thermostats and the connected weather stations. These items usually monitor the indoor temperature and indoor air quality. Compared to the former programmable thermostat, the "smart" thermostat aims at reliably save energy through new features, such as smartphone user interfaces, energy-use feedback, occupancy sensing, fault detection, demand response or control of multiple zones [PWM16]. The *white goods*, i.e. the large domestic electrical goods such as a refrigerator or a washing machine, form the third category. These appliances are starting to get connected, with applications such as distance control or energy monitoring. We include also a potentially smart electric hot water tank (see an investigation in [Bee16]) in this category. The fourth category involves the items that enable the *home automation*: wireless door or window contacts, smart light bulbs, etc. The *security* appliances include indoor security cameras or even connected locks. Some connected cameras are for instance able to handle face recognition and send notifications accordingly. Finally, the most widely spread connected object is simply the *smartphone*, which may serve as an interface between the user and other connected appliances.

Following this brief description of the ecosystem of the connected objects, the key assumption of the thesis is that an energy utility may use the information brought by these objects in order to improve the observation of the thermal behaviour of the building and perform a model identification. Starting with the heating loads, the smart meter gives at least an upper bound of this contribution. Moreover, connected electrical appliances may help lower this bound. This is further detailed in Chapter 7. Should the utility decide to provide the client with a connected weather station, they

would then have access to *a* measurement - not *the* measurement - of the outdoor temperature, and similarly for the indoor temperature. Note that the connected weather stations do not measure the received solar radiation, which is therefore completely unobserved *a priori*. Finally, the devices in relation to the automation or the security may give information about the number of occupants in the building, and thereby about the metabolic heat gains. However, this seems very tentative at this time (see Chapter 7).

Hence, from an energy utility perspective, it seems reasonable to assume that only the total load curve is accurately measured, that the solar radiation and the metabolic gains are not measured at all, whereas the heating load curve as well as the outdoor and indoor temperatures are measured with some degree of uncertainty. This results in the simplified input-output represention of a building model in Figure 1.3. The constraints that arise in practice alter thus considerably the conditions of application of the standard building identification algorithms. Therefore, tackling these challenges raises a set of new questions for the building identification. What information do the connected objects bring in order to help mitigating the practical constraints? How to model these information mathematically? How can one adapt the identification procedure such that it accounts for some level of uncertainty of the inputs of the system? Can we perform an accurate identification from a reduced set of cheap connected sensors, despite their inherent uncertainty? What sensors are key to a good estimation? The purpose of the thesis is to address such questions. The next section describes the outline of the document and the main contributions to solving the issues raised in this introduction.

1.4 Outline of the document and contributions

As a preamble, the structure of the building model and the data that support the numerical experiments throughout the different chapters are presented in Chapter 2. Consistently with the chosen data-driven approach, this model is a low-order R3C2 equivalent electrical network, made of three resistors and two capacitors. From a mathematical point of view, the corresponding modelling framework is that of state-space models in closed loop. We emphasize the fact that our purpose in this thesis is not to find the best representation of a building, but instead to show how to adapt the identification procedure depending on some practical constraints. Hence, the algorithms suggested in the following chapters could easily be adapted to some refinements of the R3C2 model. Subsequently to this introductory chapter, the contributions of the thesis are two-fold.

First of all, Part I of the document assumes that every input and output of the R3C2 network is accurately measured, and looks into the algorithmic side of the identification process. Indeed, although many estimation algorithms are suggested in the literature, it is not clear whether the model should be identified in open loop, or within a control loop, the latter case corresponding to the real exploitation of buildings (see Figure 1.2). Hence, Chapter 3 gives some statistical and empirical arguments supporting the claim that an open-loop building model may be identified from on-site measurements and with a maximum likelihood estimator, even if the R3C2 system actually operates in closed loop. This result enables us not to model the feedback loop in general, which simplifies the problems studied in the second part of the thesis. We choose then a maximum likelihood identification procedure of the open-loop R3C2 network, with the Expectation-Maximization (EM) algorithm.

The second part deals more specifically with the issues raised in this introduction, related to the actual availability of the data in practice. Consequently, the different chapters of Part II are dedicated to the adaptation of the model and algorithm presented in Chapters 2 and 3 whenever each one of the inputs or output of the building model is successively assumed to be inaccurately

observed or not observed at all. To begin with, we assume in Chapter 4 that the indoor temperature measurements are not available at all, whereas the other data are available. This is a challenging framework, because the indoor temperature is the actual output of the building system. Since this quantity is not measured, we aim instead at identifying a closed-loop model whose observed output is the load curve. Relying in particular on the information contained in the setpoint temperature, we suggest a Bayesian estimation algorithm, based on the Hamiltonian Monte Carlo sampler. The numerical simulations show that the chosen algorithm achieves a good accuracy, which makes it possible for an utility to consider least intrusive identification methods.

Then, Chapter 5 introduces an estimation procedure that takes into account the uncertainty on the outdoor temperature. The underlying hypothesis is that the utility does not know the specific location of the temperature sensor, which is up to the end-user. Depending on its actual position, the sensor might thus give a misleading measurement. By appending a linear dynamic model of the outdoor temperature to the R3C2 network, we are able to reduce the possible bias part of the measurement provided by the sensor, and integrate thus the uncertainty inherent to connected weather stations. This hints moreover at the fact that the "true" outdoor temperature may be recovered from the measurements of an arbitrary temperature sensor and of the solar radiation.

Reversing the reasoning, we propose in Chapter 6 to replace the expensive pyranometer measuring the solar flux by a cheap temperature sensor purposefully exposed to the sun. Our assumption is thus that the temperature bias between this sensor and a second sheltered sensor is an indirect observation of the solar flux. The estimation procedure is adapted consequently, and includes in particular a nonlinear representation of the solar flux dynamics. The resulting model is a conditionally linear Gaussian state-space, whose identification is based on a Rao-Blackwellised version of the EM algorithm. Interestingly and despite this undirect observation of the solar flux, we are able to correctly learn both the dynamics of the solar flux and the physical parameters of the building. This confirms that ordinary connected objects may usefully replace the regular, cumbersome, instrumentation .

Thereafter, Chapter 7 addresses the case of the heating load curve. We use the information provided by the connected objects of the building in order to model the load curve. The main assumption is that the connected objects give access to a timeline of non-heating events, such that it is known whenever the measured aggregated load curve contributes either entirely or not to heating the building. We modify therefore the estimation of the R3C2 network such that it accounts precisely for the heating loads only. Besides, we design a statistical aggregation strategy for learning the parameters of the regulation as well as an iterative updating procedure of the R3C2 parameters. In the end, although still prospective, the suggested algorithm is shown to improve the naive identification that does not use the information of the connected objects.

Finally, we conclude and discuss the perspectives of this work in Chapter 8.

Note Some parts of this thesis have been presented in conferences or have been the subject of a patent. These publications are listed hereinbelow.

List of conference papers

[NMJ+17] Tahar Nabil, Éric Moulines, Jean-Marc Jicquel, Alexandre Girard, and Christian Lajaunie. Méthode de Monte Carlo à dynamique hamiltonienne pour estimation d'un modèle thermique de bâtiment. In *XXVIe Colloque GRETSI*, 2017.

[NMR+16] Tahar Nabil, Éric Moulines, Francois Roueff, Jean-Marc Jicquel, and Alexandre Girard. Maximum likelihood estimation of a low-order building model. In *24th European Signal Processing Conference (EUSIPCO)*, pages 702–707, 2016.

[NRJG17] Tahar Nabil, Francois Roueff, Jean-Marc Jicquel, and Alexandre Girard. Identification of a thermal building model by learning the dynamics of the solar flux. In *Machine Learning for Signal Processing (MLSP), 2017 IEEE 27th International Workshop on*, pages –. IEEE, 2017.

List of patents

[NJGR16] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Système et procédé d'estimation du comportement thermique d'un bâtiment, pour un contrôle optimal de chauffage. Patent applied for, 2016. FR 1657902 / EP 17187160.

[NJGR17] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Procédé et dispositif de détermination indirecte d'un flux solaire incident. Patent applied for, 2017. FR 1758788.

[NJGR18] Tahar Nabil, Jean-Marc Jicquel, Alexandre Girard, and Francois Roueff. Estimation d'un circuit électrique équivalent, à résistances et capacités thermiques, d'un bâtiment pour le contrôle optimal du chauffage du bâtiment. Patent applied for, 2018. FR 1850030.

Part I

Perfect observation of the building

Chapter 2

Modelling framework and available data

This chapter is devoted to the introduction of the open-loop building model that will be used in the subsequent chapters of this document. The chosen framework is that of the grey-box models based upon the analogy between electrical and thermal quantities. The building model is thus an electrical network, made only of resistors and capacitors. This network is usually denoted *RxCy*, where x and y are the number of resistors and capacitors, respectively. In the analogy, the temperatures (°C) are associated to the potentials of the network, whereas the heat fluxes (W) are represented by electric currents. Electrical resistances stand for heat transmission resistances, and electrical capacities for thermal capacities. These linear models with lumped parameters have been widely used in the literature, see for instance the review [AAC+15], T. Berthou's thesis [Ber13], or other articles such as [FVLA02, JMA08, LMDP14, NGRC15].

The first part of this chapter justifies thus the use of simplified *RxCy* models in the thesis, before describing more precisely the chosen model, putting emphasis on the physical meaning of its parameters. This chosen model is a R3C2 network, although all the algorithms derived in the subsequent chapters may easily be extended to other *RxCy* networks. Hence, we insist on the fact that we do not claim that the R3C2 network is the "best" building model: instead, Section 2.1.2 shows that it is a "good enough" model, suitable to our application. We show moreover how to translate mathematically this model into a linear discrete-time state-space representation.

Besides, defining a modelling framework entails defining a set of physical quantities that ought to be measured. Here, the indoor temperature is seen as the outcome of three factors, whose influence is mediated by the building characteristics, namely the outdoor temperature, the solar gains and the internal gains (heating, occupancy, etc). In order to validate the identification algorithms, it is necessary to have access to real or simulated datasets. We will use two kind of datasets: on one hand, a dataset corresponding to a physical experiment that has been conducted at EDF R&D Lab Les Renardières in order to provide on-site records of every input and output of the R3C2 model, and on the second hand, an hybrid dataset, with on-site measurements of the weather conditions, but numerical simulation of the indoor temperature and heating power. The second part of this chapter describes more specifically both the experiment and the data generating process for the simulated dataset.

2.1 A low-order builing model

2.1.1 The choice of a grey-box model

In this paragraph, we motivate the focus of this thesis on the class of models represented by electrical networks. In general, the building models are categorized as purely physics-based models (*white box* models), purely empirical models (inverse, *black box* models) and in-between models (*grey box* models). The advantages and drawbacks of each approach are well studied in the building identification literature. We review briefly these properties, and refer to [ZM12, LW14, Ber13, AAC+15] for a thorougher description.

The white box models are based on detailed physics equations. They require thus an accurate description of the building, in terms of geometry, materials, etc. This approach is developed in many softwares such as Energy-Plus or TRNSYS. Their main assets lie in the fact that they can capture precisely the building dynamics and that the model is physically interpretable. However, these models are of high order, which results in a poor computation time. Besides, they are heavily parameterized and their inputs are complex to obtain in practice, since they include detailed specific information about the building structure, such as the thermal properties of its materials.

The black box models on the other hand adopt an opposite standpoint, relying on the inverse modelling. Examples of such models are the artificial neural network [MZ04] or the ARX model [JMA08, MCPF10]. Contrary to the white box models, they are purely data-driven and require a set of historical on-site measurements over a certain period of time, but no *a priori* knowledge about the building. These empirical models are easier to build with fewer parameters and a reduced computational cost. Nevertheless, their physical interpretability is less obvious, and their performances depend heavily on the training dataset. In particular, the training period must be sufficiently long in order for the training data to cover the forecasting range.

It follows from this description that choosing a class of building models depends first of all on the targeted application of the model. Here, given the context developed in the Introduction, our interest is not on the building model *per se*, but rather on one hand on the thermal characteristics (e.g. well or poorly insulated) and on the other hand on the services and actions that can be designed from the identified model (e.g. model predictive control, load shifting, etc). As a consequence, the white box models are discarded, because of their high complexity and computational cost. Moreover, the black box models lack of physical interpretability, which is why we choose a third option, called *grey box* models, which are hybrid methods in between the white and black box models. Grey box models are semi-physical models, that use simplified physical descriptions to represent the thermal behaviour of the buildings. They are also learned from empirical input-output data. They offer thus a good trade-off between the previous two methods, and are consequently widely used for energy efficiency or optimal control purposes. The equivalent electrical RC network is an example of such semi-physical modelling combining computational efficiency and physical insights. Another advantage of the RC network that will be illustrated in Part II is the flexibility offered by the modelling framework of the chosen grey box model.

2.1.2 Model description

We consider in this thesis that a building is modelled as a R3C2 network, depicted in Figure 2.1. This model was initially presented by EDF in [MR91], and is used more recently in [Zay11].

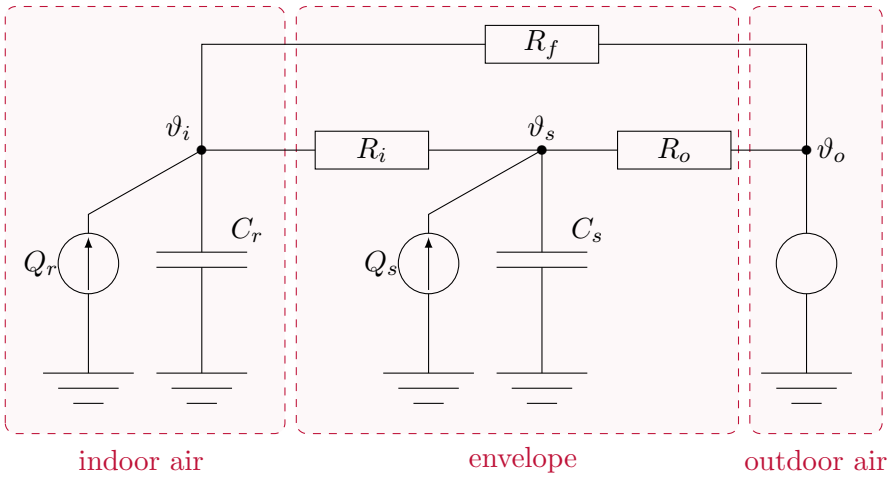

Figure 2.1: Building model: a single-zone equivalent R3C2 electrical network.

The first characteristic of the model is that it represents the indoor air of the building as a unique thermal zone. This is a standard approximation which implies that the model indoor temperature at node ϑ_i of the network should be representative enough of the indoor climate. A second feature is that the only climatic variables considered in the model are the outdoor temperature and the solar radiation. This is consistent with the findings of [Gof13], which shows on a case study that the two meteorological variables that explain most of the variance of the thermal model are the outdoor temperature and the direct solar radiation. Regarding the meaning of the parameters, [BM11] emphasizes the fact that, depending on the network, the resistances and capacities may have different physical interpretations. Here, the thermal resistances of the heavy materials forming the walls of the building are captured by R_i and R_o. A third resistance R_f represents the thermal losses through the lighter components of the envelope, as well as the losses due to ventilation [Zay11]. Moreover, the indoor air mass has capacity C_r while the envelope has its own capacity C_s, in order to account for the slow dynamics of the building. Hence, in general, one expects that the order relation $C_s > C_r$ holds. This envelope separates the indoor temperature ϑ_i from the outdoor temperature ϑ_o. It has also its own temperature ϑ_s, that is representative of the thermal state of the structure. Finally, the building is heated by two fluxes: on one hand, the internal gains Q_r (in W), including the heating system Q_h and the free internal gains due to the occupancy (metabolic Q_m as well as appliance-related Q_{fg} gains), and on the other hand, the solar gains Q_s through the windows. Usually, we have $Q_s = A_w \phi_s$, where ϕ_s is the solar flux, in W/m^2, and A_w is the effective window area, in m^2. The controllable input is the heat flux Q_h only. A summary of the inputs and the output of the building model is given in Table 2.1.

Pursuing the analogy with electrical quantities, we may apply Kirchhoff's laws to the R3C2 network, and derive the following linear continous-time differential equations:

$$Q_r = C_r \frac{d\vartheta_i}{dt} + \frac{1}{R_i}(\vartheta_i - \vartheta_s) + \frac{1}{R_f}(\vartheta_i - \vartheta_o), \tag{2.1a}$$

$$Q_s = C_s \frac{d\vartheta_s}{dt} + \frac{1}{R_i}(\vartheta_s - \vartheta_i) + \frac{1}{R_o}(\vartheta_s - \vartheta_o). \tag{2.1b}$$

Equations (2.1a)-(2.1b) are the dynamic equations of the system, representative of its thermal behaviour according to the R3C2 model.

Table 2.1: Inputs and output of the R3C2 open loop model.

Inputs	**Output**
solar gains heating flux internal gains outdoor temperature	indoor temperature

This model has the advantages of any grey-box model, that is low-order (2) and small number of parameters (5). The order 2 has been repeatedly reported to be the minimal sufficient order required to reproduce the dynamics in an accurate fashion [MCPF10, Zay11, HGP12]. For instance, [MCPF10] highlights the fact that two time scales, a fast one and a slower one, drive the dynamics of the indoor temperature. For computational purposes, the network is parameterized by θ formed by the inverse of the resistances and capacities:

$$\theta = \begin{pmatrix} z_r & z_s & z_f & z_o & z_i \end{pmatrix}^\intercal, \tag{2.2}$$

with $z_r := 1/C_r$, $z_s := 1/C_s$, $z_f := 1/R_f$, $z_o := 1/R_o$ and $z_i := 1/R_i$. We will assume in the rest of this thesis that the effective window area A_w is known. In addition, this model is also convenient because four physical parameters can be extracted from the representation. The first one, called the heat transfer coefficient, denoted UA and expressed in W/K, is representative of the static behaviour of the buidling. In the field of thermal science, UA quantifies the heat losses in a building. However, it is a tedious task to compute it for the thermician, since it requires the comprehensive knowledge of the geometry of the building, all its materials and their thermal properties, etc. On the other hand, one of the benefits of RC networks is that they provide a simple formula for UA, as the total conductance of the network, that is here

$$UA := \frac{1}{R_f} + \frac{1}{R_i + R_o} = z_f + \frac{z_o z_i}{z_o + z_i}. \tag{2.3}$$

Let us briefly explain the reasoning leading to this formula. Suppose that the system is in steady-state. Using the representation (2.1) at equilibrium, such that the derivatives are set to 0, we obtain the steady-state equation

$$Q_r = UA(\vartheta_i - \vartheta_o) - \frac{R_o}{R_o + R_i} Q_s, \tag{2.4}$$

where UA is defined as in (2.3). This indeed corresponds to the definition of UA as the static coefficient for the amount of losses through the envelope of the building per unit of the temperature difference between the indoor air and the outdoor air $\vartheta_i - \vartheta_o$. Hence, larger values of UA correspond to worse energy performances of the buildings. Likewise, the static equation (2.4) shows that, under the RC network model, the dimensionless ratio

$$g := \frac{R_o}{R_o + R_i} = \frac{z_i}{z_o + z_i} \tag{2.5}$$

defines the solar energy transmittance. By construction, we have $0 < g < 1$, with values closer to 1 characterizing those buildings that maximize the free heat gains received from the solar radiation. These efficient buildings are such that $R_i < R_o$, that is the heat transfers between the indoor air and the envelope are facilitated, compared to the transfers between the envelope and the outdoor air which shall be avoided.

Finally, two time constants respectively denoted τ_1 and τ_2 can be computed from the representation (2.1a) - (2.1b). They are also some physical parameters of interest, since they charaterize the inertia of the building, that is its dynamic behaviour. These four coefficients UA, g, τ_1 and τ_2 allow one to compare the R3C2 network with other models, whether they are other types of networks (see e.g. [BM11, BSSM14]) or higher-order physical models for instance.

2.1.3 Linear discrete-time state space representation

In order to derive the mathematical representation of the R3C2 network, we assume in this chapter that both the indoor temperature and the inputs $U := \begin{pmatrix}\vartheta_o & Q_r & Q_s\end{pmatrix}^\intercal$ are measured. The measured indoor temperature T_i is a noisy observation of the model temperature ϑ_i, which is not observed. Let $Y = T_i$ denote the observed output and $X = \begin{pmatrix}\vartheta_i & \vartheta_s\end{pmatrix}^\intercal$ the latent (hidden) state vector. (2.1) can be written as a linear continuous time state-space equation:

$$\frac{dX_t}{dt} = A(\theta)X_t + B(\theta)U_t + V_t, \tag{2.6}$$

$$Y_t = CX_t + W_t, \tag{2.7}$$

where V_t is a model noise and W_t an observation noise. V_t and W_t are independant white Gaussian processes of respective covariance matrices Q and R. The state matrices are

$$A(\theta) = \begin{bmatrix} -z_r(z_f + z_i) & z_r z_i \\ z_s z_i & -z_s(z_o + z_i) \end{bmatrix}, \quad B(\theta) = \begin{bmatrix} z_r z_f & z_r & 0 \\ z_s z_o & 0 & z_s \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \end{bmatrix}. \tag{2.8}$$

Applying the constant variation method and the theory of stochastic differential equations (see e.g. [Øks03]), it is standard to discretize the state equation (2.6) and obtain for a time step δ and at time $t = t\delta$:

$$X_t = A_\delta(\theta)X_{t-1} + B_\delta(\theta)U_t + V_t, \tag{2.9}$$

with $V_t \sim \mathcal{N}(0, Q_\delta)$, $Q_\delta = \int_0^\delta e^{sA} Q e^{sA^\intercal} ds$, $A_\delta = e^{\delta A}$, $B_\delta = A^{-1}(e^{\delta A} - I)B$. [VL78] gives an exact computation of Q_δ. Unless otherwise stated, the index t refers to a discrete time index in the rest of the thesis. Finally, the observation equation (2.7) is straightforward to discretize. In particular, the C matrix remains unchanged.

The R3C2 network is therefore written as a discrete time-invariant linear Gaussian state-space model (LGSSM), in open loop, with five parameters. For a given building, the five parameters in θ are unknown and estimated from historical on-site measurements of the inputs U and the output Y. The estimation procedure corresponding to the case where all the quantities in U and Y are correctly observed is postponed to Chapter 3, whereas alternative algorithms taking into account practical constraints are suggested in Part II. In the next section, we present the dataset that supports the evaluation of all these algorithms in the thesis.

2.2 Available data

2.2.1 BESTLab experiment

Presentation of the laboratory

An experiment was conducted during the thesis at BESTLab (Building Envelope & Solar Technologies) laboratory (EDF R&D, Les Renardières), from the end of 2015 to 2017. This laboratory

Figure 2.2: Location of the two cells used in our experiment, with western (CBL-Ouest) and southern (CBl-S1) orientations, respectively.

is a two-storey building (ground level and attic) made for testing envelope components and solar components integrated systems. Each level is therefore made of six cells, and each cell is made of one removable outside wall and five over-insulated walls. The removable wall is vertical for the cells at ground level, and tilted (30° inclination) for the attic. It was created in order to test for the thermal properties of innovative materials forming the envelope of a building. Consequently, the cells are identical, well insulated, and independent of the others, with heat losses occurring only through the removable outside wall, in contact with the outdoor air.

Two cells from the attic are used in the thesis, with a southern and a western orientation respectively. Their location inside the building is depicted in Figure 2.2, and several changes have been operated in order to adapt the cells to the context of the thesis. These changes enable to account for the influence of the solar radiation on the indoor climate, and increase the eventual heat losses of the cells:

1. a 1340 mm× 978 mm skylight with an electric shutter is installed on the outside wall of each cell. A small tablet controls the angle of the window opening and the percentage of occultation caused by the shutter;

2. an air duct is placed between the two cells, with a ventilation of controllable intensity.

In terms of instrumentation, the laboratory is equipped with a climate station, called ETNA and located a few meters away from the main building BESTLab. The main measured variables are summarized in Table 2.2. This station measures the direct beam (pyrgeometer, accuracy 5%), sky diffuse (pyranometer CM11, accuracy 0.5%) and total solar radiations (pyranometer CM21, accuracy 0.2%), as well as the outdoor temperature and other meteorological variables not used in the thesis (atmospheric pressure, hygrometry, wind speed and direction, rainfall), at a one-minute resolution. The outdoor temperature sensors are Pt100 1/3 Din sensors, with a 0.15°C accuracy. Each cell possesses also its own set of indoor and outdoor temperature sensors, as well as solar radiation sensors alongside the tilted removable wall surface (pyranometer CMP21, accuracy less than 2%). There is also a flow meter inside the air duct between the two cells. All these data are collected on a daily basis with a one-minute time-step.

Table 2.2: Measured variables at BESTLab and sensor characterics. The top section corresponds to the native instrumentation in the nearby climate station ETNA, the middle one to the native instrumentation in the cells, the bottom one to additional sensors installed in the cells.

Measured variable	Unit	Sensor	Range of device	Accuracy
Outdoor temperature	°C	Pt 100 1/3 Din	−100 to +200°C	±0.15°C
Global solar flux	W/m^2	pyranometer CM21	0 to 4000 W/m^2	0.2%
Outdoor temperature	°C	Pt 100 4-wire	−100 to +200°C	±0.1°C
Tilted solar flux	W/m^2	pyranometer CMP21	0 to 4000 W/m^2	< 2%
Indoor temperature	°C	Pt 100 4-wire	−100 to +200°C	±0.1°C
Outdoor temperature	°C	Thermokon SR65	−20 to +60°C	1%
Indoor temperature	°C	Thermokon SR04	0 to +40°C	±0.4°C
Heating consumption	Wh	Pulse counter	0 to 10^5 kWh	1 Wh

Command and control of the indoor temperature

However, the initial setup of the laboratory does not allow for the simulation of the real exploitation of a building by controlling the indoor temperature. In particular, it is not possible to force the indoor temperature to follow a user-defined setpoint signal. As a result, two 500W electric radiators are installed in each cell. For each cell, one radiator will have to maintain the setpoint temperature while the other will be used to model the free internal gains. Since these radiators do not come with any thermostat, we built our own, as described by the scheme in Figure 2.3(b). The key tool in this architecture is the multi-purpose controller and gateway from Schneider Electrics' SmartStruxure™Lite solution. This box has a web interface that enables one to manage the heating, ventilation and cooling system, the lighting and the other metering applications. Moreover, the manager is able to establish wired or wireless communications with different types of sensors and actuators, and includes the possibility for the user to control and monitor these apparatuses by writing its own code in programming language Lua.

Our architecture is therefore the following: on one hand, a wireless temperature sensor sends its measurement to the manager, while on the other hand the user specifies its prefered setpoint. Then, by comparing the setpoint to the measurement, at the level of the web interface, a program computes the required amount of power. This power is converted into a proportion of time during which the radiator is On at full power. We use two-minute intervals: if the control signal is 250W, then the radiator will be On during the first minute and off during the second. The manager sends thus an On/Off command to a Thermokon actuator located ahead of the radiator.

The measurement of the indoor temperature is provided by a set of 12 wireless and autonomous temperature sensors Thermokon SR04 in each cell (see Figure 2.3(a)). These sensors have an accuracy of ±0.4°C at 21°C. They communicate minutely using the *EnOcean* protocol, while their energy is provided by a photovoltaic cell. The sensors are distributed along three levels, and we use the average of the four sensors on the bottom level as feedback for the command.

Regarding the command, we choose a saturated PI controller with a simple anti-windup strategy [PR15]:

$$\begin{aligned} \mathrm{cmd}_{t+1} &= \mathrm{sat}(K_p\varepsilon_t + \Sigma_t), \\ \Sigma_{t+1} &= \Sigma_t + \left(K_i\varepsilon_t + K_s(\mathrm{cmd}_{t+1} - K_p\varepsilon_t - \Sigma_t)\right), \end{aligned} \tag{2.10}$$

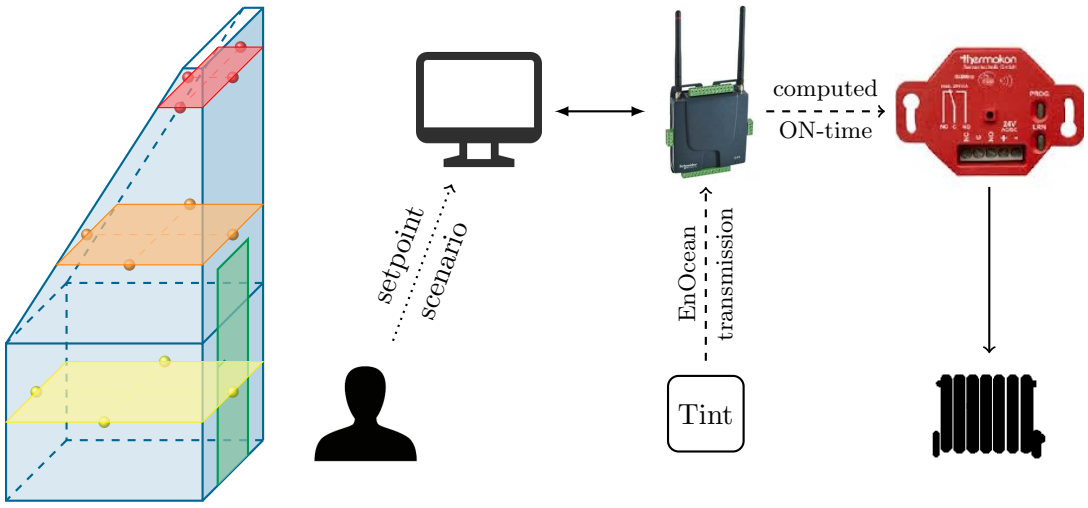

Figure 2.3: (a) Distributions of the 12 EnOcean temperature sensors (small spheres) inside each cell of the BESTLab experiment. The first level is at a height of 0.84m, the second at 2.44m, the third at 4.04m. (b) Infrastructure for the flexible command of the indoor temperature.

where cmd_t is the command (in W) at time t, ε_t is the difference between the setpoint and the measured indoor temperature at time t, Σ_t is the sum of these errors from 0 to t, sat is the saturation function between 0 and 500W, and K_p, K_i and K_s are the gains of the controller. The parameters K_p and K_i are tuned heuristically, using the Ziegler-Nichols method. K_s is chosen empirically such that $K_s K_p > K_i$, as recommended in [PR15].

The versatility of the manager enables one to monitor the status of the skylights (open vs closed), by mounting a wireless window sensor on the skylights. Moreover, two additional wireless temperature sensors are used to measure the outdoor air temperature. These sensors are not located at the climate station, but directly on the building, with a southern and northern orientation, respectively. Their specific feature is that they are exposed to the solar radiation, unlike the reference sensor from the laboratory instrumentation which is sheltered. The interest is that we may thus study the influence of the quality of the measurement of the outdoor temperature on the identification procedure (see Chapter 5).

Finally, prior to the BESTLab experiment, the free internal gains Q_{fg} of an office building located at EDF Lab Les Renardières were recorded during 27 days. They correspond to the aggregated consumption of all the computers in the offices. An example of BESTLab measurements of every input and output of the R3C2 model is given in Figure 2.4. The top plot in particular shows that the regulation strategy is efficient, since the indoor temperature follows closely the setpoint signal, without bias.

2.2.2 Simulated data

The BESTLab experiment provides a controlled environment as well as the meteorological data suitable for testing our algorithms. However, it will not be possible to compare the estimated models to some "ground truth" data. This is why we choose first to validate our algorithms on a set of (semi-)simulated data: the meteorological data (outdoor temperature, solar radiation) are provided by BESTLab, but the indoor temperature and the heating power are simulated.

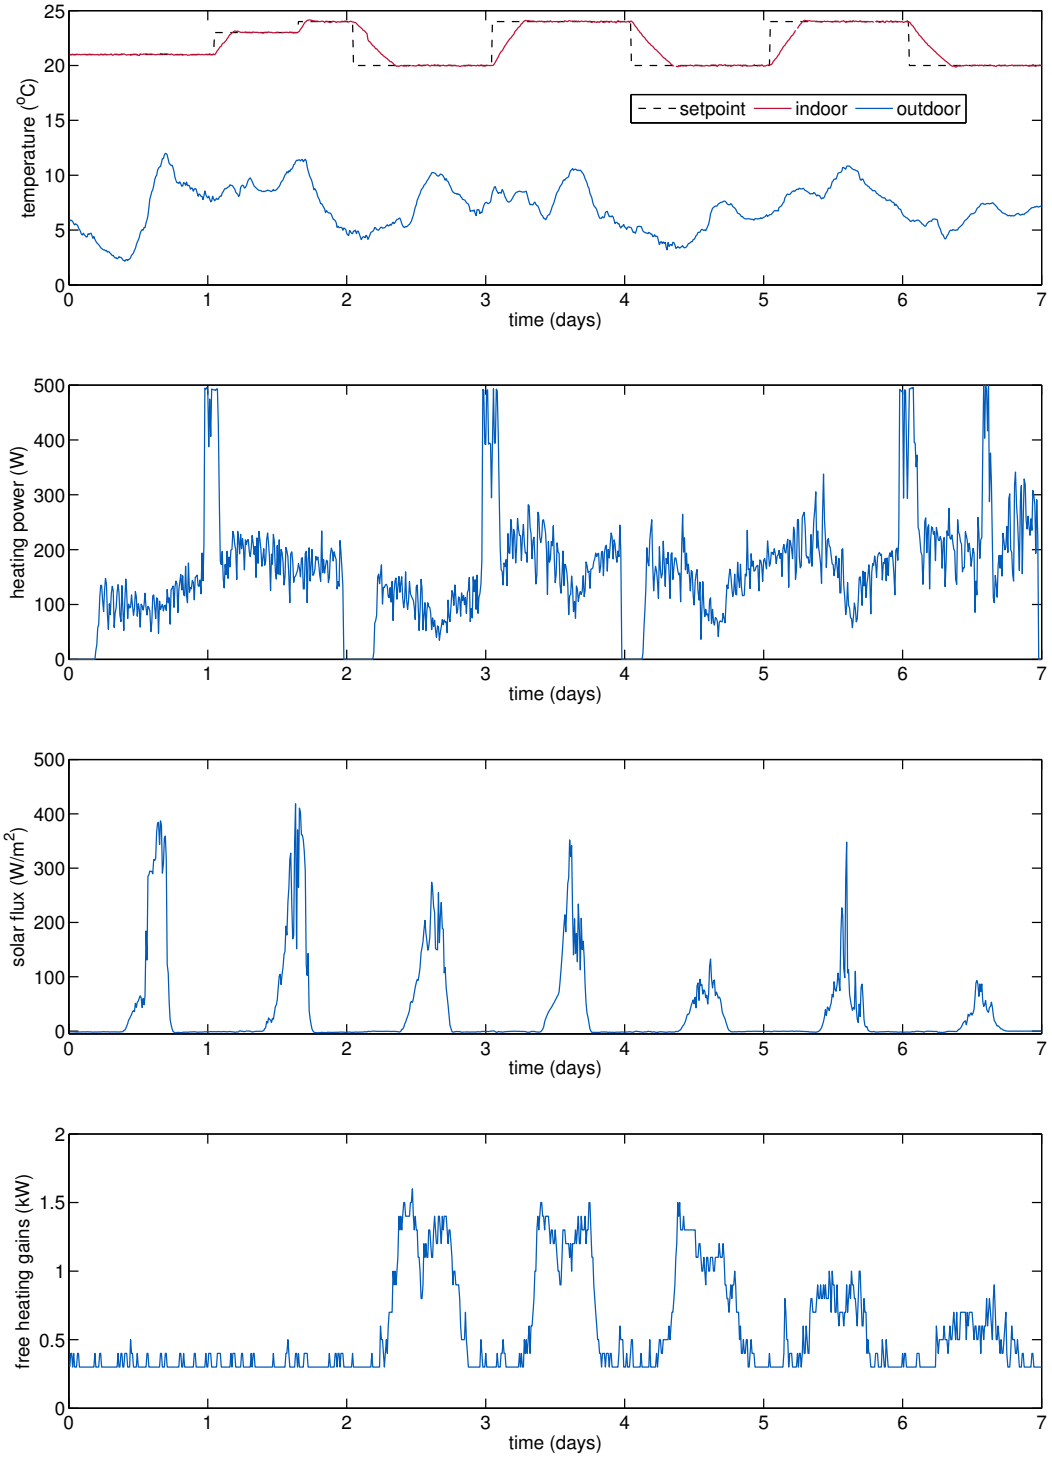

Figure 2.4: Example of one-week measurements at BESTLab. From top to bottom: indoor temperature T_i in the western cell, outdoor temperature ϑ_0, consumption Q_h of the radiator in the western cell, solar gains and free heating gains Q_{fg} from an office building located at Les Renardières. All the data, except Q_{fg}, are measured in December 2015, 23-29.

The generating process of the heating load and the internal temperature consists in selecting the true values of the five parameters of the R3C2 building model, adding a control loop and simulating a closed-loop model from the definition of a temperature setpoint sequence. This simulates the real exploitation of a building. Similarly to [Zay11], we use a Proportional-Integral controller coupled with a non-linear saturation function and an anti-windup strategy, such that the heating power is always positive and bounded. The setpoint has typically two levels (for instance cooling, 16°C, heating, 21°C, and matches the diurnal activity hours of the free internal gains Q_{fg}.

The semi-simulated dataset will be used for evaluating the performances of the algorithms in Chapters 3 to 7, whereas the real dataset from BESTLab will serve for evaluating the algorithms in Chapters 3, 6 and 7.

2.3 Summary

In this chapter, we have presented the building model that will be used in the rest of this thesis. Note that it will not be the purpose of the thesis to claim that the chosen model is the "best" representation of the thermal behaviour of a building. The R3C2 network was chosen and described in this chapter, mainly because of its low order and small number of parameters. This model, or any equivalent electrical network, takes the form of a discrete linear Gaussian state-space representation.

We have also presented the dataset that will be used for testing the algorithms in the rest of the thesis. Whether it be on simulated or real data from the BESTLab experiment, the data generating process is in closed loop, with the indoor temperature that follows a certain setpoint signal. This corresponds to the real exploitation of a building, in winter.

Chapter 3

Open-loop identification of a building observed in closed loop

In this thesis, we are mainly interested in the identification of a thermal building model under practical constraints. The first implicit constraint is that the building operates in closed loop. This means that the R3C2 open-loop representation of its thermal behaviour introduced in Chapter 2 is included in a feedback loop (see also Figure 1.2), the indoor temperature following a certain setpoint signal. It corresponds to the real exploitation of most buildings during the heating (or cooling) periods. Several types of estimation methods are suggested in the system identification literature, depending first of all on the available dataset and knowledge about the feedback [Lju87, FL99]. In particular, if the command (input) and the output of the open-loop system are measured, knowing the controller does not bring any additional information. The system should therefore be identified in open loop from input/output data, this is the so-called *direct approach*.

The direct approach is preferable in the context of building identification, because it ignores the feedback mechanism. This indeed is an advantage, because the controller would in fact introduce a non-linearity in the model (see Chapter 4): the heating power is saturated. Yet, in Section 3.1, we highlight the fact that this discussion is often circumvented in the building identification literature. Most contributions estimate either an open-loop model from open-loop generated data, which is unfeasible in practice, while the closed loop approaches face the problem of the nonlinearity.

In this chapter, we suggest in line with the system identification literature, that although the system operates in closed-loop, an open-loop identification algorithm is suitable, provided that the inputs and output are recorded. Hence, our purpose is to identify straightforwardly the open-loop R3C2 model, and study the performances of the estimator. Consequently, Section 3.2 is then devoted to the estimation procedure of the R3C2 network: this open-loop model will be identified directly from the data actually collected in closed loop by a maximum likelihood estimator (MLE). Some statistical arguments regarding the asymptotic properties of the MLE and numerical experiments supporting this approach are also given respectively in Section 3.3 and Section 3.4.

3.1 Literature review

Grey-box building models are widespread in the literature, their diversity lying mostly in the implicit physical phenomena that are taken into account. This is visible on one hand through the order and the number of parameters of the models, and on the other hand on the inputs and

Table 3.1: Comparison of different grey-box approaches towards building identification in terms of model class and data.

	Model order	Number of parameters	Time step (minutes)	Days	Data generation process	Number of variables
[BSSM14]	2	10	6	14	closed loop	7
[BM11]	1 to 5	6 to 23	5	6	open loop	4
[HGP12]	2+white box	6	1	60	closed loop	4
[MCPF10]	2	9 or 21	60	365	open loop	5
[Zay11]	2	7	10	14	closed loop	5

outputs of the model. In other words, the identification procedure is not only an algorithmic task, but depends also on the availabilty of the measurements.

In order to strengthen the link between the building model and the data at hand, some authors have for instance studied the influence of the number of days required for the identification and the discrete time step ([Ric91, BSSM14]). Here, we focus rather on the data generating process. In particular, the indoor temperature may be generated either in open loop, for a given heat load sequence, or within a feedback loop, with a controlled heat load sequence, which corresponds to the real exploitation of buildings. Other inputs, such as the weather conditions (outdoor temperature, solar radiation), are not controllable, and therefore merely seen as measurable perturbations.

From the perspective of optimal system identification, it is advisable to collect the data in open loop. In such case, it is indeed possible to choose and tune the heat load signal. Setting this as a Pseudo Random Binary Sequence (PRBS) enables to maximize the amount of information contained in the data. In [Ric91], it was shown that the heat load sequence should be such that the indoor temperature is highly unsteady. The system is therefore sufficiently excited at all frequencies, which means that it is identifiable [Zay11]. Regarding our application, this approach was for instance used in [BM11], where the authors propose a model selection procedure based on experiments in open loop for empty buildings. However, the drawback is that this method is intrusive and causes thermal incomfort whenever the building is inhabited.

More subtly, in [HGP12], the authors use a grey-box building model and include moreover a control loop for the heat load, based on the water inlet temperature of the radiator. Yet, they acknowledge that comfort conditions on inhabited buildings prevent the use of PRBS sequences for the setpoint. They resort then to use in parallel a detailed, white-box, building model in order to generate the input-output records corresponding to a pseudo random binary setpoint. This cosimulation approach is also used in [PCV+13]. Hence, the identification depends entirely on the white-box model, whereas the grey box model is kept for control purposes. The drawbacks are those of any white box model, as outlined in introduction: an indepth physical knowledge is needed and the computation time is significantly increased. We understand thus that using a very informative input signal, whether it is directly the heat flux or the temperature setpoint, is theoretically the best choice but meets with some practical limitations.

On the other hand, recent works have focused on closed loop approaches, without assuming any structure on the input signals. Closed loop implies to model the controller. The difficulty is essentially due to the fact that the controller is not linear: the heat flux is saturated, being upper bounded by the subscribed power - it is also necessarily positive, since only heating, and not cooling is considered in this thesis. This saturation has two consequences, namely the poorer informative content of the heat flux, and the algorithmic complexity induced by the overall nonlinearity. A first promising approach to tackle closed loop identification was based on a linear Proportional-Integral

(PI) controller [Zay11]. The global model is thus kept linear, and the estimation procedure, based on a Bayesian algorithm, gives fine results on simulated data. However, it lacks accuracy when taking into account the saturation of the controller. Another line of work, in T. Berthou's thesis [Ber13] (see also [BSSM14]), uses a simplified, ideal, regulation: the temperature reaches the setpoint as soon as the required power is available. This is a strong assumption, especially at small time steps, such as 6 minutes in [BSSM14]. Hence, closed-loop models are faced with algorithmic obstacles.

In Table 3.1, we summarize this quick literature review, and provide a few more examples. The table illustrates the heterogeneity of different grey box approaches for building identification, with an emphasis on the experimental conditions.

3.2 A maximum likelihood estimator

Given the above literature view, we suggest to adopt an open-loop identification strategy but without constraining the data generating process: the building system may operate in closed loop. The statistical rationale behind this choice is postponed to Section 3.3. In this section, we show first how to estimate the parameters R3C2 network described in Chapter 2, provided that every input and output is correctly measured. As previously shown in Section 2.1.3, the R3C2 network can be written as a discrete-time linear Gaussian state space model of the form

$$X_t = A(\theta)X_{t-1} + B(\theta)U_t + V_t, \tag{3.1}$$
$$Y_t = CX_t + W_t, \tag{3.2}$$

where θ is the parameter vector, $X := \begin{pmatrix}\vartheta_i & \vartheta_s\end{pmatrix}^\intercal$ the hidden state, $U := \begin{pmatrix}\vartheta_o & Q_r & Q_s\end{pmatrix}^\intercal$ the input vector, Y the noisy observation of ϑ_i and under the assumption that the initial state is normally distributed. Recall that ϑ_i (respectively ϑ_o) is hidden indoor (respectively structure) temperature of the building, whereas ϑ_o is the observed outdoor temperature. Q_r denotes the heating power, inluding free internal gains, Q_s the solar gains.

Standard statistical estimation methods already exist in the literature for such linear systems. We choose the Expectation-Maximization (EM) algorithm, spread by the seminal paper [DLR77], to estimate the parameters θ. In-depth studies of the EM algorithm are provided for instance. in [DMS14, Appendix D] or [MK08]. EM is a standard maximum likelihood approach for estimation in the special case of a linear Gaussian state-space models, as described in [SS82]. This means that the estimator $\widehat{\theta}$ of the parameter vector θ is such that

$$\widehat{\theta} = \arg\max_{\theta} L_\theta(y_{1:\mathrm{T}}), \tag{3.3}$$

where $L_\theta(y_{1:\mathrm{T}})$ is the log-likelihood of the T observations $y_1, \ldots, y_\mathrm{T}$:

$$L_\theta(y_{1:\mathrm{T}}) = \log p_\theta(y_{1:\mathrm{T}}) = \log\left(p_\theta(y_1)\prod_{t=2}^{\mathrm{T}} p_\theta(y_t|y_{1:t-1})\right). \tag{3.4}$$

Note that in the rest of this thesis, the generic notation $X_{k:l}$ stands for the set $\{X_i\}_{k=i}^{l}$. Besides, with a slight abuse of notation, p will be a generic notation for any probability density, in this chapter and the subsequent ones.

The strategy of the EM algorithm consists in iterating two steps, sketched in Algorithm 3.1. Firstly, the Expectation step calculates $\ln L_\theta(X, Y)$ the log-likelihood of the complete data problem,

Algorithm 3.1 EM for linear Gaussian state space models

0. Initialize the parameters to θ_0 and set $k = 0$.

1. *E-step:* compute

$$\mathcal{Q}(\theta, \theta_k) := \mathbb{E}_{\theta_k} \left[\ln L_\theta(X, Y) | Y = y_1, \dots, y_T \right]. \tag{3.5}$$

2. *M-step:* compute θ_{k+1}:

$$\theta_{k+1} = \arg\max_\theta \mathcal{Q}(\theta, \theta_k). \tag{3.6}$$

3. If convergence is reached, terminate, otherwise set $k \leftarrow k + 1$ and return to step 1.

adding the state to the observations. The states X being actually hidden, the true log-likelihood $\ln L_\theta(Y)$ is approached under the hypothesis θ_k by the quantity

$$\mathcal{Q}(\theta, \theta_k) := \mathbb{E}_{\theta_k} \left[\ln L_\theta(X, Y) | Y = y_1, \dots, y_T \right]. \tag{3.7}$$

This expectation is computed with a standard Kalman smoother (see e.g. [DMS14, Chapter 2]) applied to the model parameterized by θ_k. The detailed computations involved in solving the E-step are described in Algorithm 3.2. Equation (3.7) suggests then the iterative nature of the EM algorithm. Instead of maximizing directly the log-likelihood function (3.4), the M-step maximizes $\mathcal{Q}(\theta, \theta_k)$, which is an approximation of $L_\theta(x_{1:T}, y_{1:T})$.

A closed-form solution to the Maximization step can be derived when no constrain applies to the structure of the state matrices, see e.g. [DMS14]. Yet, the matrices A and B corresponding to the R3C2 network are structurally defined by $\theta = \begin{pmatrix} z_r & z_s & z_f & z_o & z_i \end{pmatrix}^\mathsf{T}$, as can be seen from (2.8). In order to guarantee that the structure of these matrices is kept invariant by the M-step, we need thus to use a general optimization solver such as Matlab's nonlinear programming solver `fminunc`. In practice, `fminunc` implements the trust-region algorithm, which requires the computation of $\nabla_\theta \mathcal{Q}(\theta, \theta_k)$. Moreover, we choose to express the parameters θ as $\theta = \exp(\gamma)$ and optimize with respect to γ. This is a simple way to make sure that estimated resistors and capacitors are positive while keeping the problem solved by `fminunc` unconstrained.

Finally, EM terminates either when a maximum number of iterations (here, 100) is reached or when a stopping criterion is satisfied. We decide that the normalized slope of the likelihood function being under a specified threshold is this criterion.

Evidently, in order for this strategy to be valid, there must be a "good" relationship between $\mathcal{Q}(\theta, \theta_k)$ and $L_\theta(y_{1:T})$. Hence, it can be shown that an increase in the $\mathcal{Q}$ quantity induces that the likelihood increases too: $\mathcal{Q}(\theta, \theta_k) \geq \mathcal{Q}(\theta_k, \theta_k) \Rightarrow L_\theta(y_{1:T}) \geq L_{\theta_k}(y_{1:T})$. See e.g. [DMS14, Proposition D.2] for a proof. Though it does not prove the convergence of the algorithm, it indicates that at least EM guarantees a non-decreasing likelihood at each step.

3.3 Asymptotic properties of the maximum likelihood estimator

In this section, we consider the class of maximum likelihood estimators (MLE) of θ. For the terminology, we are interested in linear hidden Markov models (HMMs), with additional inputs. See for instance the ARMAX model [DMS14, Chapter 2]. An extensive study of HMMs without

Algorithm 3.2 Resolution of the *E-step* for the R3C2 model

Inputs: series $\{Q_r(t),\, \vartheta_0(t),\, Q_s(t),\, T_i(t)\}_{t=1}^{\mathrm{T}}$, hypothesis θ', initialization $x_0^0 = \mu_0$, $P_0^0 = \Sigma_0$.

1. Kalman filtering equations, recursively forward in time for $t = 1, \ldots, \mathrm{T}$:

$$x_t^{t-1} = A(\theta')x_{t-1}^{t-1} + B(\theta')u_t, \tag{3.8}$$
$$P_t^{t-1} = A(\theta')P_{t-1}^{t-1}A(\theta')^\intercal + Q, \tag{3.9}$$
$$K_t = P_t^{t-1}C^\intercal(CP_t^{t-1}C^\intercal + R)^{-1}, \tag{3.10}$$
$$x_t^t = x_t^{t-1} + K_t(y_t - Cx_t^{t-1}), \tag{3.11}$$
$$P_t^t = (I - K_tC)P_t^{t-1}, \tag{3.12}$$

   where $x_t^s := \mathbb{E}_{\theta'}[x_t|y_{1:s}]$ and $P_t^s := \mathrm{cov}_{\theta'}(x_t|y_{1:s})$ for any $1 \leq s, t \leq \mathrm{T}$.

2. Kalman smoothing equations, recursively backward in time for $t = \mathrm{T}-1, \ldots, 0$:

$$J_{t-1} = P_{t-1}^{t-1}A(\theta')^\intercal(P_t^{t-1})^{-1}, \tag{3.13}$$
$$x_{t-1}^{\mathrm{T}} = x_{t-1}^{t-1} + J_{t-1}(x_t^{\mathrm{T}} - x_t^{t-1}), \tag{3.14}$$
$$P_{t-1}^{\mathrm{T}} = P_{t-1}^{t-1} + J_{t-1}(P_t^{\mathrm{T}} - P_t^{t-1})J_{t-1}^\intercal, \tag{3.15}$$

   initialized with the outputs $x_{\mathrm{T}}^{\mathrm{T}}$ and $P_{\mathrm{T}}^{\mathrm{T}}$ of the Kalman filter.

3. Lag-one covariance smoother:

$$P_{t-1,t-2}^{\mathrm{T}} = P_{t-1}^{t-1}J_{t-2}^\intercal + J_{t-1}(P_{t,t-1}^{\mathrm{T}} - A(\theta')P_{t-1}^{t-1})J_{t-2}^\intercal \tag{3.16}$$

   initialized with $P_{\mathrm{T},\mathrm{T}-1}^{\mathrm{T}} = (I - K_{\mathrm{T}}C)A(\theta')P_{\mathrm{T}-1}^{\mathrm{T}}$, where $P_{t_1,t_2}^s := \mathrm{cov}_{\theta'}(x_{t_1}, x_{t_2}|y_{1:s})$

4. With tr denoting the trace operator, compute

$$\begin{aligned}-2\mathcal{Q}(\theta,\theta') = &\ln|\Sigma_0| + \mathrm{tr}\left\{\Sigma_0^{-1}\left(P_0^{\mathrm{T}} + (x_0^{\mathrm{T}} - \mu_0)(x_0^{\mathrm{T}} - \mu_0)^\intercal\right)\right\} + \mathrm{T}\ln|Q| + \mathrm{tr}\left\{Q^{-1}\mathcal{S}(\theta,\theta')\right\} \\ &+ \mathrm{T}\ln|R| + \mathrm{tr}\left\{R^{-1}\sum_{t=1}^{\mathrm{T}}\left(CP_t^{\mathrm{T}}C^\intercal + (y_t - Cx_t^{\mathrm{T}})(y_t - Cx_t^{\mathrm{T}})^\intercal\right)\right\},\end{aligned} \tag{3.17}$$

   with $\mathcal{S}(\theta,\theta') := S_{00} - S_{01}A(\theta)^\intercal - A(\theta)S_{10} + A(\theta)S_{11}A(\theta)^\intercal - \Phi_{01}B(\theta)^\intercal - B(\theta)\Phi_{10} + B(\theta)\Phi_{11}B(\theta)^\intercal + A(\theta)\Psi_{11}B(\theta)^\intercal + B(\theta)\Psi_{11}^\intercal A(\theta)^\intercal$ and

$$\begin{aligned} S_{00} &:= \textstyle\sum_{t=1}^{\mathrm{T}}\left(x_t^{\mathrm{T}}(x_t^{\mathrm{T}})^\intercal + P_t^{\mathrm{T}}\right), & \Phi_{01} &:= \textstyle\sum_{t=1}^{\mathrm{T}} x_t^{\mathrm{T}}u_t^\intercal, \\ S_{01} &:= \textstyle\sum_{t=1}^{\mathrm{T}}\left(x_t^{\mathrm{T}}(x_{t-1}^{\mathrm{T}})^\intercal + P_{t,t-1}^{\mathrm{T}}\right), & \Phi_{10} &:= \textstyle\sum_{t=1}^{\mathrm{T}} u_t(x_t^{\mathrm{T}})^\intercal, \\ S_{10} &:= \textstyle\sum_{t=1}^{\mathrm{T}}\left(x_{t-1}^{\mathrm{T}}(x_t^{\mathrm{T}})^\intercal + P_{t-1,t}^{\mathrm{T}}\right), & \Phi_{11} &:= \textstyle\sum_{t=1}^{\mathrm{T}} u_tu_t^\intercal, \\ S_{11} &:= \textstyle\sum_{t=1}^{\mathrm{T}}\left(x_{t-1}^{\mathrm{T}}(x_{t-1}^{\mathrm{T}})^\intercal + P_{t-1}^{\mathrm{T}}\right), & \Psi_{11} &:= \textstyle\sum_{t=1}^{\mathrm{T}} x_{t-1}^{\mathrm{T}}u_t^\intercal. \end{aligned}$$

Outputs: E-quantity $\mathcal{Q}(\theta,\theta')$.

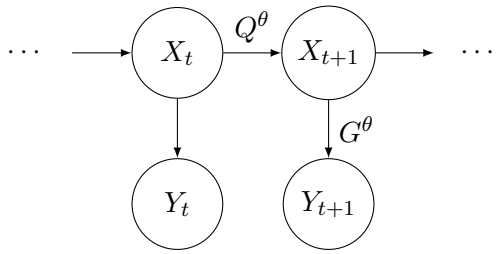

Figure 3.1: Graphical representation of the dependency structure of a hidden Markov model with observable process $\{Y_t\}$, hidden chain $\{X_t\}$, transition kernels Q^θ and G^θ parameterized by some vector θ.

additional input is provided in the reference textbook [CMR05]. Let us introduce some general notations and framework. The parameter space Θ is such that $\Theta \subseteq \mathbb{R}^d$, for some positive integer d. Apart from being a subspace of $\mathbb{R}^d$, no further assumption is made on the structure of Θ. The true parameter vector is denoted θ_0. For a model parameterized by $\theta \in \Theta$, l_n^θ denotes the log-likelihood function of θ given n observations of the latent system. Moreover, the symbol ∇_θ is the gradient operator (differentiation with respect to θ), and similarly ∇_θ^2 forms the Hessian. Hence, $\nabla_\theta l_n^\theta$ forms the so-called *score* function.

We let $\{X_k\}_{k=0}^\infty$ be a Markov chain on the measurable space $(\mathsf{X}, \mathcal{X})$. X is a separable state space, $\mathcal{X}$ the associated Borel σ–algebra. We denote by $\{Q_\theta(x, A), x \in \mathsf{X}, A \in \mathcal{X}\}$ the Markov transition kernel of the chain. $\{Y_k\}_{k=0}^\infty$ is a sequence of random variables taking values in the measurable space $(\mathsf{Y}, \mathcal{Y})$. Y is a subset of $\mathbb{R}^{n_y}$ and $\mathcal{Y}$ the associated Borel σ–algebra. Conditional on $\{X_k\}_{k=0}^\infty$, the Y_k are independent, with a conditional density $g_\theta(x_k, y_k)$ with respect to some dominating measure μ. This model is usually graphically depicted as in Figure 3.1, with arrows representing the dependencies between the different variables.

In this section, we are more particularly interested in the statistical properties of the maximum likelihood estimator (MLE) of the parameters θ of the model, namely the (strong) consistency (convergence to the true parameters as the sample size grows to infinity) and asymptotic normality (via a Central Limit Theorem, CLT). Proving these properties is an involved task that depends on both the model structure and the data generating process, and several contributions in the litterature already prove them in certain contexts. For instance, the asymptotic properties of the MLE of hidden Markov models without input process have been extensively studied, starting from the seminal papers [BP66, Pet69] (finite state and observation spaces) till more recent articles such as [DM01] (separable and compact state space), [GCL06] (state variable evolving in an open interval of the real line).

However, adding a random exogenous input to this general setup of HMMs complicates the establishment of the asymptotic properties of the MLE. Yet, some results exist in case of a linear model, which is the case of the R3C2 network. Hence, the next section mentions the open loop case, whereas the following one will address the closed loop case.

3.3.1 Data generated in open loop

The dependency structure in the open loop case is depicted in Figure 3.2. It is such that the data generating process is not known, and assumed to be independent of the state space at study. Whenever the model is linear, the statistical properties of the MLE have already been studied, for instance in [HD88, Cai87]. See also [DMS14, Chapter 2]. In particular, the consistency and

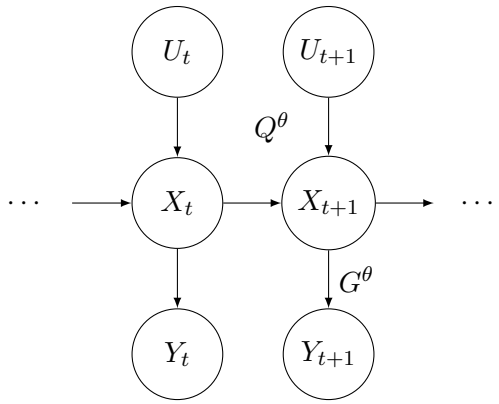

Figure 3.2: Graphical representation of the dependency structure of a hidden Markov model with observable process $\{Y_t\}$, hidden chain $\{X_t\}$ and exogenous input process $\{U_t\}$.

asymptotic normality of the MLE hold, under some general assumptions. One key assumption is the stability of the filter, which is ensured if the eigvenvalues of the state matrix A are inside the unit circle.

3.3.2 Data generated in closed loop

We are now interested in the statistical asymptotic properties of the MLE of state space models that are *commanded*. Accordingly, and unlike Section 3.3.1, the input process is not uncorrelated with the observable output of the system. On the contrary, the input process is generated from the past measurement(s) of the output. This new dependency structure of the variables is represented in Figures 1.2 and 3.3.

Model and hypotheses

The model at study here belongs to a parameterized family of state space models with commanded input and parameter space Θ. The state space model is not necessarily linear. The input process $\{U_t\}_0^\infty$, a real-valued random variable defined on the measurable space $(\mathsf{U},\mathcal{U})$, conditions the transition dynamics of the regime $\{X_t\}_0^\infty$, such that the transition kernel Q^θ is actually defined on $(\mathsf{U}\times\mathsf{X})\times\mathcal{X}$ rather than $(\mathsf{X}\times\mathcal{X})$. However, $\{U\}$ does not affect the observation equation, and the conditional density with respect to some dominating measure μ of Y_t given the state X_t is $g^\theta(X_t,Y_t)$. Several additional hypotheses are made:

(CL1) The model is fully dominated: there exists a probability measure λ on $(\mathsf{X},\mathcal{X})$ that dominates $Q^\theta(u,x,\cdot)$ for all $(u,x)\in\mathsf{U}\times\mathsf{X}$. The corresponding transition density is denoted $q^\theta(u,x,x')$.

(CL2) The parameters of the model are supposed to be *separated* between $\theta\in\Theta$ for the system and $\eta\in\Xi$ for the feedback loop (see Figure 3.3). That is the control is parameterized independently from the state space model.

(CL3) The controlled input U_t depends upon the past states and measurements of the output only through Y_{t-1}. The (possibly time-heterogeneous) conditional density with respect to some dominating measure ν of U_t given Y_{t-1} is $h_t^\eta(Y_{t-1},U_t)$.

In (CL3), we assume that the input U_t is generated only from the last measurement of the output Y_{t-1}. This is because it simplifies the notation, neverthess the discussion that follows in this

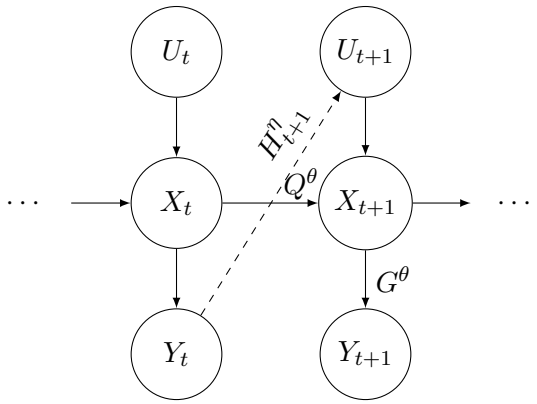

Figure 3.3: Graphical representation of the dependency structure of a hidden Markov model with observable process $\{Y_t\}$, hidden chain $\{X_t\}$ and input process $\{U_t\}$ generated from the past measurement of $\{Y_t\}$.

section would apply all the same if U_t depended on $Y_{s:t-1}$, $s < t$. Similarly, we could have added a dependency on past inputs $U_{1:t-1}$. Note also that if the controlled input does depend upon the past observable process $\{Y\}$, it does not depend upon the past hidden chain $\{X\}$: on Figure 3.3, the feedback arrow is from Y_t to U_{t+1}, and not from X_t to U_{t+1}. Indeed, should we have for instance $U_t = f(X_{t-1})$, then the form and the asymptotic properties of the maximum likelihood estimator would be different.

Some properties

Our interest is in the maximum likelihood estimation of the parameters θ of the system with the data generated in closed loop, as described in the previous section. In the considered context, the parameters η of the feedback loop need not to be estimated due to the chosen direct approach for the identification. Proposition 3.3.1 justifies the use of the direct approach when the data are generated in closed loop.

Proposition 3.3.1. *Consider a state space model with observed inputs generated in closed loop. Assume that conditions (CL1)-(CL2)-(CL3) hold true. Let θ_0 be the true parameter value of the system and $\hat{\theta}_n$ be the maximum likelihood estimator from the n data $y_1, \ldots, y_n$. Then whether or not the data $y_1, \ldots, y_n$ have been generated in closed loop, $\hat{\theta}_n$ can be computed from the likelihood function obtained from the same data, as if they were generated in open loop.*

Proof. For $t = 1, \ldots, n$, let $z_t := \{u_t, x_t, y_t\}$ and $z_0 = \{x_0, y_0\}$. For $t = 0, \ldots, n-1$, the dependencies of the model imply that

$$p^{\theta,\eta}(z_{t+1}|z_{0:t}) = h^\eta_{t+1}(y_t, u_{t+1}) q^\theta(u_{t+1}, x_t, x_{t+1}) g^\theta(x_{t+1}, y_{t+1}). \tag{3.18}$$

Moreover, it is straightforward to write

$$p^{\theta,\eta}(z_{1:n}|z_0) = \prod_{t=0}^{n-1} p^{\theta,\eta}(z_{t+1}|z_{0:t}), \tag{3.19}$$

and by marginalizing out $x_{1:n}$

$$p^{\theta,\eta}(u_{1:n}, y_{1:n}|x_0, y_0) = \int p^{\theta,\eta}(z_{1:n}|z_0) \lambda(dx_{1:n}). \tag{3.20}$$

Hence, combining these three expressions, we get a formula for the joint log-likelihood:

$$p^{\theta,\eta}(u_{1:n}, y_{1:n}|x_0, y_0) = \int \prod_{t=0}^{n-1} \left\{ h_{t+1}^{\eta}(y_t, u_{t+1}) q^{\theta}(u_{t+1}, x_t, x_{t+1}) g^{\theta}(x_{t+1}, y_{t+1}) \lambda(dx_t) \right\} \tag{3.21}$$

$$= \left\{ \prod_{t=0}^{n-1} h_{t+1}^{\eta}(y_t, u_{t+1}) \right\} \times p_o^{\theta}(y_{1:n}, u_{1:n}|x_0), \tag{3.22}$$

where $p_o^{\theta}(y_{1:n}, u_{1:n}|x_0) = \int \prod_{t=0}^{n-1} q^{\theta}(u_{t+1}, x_t, x_{t+1}) g^{\theta}(x_{t+1}, y_{t+1}) \lambda(dx_t)$ is the density that would be obtained if the data were actually generated in open loop. Therefore, the joint log-likelihood is

$$\ln p^{\theta,\eta}(u_{1:n}, y_{1:n}|x_0, y_0) = \sum_{t=0}^{n-1} \ln h_{t+1}^{\eta}(y_t, u_{t+1}) + l_{o,n}^{\theta}(y_{1:n}, u_{1:n}|x_0) \tag{3.23}$$

$$= C + l_{o,n}^{\theta}(y_{1:n}, u_{1:n}|x_0). \tag{3.24}$$

Here, the term C does not depend on θ and $l_{o,n}^{\theta}$ denotes thus the log-likelihood for the open loop model only. (3.24) shows that we do not need to deal with the control loop model when computing the MLE: maximizing in θ the log-likelihood for data generated in closed-loop is the same as maximizing in θ the log-likelihood $\ln L_o^{\theta}(y_{1:n}, u_{1:n}, x_0)$ of the open loop state space model. □

The point of the proof of Proposition 3.3.1 is to show how under assumptions (CL2)-(CL3), the densities h_t^{η} for the control function are excluded from the objective function. Notice the importance of assumption (CL3): if the controlled input depends on past state variables, then (3.21) is intractable, since in particular the densities h_t^{η} remain under the integral sign. Nevertheless, proving the asymptotic properties of the MLE in this context remains beyond the scope of the thesis. We will therefore posit that these properties hold true.

(CL4) the true parameter θ_0 is an interior point of Θ, and its MLE $\hat{\theta}_n$ verifies:

i) $\hat{\theta}_n$ is strongly consistent;

ii) $\hat{\theta}_n$ is asymptotically normal with covariance matrix the inverse of the Fisher information $\mathcal{I}(\theta)$ defined by

$$\mathcal{I}(\theta) = \lim_{n\to\infty} \frac{1}{n} \mathbb{E}_{\theta} \left[-\nabla_{\theta}^2 \ln L_{\theta}(y_{1:n}) \right]; \tag{3.25}$$

iii) there is a law of large numbers for the observed information matrix.

These properties hold whether the data are generated in open or closed loop.

Although (CL4) could appear as a strong assumption, it seems appropriate, because the statistical asymptotic properties of the MLE have been demonstrated for many different models, as highlighted by this document: input-ouput system in [FL99], general hidden Markov models (without input process), linear state space models with exogenous inputs,... Besides, in order to justify furthermore this hypothesis, we may show that under assumptions (CL1) to (CL4), the empirical approximations of the score function and of the Fisher information matrix $\mathcal{I}(\theta)$ is the same regardless of the data generating process (open or closed loop). The computations are similar to the proof of Proposition 3.3.1.

Consider thus a state-space model with observed inputs generated in closed loop, and assume that conditions (CL1)-(CL2)-(CL3)-(CL4) hold true. We have:

$$p_n^{\theta,\eta}(y_{1:n}|x_0,y_0) = \frac{p_n^{\theta,\eta}(y_{1:n},u_{1:n}|x_0,y_0)}{p_n^{\theta,\eta}(u_{1:n}|x_0,y_{0:n})} \tag{3.26}$$

$$= \prod_{t=0}^{n-1}\left\{\frac{p_n^{\theta,\eta}(y_{t+1},u_{t+1}|x_0,y_{0:t},u_{1:t})}{h_{t+1}^{\eta}(y_t,u_{t+1})}\right\}. \tag{3.27}$$

Hence the marginal-log likelihood is

$$l_n^{\theta} = \sum_{t=0}^{n-1} \ln \frac{p_n^{\theta,\eta}(y_{t+1},u_{t+1}|x_0,y_{0:t},u_{1:t})}{h_{t+1}^{\eta}(y_t,u_{t+1})}. \tag{3.28}$$

Its gradient is therefore

$$\nabla_\theta l_n^{\theta} = \sum_{t=0}^{n-1} \nabla_\theta \ln p_n^{\theta,\eta}(y_{t+1},u_{t+1}|x_0,y_{0:t},u_{1:t}), \tag{3.29}$$

which does not depend on the control loop. Similarly for the Hessian matrix, we obtain $\nabla_\theta^2 l_n^{\theta} = \nabla_\theta^2 l_{o,n}^{\theta}$. We conclude that, under conditions (CL1) to (CL4),

1. $\hat{\theta}_n$ can be computed as

$$\hat{\theta}_n = \arg\max l_{o,n}^{\theta}; \tag{3.30}$$

2. the central limit theorem reads

$$\sqrt{n}(\hat{\theta}_n - \theta_0) \xrightarrow{d} \mathcal{N}(0, \mathcal{I}(\theta_0)^{-1}); \tag{3.31}$$

3. and moreover

$$\frac{1}{n}\nabla_\theta^2 l_{o,n}^{\theta} \xrightarrow{P} -\mathcal{I}(\theta_0). \tag{3.32}$$

In other words, although the data is generated in closed loop, the system parameter can be identified as if in open loop, and its asymptotic covariance estimated using the Hessian of the open-loop likelihood. Observe, however, that this asymptotic covariance $\mathcal{I}(\theta)^{-1}$ is different from the open-loop case.

3.4 Numerical experiments

3.4.1 Case 1: simulated dataset

The numerical experiments are conducted on several datasets, with length varying from 7 to 14 days. The data are generated *in closed loop*, according to the procedure described in Chapter 2. In the sequel, the base scenario is when the data are generated in closed loop, with a cooling setpoint of 16 °C. $M = 50$ initializations are performed, uniformly in $[0, 1]$ for every parameter (all true values are smaller than 1), except for z_s which is drawn between 0 and z_r. The EM algorithm is iterated until convergence, with a maximum number of iterations set to 100. Besides, we check at iteration that the Kalman filter is stable, and reinitialize the procedure if one eigenvalue of the state matrix $A(\theta_k)$ is outside the unit circle. Similarly, if $A(\theta_k)$ appears to be ill-conditioned, the procedure is also reinitiliazed.

Table 3.2: Estimated parameters of the R3C2 network and twice their standard errors (2SE), from the best estimator among $M = 50$ initializations. The data are generated in closed loop, for 14 days at time step 10 minutes.

Parameter	True value	Estimation (± 2SE)	Error (%)
z_r	0.1	0.090 (±0.001)	10
z_s	0.01	0.0070 (±0.0023)	30
z_f	0.2	0.207 (±0.006)	3.5
z_o	0.4	0.430 (±0.080)	7.5
z_i	0.25	0.221 (±0.012)	11.6
UA	0.354	0.352 (±0.009)	0.6
g	0,385	0.346 (±0.040)	10.1
τ_1	33.7	49.9 (±6.0)	48.1
τ_2	3.6	4.2 (±0.2)	16.7

An example of point and standard error estimates is given in Table 3.2. Given the relative small amount of data (14 days), we see that the EM algorithm can perform well on estimating all four physical parameters. UA and g are accurately estimated, whereas the time constants are harder to retrieve, but their order of magnitude is good enough from a thermicist point of view. Besides, the closed-loop simulation of the indoor temperature and heating power on a test dataset and with these estimated parameters is shown in Figure 3.4. It emphasizes the good agreement between the model and the observations. Furthermore, Table 3.3 shows that compared to a Monte-Carlo procedure on 100 experiments, using the CLT gives similar but narrower intervals. It seems therefore legitimate to make the approximation of using the CLT even for closed-loop generated data.

In order to further illustrate the performances of our approach, we also implemented a standard algorithm of the literature. A widespread approach consists for instance in identifying an open-loop Multi-Input Single-Output (MISO) discrete transfer function by means of the least squares algorithm, as in [HGP12]. The transfer function takes the form

$$H(z^{-1}) = \begin{bmatrix} \frac{b_{11}z^{-1}+b_{12}z^{-2}}{1+a_1z^{-1}+a_2z^{-2}} \\ \frac{b_{21}z^{-1}+b_{22}z^{-2}}{1+a_1z^{-1}+a_2z^{-2}} \\ \frac{b_{31}z^{-1}+b_{32}z^{-2}}{1+a_1z^{-1}+a_2z^{-2}} \end{bmatrix}. \tag{3.33}$$

In such case, the physical parameter UA and the time constants can be defined from the coefficients $a_1, \dots, b_{32}$. Indeed, the two time constants of the system are classically obtained from the poles of the transfer function, whereas the parameter UA is the inverse of the static gain with respect to the input Q_r. Confidence intervals can also be constructed for the identified parameters of the transfer function (see e.g. [Lju87]) The $\delta-$method is used to derive the confidence intervals for UA, τ_1 and τ_2 from those of $a_1, \dots, b_{32}$. The different approaches shall therefore be compared through the accuracy of the estimated physical parameters.

We performed the identification with this approach on the same data set generated in closed loop. The point estimates and twice their estimated standard errors are given in Table 3.4, together with the estimation errors. It turns out that this simpler approach gives good results regarding the accuracy of the static parameter UA. However, it fails to reproduce the dynamic behaviour of the system, the two time constants being largely underestimated.

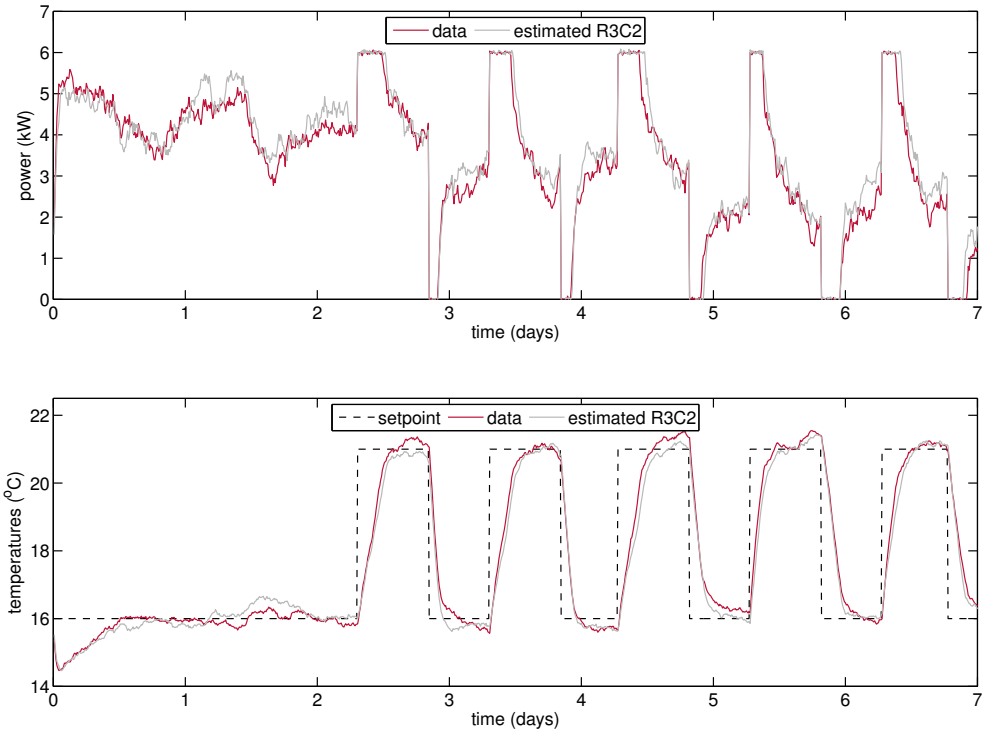

Figure 3.4: Closed loop simulation of the heating power (top) and the indoor temperature (bottom) with the true (red) and estimated (grey) parameters, with true regulation parameters.

Table 3.3: Estimated standard errors of the parameters, compared to a Monte-Carlo procedure on $N = 100$ data sets. The data are generated in closed loop, for 14 days at time step 10 minutes.

Parameter	CLT	Monte-Carlo
z_r	0.0004	0.0013
z_s	0.0012	0.0015
z_f	0.0030	0.0073
z_o	0.040	0.0871
z_i	0.0058	0.0084
UA	0.0046	0.0059
g	0.0201	0.0496
τ_1	2.98	19.21
τ_2	0.1106	0.127

Table 3.4: Estimated parameters (± 2 SE) and their corresponding absolute errors for the EM and the least squares algorithms. The data are generated in closed loop, for 14 days at time step 10 minutes.

Parameter	EM	Transfer function
UA	0.352±0.009 (0.6%)	0.318±0.005 (10%)
τ_1	49.9±6.0 (48.1%)	6.2±13.0 (82%)
τ_2	4.2±0.2 (16.7%)	0.18±0.04 (95%)

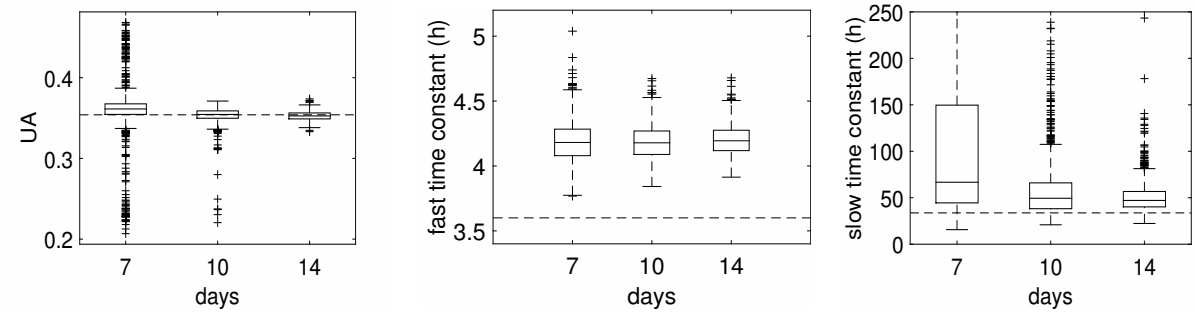

Figure 3.5: Boxplots of the $N = 1000$ estimations of UA (left), the fast time constant (center) and the slow time constant (right), against the number of days in the learning data set at sample rate 10 minutes. The true values are the horizontal dashed lines.

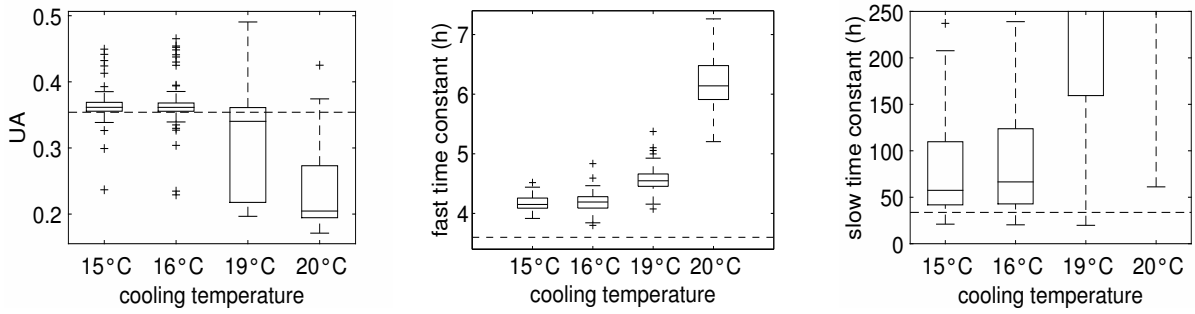

Figure 3.6: Boxplots of the $N = 100$ estimations of UA (left), the fast time constant (centerr) and the slow time constant (rights), against the cooling setpoint. The heating setpoint is 21 °C, and the data length is 7 days. The true values are the horizontal dashed lines.

Finally, some simulations are conducted in order to quantify the influence of the inputs on the quality of the estimation procedure. Indeed, it is well-known that a necessary condition for the identification to be efficient is that the inputs of the system comprise as much information as possible. In this view, we explored the influence of the data length and of the gap between the cooling and heating setpoint. The intuition behind the latter point is that under nonlinear control, the bigger the gap is between the two setpoints, the more useful information there is, and the better the time constants are estimated. Figure 3.5 and Figure 3.6 confirm this intuition, since at least 10 days of data sampled at rate 10 minutes and a gap greater than 2 °C are needed to estimate the slow time constant. Whereas UA and g (not displayed) are robustly estimated, the slow time constant estimate may indeed explode when violating these conditions. Table 3.2 reinforces the fact that it is difficult to get a very accurate estimate of this parameter in any case, while a small bias remains for the fast time constant. The time-scaled method [MCPF10] is a promising track to address this issue.

3.4.2 Case 2: BESTLab data

Real data from BESTLab are used in this section. In particular, the indoor temperature and heating power are no longer simulated from nominal values of the R3C2 model, but measured in the two cells of the BESTLab experiment described in Section 2.2.1. Note that these real data are also generated in closed loop. The setpoint follows a predetermined schedule, and switches between 4 possible discrete levels, respectively at 19°C, 20°C, and 24°C. Furthermore, the dataset contains an 18-day period during which the setpoint is constant at 22°C. Several R3C2 models are identified with the EM algorithm, where each model corresponds to a random excerpt of BESTLab's database with varying length, from 3 days to 17 days.

Figure 3.7 shows the estimated physical parameters of the western cell. There are 80 datasets, each of them with 14 days of measurements. The results are displayed against the proportion of time during which the setpoint remains constant at 22°C: respectively less that 10% of the time, between 10% and 30%, 30% and 60%, 60% and 100%. Although there is no ground truth data to which our results can be compared, some conclusions may still be drawn. Indeed, the five-over insulated walls of either one of the two cells have very small heat transfer coefficients $U < 0.1$ W/(m^2.K). Hence, the thermal losses of the cells are caused mostly by the removable outside wall. This is why we report the U coefficient, which is simply the UA coefficient divided by the surface $\mathcal{S} = 13.5\text{m}^2$ of the tilted wall. Then, Figure 3.7 emphasizes the fact that the temperature setpoint needs to vary in order to get a plausible estimation. Indeed, the training set contains a long period of constant setpoint, the estimation of g degenerates (values very close to 1), whereas the slow time constant is unreasonably high. On the other hand, when the setpoint alternates frequently between several levels, the variance of the estimations decreases significantly. Moreover, the estimated parameters are physically plausible in such case. For instance, the estimated U value of a simple homogeneous wall range from 1.5 W/(m^2.K) to 2.3 W/(m^2.K) in [NBR+14], or from 1.62 W/(m^2.K) to 2.17 W/(m^2.K) in [NGRC15]. Given that BESTLab was originally created to design solutions for low-energy buildings, characterized by a small $U-$value, our estimation of U seems therefore credible and is typical of high energy performance buildings. Besides, the fact that the estimated solar transmittance coefficients g are closer to 1 than to 0 is consistent with the presupposed high energy performances of the cells, since g represents the proportion of the solar flux transmitted inside the cell. Finally, we note that the estimated time constants are somewhat elevated, which characterises a cell with strong inertia. This seems conceivable, because the high level of insulation of the cell makes it quite slow to cool the indoor air. Similar trends are observed for the southern cell (plots

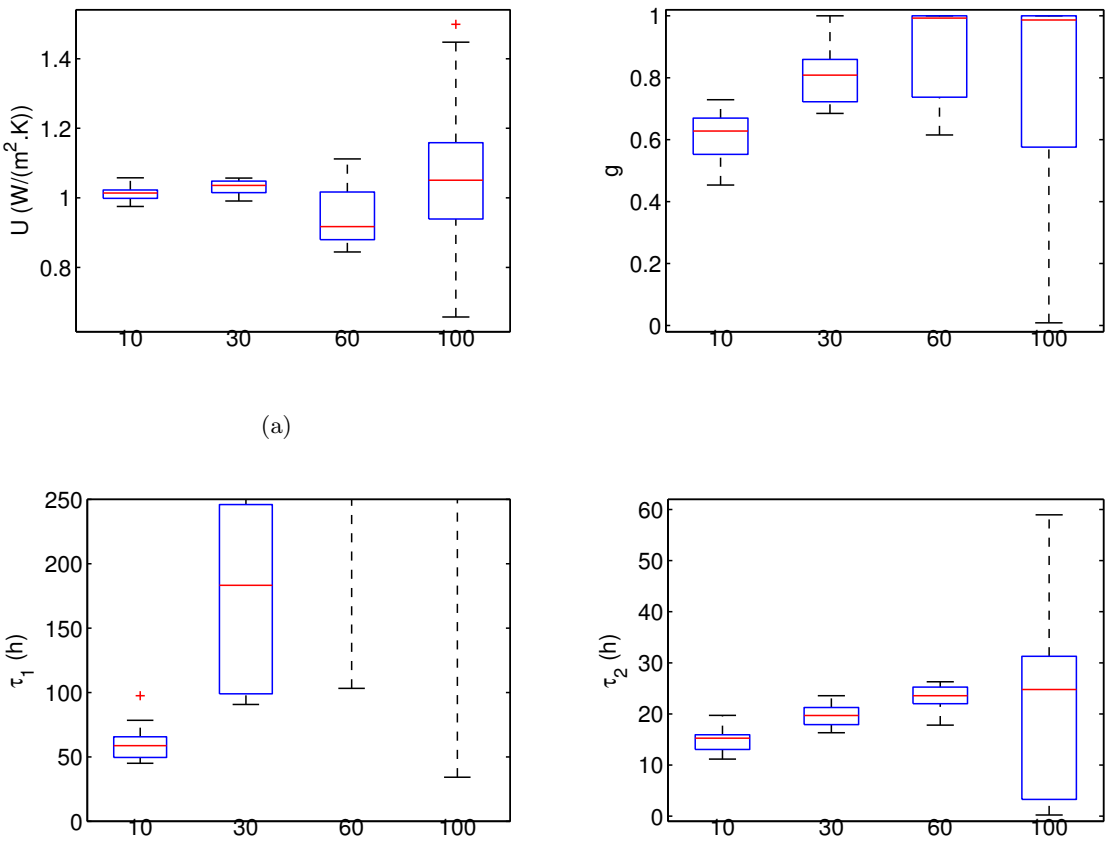

(a)

Figure 3.7: Estimation of the physical parameters of a R3C2 network for BESTLab's western cell, against the proportion of time during which the setpoint remains constant at 22°C: From left to right, top to bottom: heat loss coefficient U, solar transmission g, slow time constant τ_1, fast time constant τ_2. 80 experiments of duration 14 days.

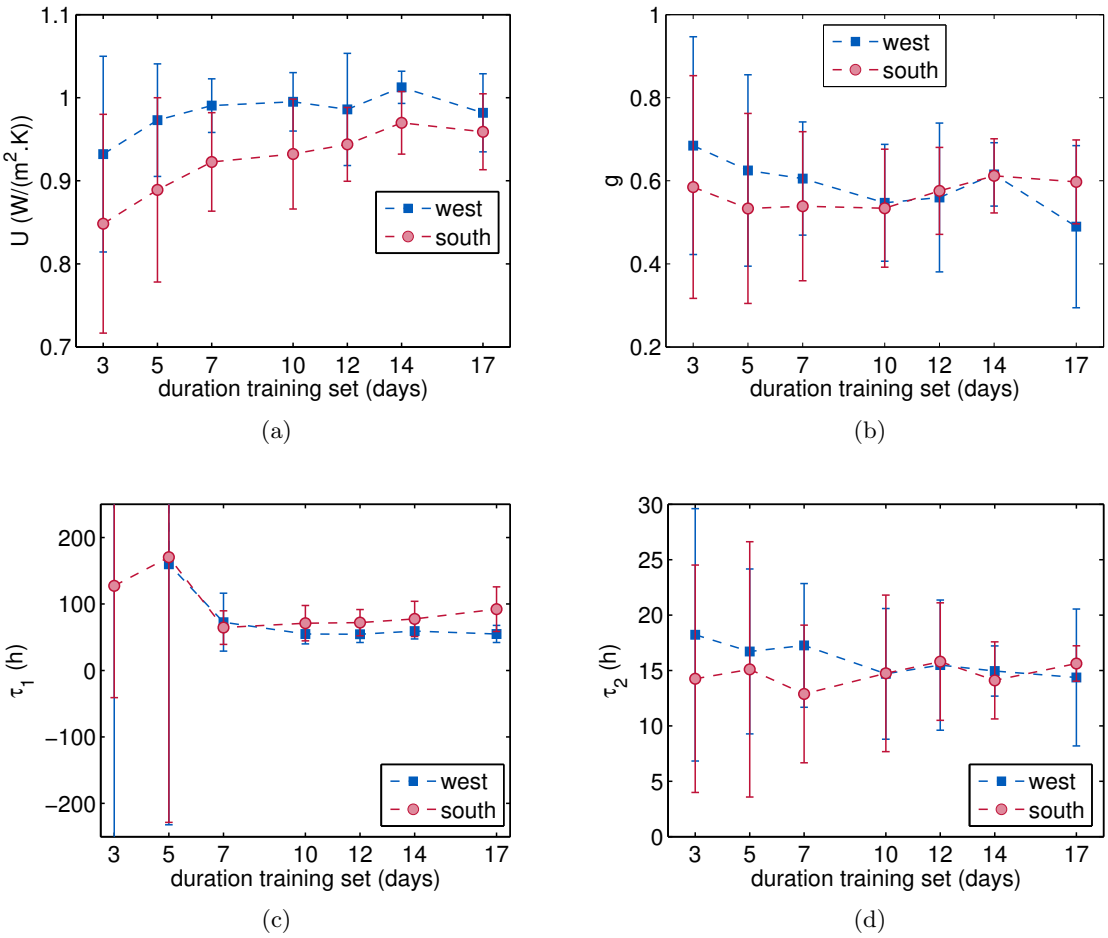

Figure 3.8: Estimated parameters of BESTLab's cells against length of the dataset. The error bars are centered around the mean estimations, their length is twice the standard deviation.

not shown).

The comparison between the estimated physical parameters of the western and southern cells is shown in Figure 3.8, with for datasets of varying length and with a setpoint that does not remain constant for a long time. Consistently with our findings on the simulated dataset, we find that at least one week of measurements are necessary before getting reliable estimations. Moreover, Figure 3.8(a) is in agreement with the physical intuition that the U value of the southern cell should be slightly lower than that of the western cell, since its orientation optimizes the received free solar gains. Finally, Figure 3.9 3.10 show the average of the simulated trajectories of the heating power and indoor temperature on a test data set for the western cell, in two different settings: with closed or open shutter, i.e. with or without solar radiation. These trajectories are simulated in closed loop from the estimated R3C2 model and the true parameters of the regulation. Whenever the solar flux does not affect the thermal behaviour of the cell, the simulated trajectories are accurate: root mean square error (RMSE) of 0.11°C for the average trajectory of the indoor temperature, total amount of energy consumed during the seven days overestimated by 6.1%. Figure 3.10 tends to indicate that these performances are degraded to a certain extent when adding the effect of the solar flux. In particular the estimated R3C2 models are too slightly slow when the building is in cooling mode. Nevertheless, the RMSE is 0.27°C for the indoor temperature, whereas the total amount of

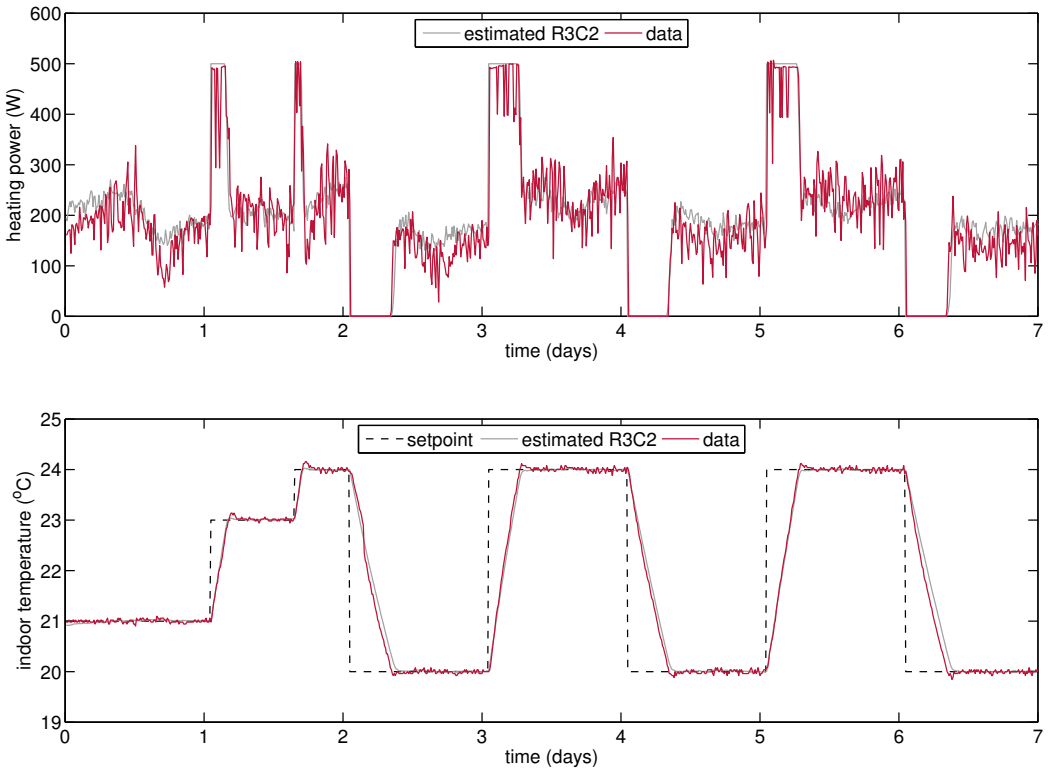

Figure 3.9: Simulated heating power and indoor temperature on test data without solar flux (closed shutter), BESTLab West.

energy consumed during the seven days is overestimated by 3.4%, which is even more accurate than without solar flux.

In conclusion, and although the BESTLab environment is a simplification of the real behaviour of an occupied building, the results presented in this section tend to validate the open-loop estimation of a R3C2 model for a building operating in closed loop.

3.5 Summary

In this chapter, we have investigated the first constraint limiting building identification in practice. Even if every input and output of the building model is correctly recorded, it remains indeed to take the data generating process into consideration. Although the data are generated in closed loop, we have given statistical and empirical arguments supporting the idea that the open-loop building model, "ignoring" the regulation. As such, we propose to learn the R3C2 model (or any equivalent electrical network) with a maximum likelihood estimator, obtained from the Expectation-Maximization algorithm. The numerical experiments, on both simulated and real data, confirm that this is a sensible approach. More specifically, in order to examine further the practical constraints, we have shown that a good estimation of the physical parameters of the R3C2 network requires at least 10 to 14 days of data, and as much variability as it may be on the setpoint signal.

In the next part, we will consider new constraints, directly linked to the actual availability of the inputs and output of the building model: how to collect these data, and how to adapt the identification procedure whenever some of them are partially or totally hidden?

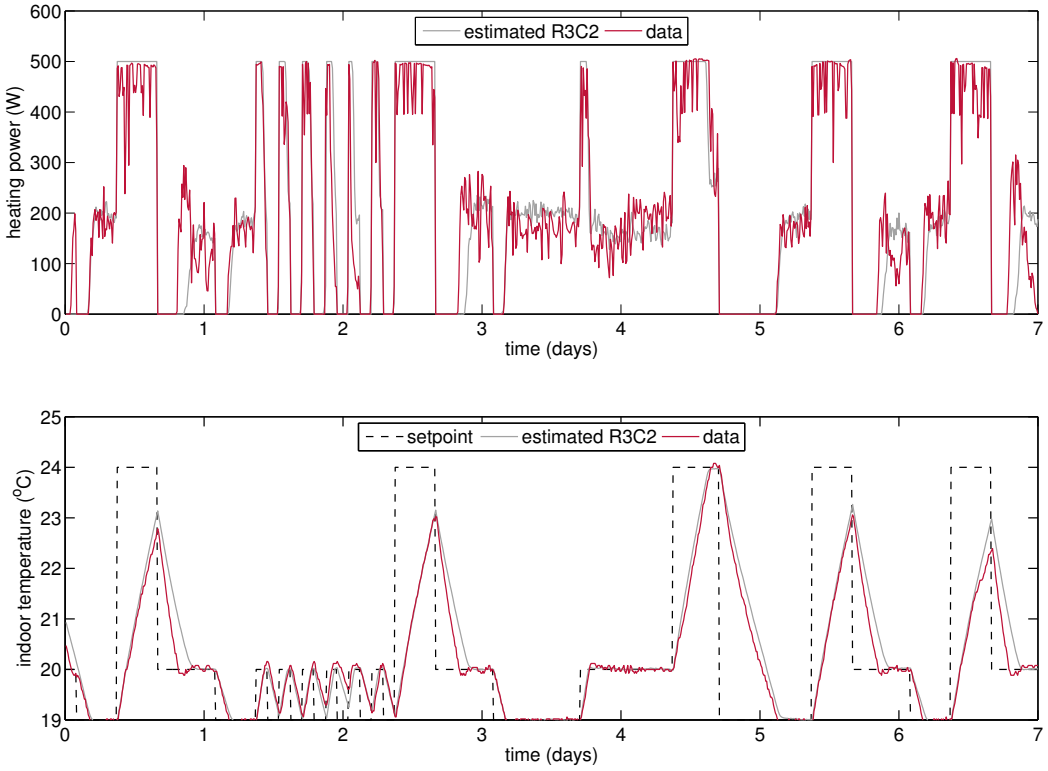

Figure 3.10: Simulated heating power and indoor temperature on test data with solar flux (open shutter), BESTLab West.

Part II

Identification under practical constraints

Chapter 4

Missing observations of the indoor temperature

In this chapter and the subsequent ones, we adopt the standpoint of an energy utility and examine the constraints that arise regarding the application of the identification method described in Chapter 3. This method relies primarily on the availability of on-site measurements of the inputs and output of the model. Yet, this perfect observability of the building is far from being obvious in practice. Hence, our aim is to question this availibility for each input and output of the building model, and adapt the identification procedure consequently.

To begin with, building identification relies first on the availability of the indoor temperature, which characterizes the thermal behaviour of the building. However, from the standpoint of an energy utility, this measure might not be available at all, since it requires a connected dedicated sensor not included in the electricity or gas meter. We assume thus in this chapter that the actual indoor temperature is hidden (while all the inputs of the building model are correctly measured). This implies that the model cannot be identified in open loop. Instead, a closed-loop model with the load curve as an alternative output is suggested. It is based on the availability of the setpoint temperature. Having access to the setpoint temperature but not to the actual indoor temperature is a reasonable hypothesis in some cases, corresponding for instance to a non-intrusive approach of the utility. It was suggested first in C. Zayane's thesis [Zay11].

The counterpart of not measuring the indoor temperature is that unlike the open loop identification described in Chapter 3, it is now necessary to assume that the regulator is of a certain structure, with unknown parameters. In the context of building identification, this introduces a non-linearity in the model, because of the saturation of the heating power, which is upper-bounded by the subscribed power and necessarily positive since heating only and not cooling is considered in this work. This non-linearity was not accounted for in [Zay11], which is why we suggest a new estimation algorithm. This is developed, together with the global closed loop model, in Section 4.1. We choose a Bayesian framework in order to solve the identification problem. The algorithm, described in Section 4.2, is based on the Hybrid Monte Carlo (HMC) sampling technique which belongs to the class of Markov chain Monte Carlo (MCMC) sampling methods. We apply HMC to the building identification problem, and show some numerical experiments, in Section 4.3.

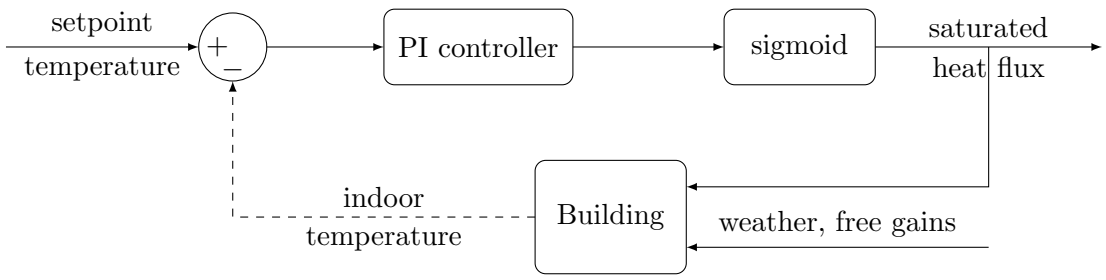

Figure 4.1: Building indoor temperature regulated in closed loop, with saturation of the heating power via the sigmoid function. The indoor temperature is not observed, and the saturated power is seen both as an input of the building model and the observed output of the closed loop system.

4.1 A modified closed loop model

Since the indoor temperature is not observed, an open-loop estimation of the building thermal model, such as suggested in Chapter 3, is not possible. Moreover, the building is actually exploited in closed-loop, the indoor temperature being regulated such that it follows a given setpoint signal. This would suggest thus to use a closed-loop identification technique as advocated in [FL99]. However, those methods are precluded because they depend on the availability of the output of the open-loop system, that is precisely the missing indoor temperature.

In order to account for the missing observation of the indoor temperature, we need thus to change our interpretation of a building system. Similarly to the previous work in [Zay11], our method consists thus into considering the system in its entirety, that is the building model, the feedback loop and the controller. The measured heating load curve Q_h is now the observed output of the sytem, and is explained by the setpoint temperature, the weather conditions (outdoor temperature, solar gains) and the free internal gains. This is depicted in Figure 4.1. The building model is the R3C2 network described in Chapter 2. The regulation is modelled as a Proportional - Integral (PI) controller, whose output is saturated via the sigmoid function:

$$Q_h = \Phi\left[K\left(T_r - \vartheta_i + \frac{1}{\tau}\int (T_r - \vartheta_i)\right)\right], \tag{4.1}$$

where T_r is the known setpoint signal, ϑ_i the hidden indoor temperature and the gain K and the time constant τ are the unknown parameters of the PI controller. The sigmoid function Φ is

$$\Phi : x \in \mathbb{R} \mapsto \frac{Q_h^{max}}{1 + \exp(-\lambda(x - Q_h^{max}/2))} \in (0,\, Q_h^{max}), \tag{4.2}$$

with Q_h^{max} the limiting power and $\lambda := 4/Q_h^{max}$. Φ is a continuously differentiable approximation of the saturation function. The coefficient λ is such that the approximation is exact at $x = Q_h^{max}/2$, that is for the linear regime (see Figure 4.2). Note that the anti-windup effect is not included in this model.

4.2 Hybrid Monte Carlo sampling method

4.2.1 Hamiltonian dynamics

Hybrid Monte Carlo, also named Hamiltonian Monte Carlo (HMC), is a Markov chain Monte Carlo (MCMC) method for sampling from complex high-dimensional distributions, introduced re-

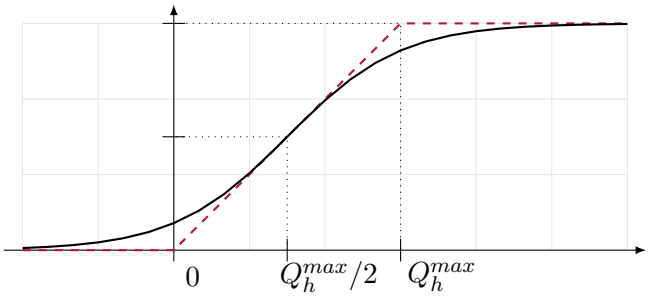

Figure 4.2: Sigmoid (solid line) and saturation (dashed line) functions.

cently in [DKPR87]. MCMC algorithms are iterative strategies aiming at generating samples from a target distribution, known up to a normalizing constant, in which the samples are the state of a Markov chain. The key is to construct a transition kernel whose invariant distribution is the distribution of interest. See e.g. [ADFDJ03] for a comprehensive introduction to the different state-of-the-art MCMC strategies, [PSC+16] for a survey of stochastic simulation methods for Bayesian inference. The main advantage of HMC over other standard MCMC techniques such as Metropolis - Hastings is that it generates non-local moves, avoiding thus the random-walk behaviour of the simulated trajectories. For simulating such trajectories, HMC uses an auxiliary variable and is based first of all on an analogy between the target probability distribution and the potential energy of a mechanical system. Let us thus consider a mechanical system of position $q \in \mathbb{R}^{n_q}$ and impulsion $p \in \mathbb{R}^{n_q}$, with total energy $H(q,p)$ given as the sum of a potential $U(q)$ and a kinetic term $K(p)$

$$H(q,p) = U(q) + K(p). \tag{4.3}$$

With ∇ denoting the gradient operator, the Hamiltonian equations describing the evolution of the sytem in the phase space are:

$$\begin{cases} \frac{dq}{dt} &= \frac{\partial H}{\partial p} = \nabla K(p) \\ \frac{dp}{dt} &= -\frac{\partial H}{\partial q} = -\nabla U(q) \end{cases}. \tag{4.4}$$

The $t-$flow φ^t of (4.4) is the map $(q(0),p(0)) \in \mathbb{R}^{2n_q} \mapsto (q(t),p(t)) \in \mathbb{R}^{2n_q}$. It can be shown (see e.g. [Liu08]) that this flow is time-reversible (φ^t is one-to-one and has an inverse), preserves the volumes $\mathrm{d}q\,\mathrm{d}p$ in the (q,p) space and keeps the Hamiltonian invariant in time ($\frac{dH}{dt} = 0$). These three properties are key to the HMC sampling method (see Section 4.2.2).

Next, (4.4) is approximated by time discretization for the numerical simulation of the dynamics. The time step, herafter denoted ε, induces a small error, and consequently the Hamiltonian is not preserved anymore. The integrator must be chosen so that it does satisfy the other two properties, time reversibility and volume preservation. Here, we employ the standard Verlet/leapfrog method, which consists of a half-step for the impulsion variable followed by a full step for the position and a second half-step for the impulsion:

$$\begin{cases} p(t+\varepsilon/2) &= p(t) - \frac{\varepsilon}{2}\nabla U(q(t)) \\ q(t+\varepsilon) &= q(t) + \varepsilon \nabla K(p(t+\varepsilon/2)) \\ p(t+\varepsilon) &= p(t+\varepsilon/2) - \frac{\varepsilon}{2}\nabla U(q(t+\varepsilon)) \end{cases}. \tag{4.5}$$

Volume preservation is easily checked by computing the jacobians of the three steps in (4.5): each of them is 1. The global error of the leapfrog method is $O(\varepsilon^2)$, and so is the error on H (see [Nea11]). This method is difficult to beat in practice, and easy to implement, which motivates our choice. See e.g. [Laj11] for other discretization schemes.

4.2.2 MCMC using Hamiltonian dynamics

Consider now a $\mathbb{R}^{n_q}$-valued random variable q with probability density $\pi(q)$ which we wish to sample from. Notions in statistical physics (see e.g. [BB06]) help us to relate the target density $\pi(q)$ to a potential energy $U(q)$: $\pi(q) = \frac{1}{Z_U}\exp(-U(q))$. Ignoring the normalization constant, that cancels out in MCMC algorithms, we get

$$U(q) = -\log(\pi(q)). \tag{4.6}$$

It is assumed in the sequel that U is differentiable and non-zero for all q. An HMC-like algorithm includes then three elements: an enlarged phase space with a Hamiltonian flow, a geometrical numerical integrator with suitable properties and an accept/reject rule.

Hamiltonian flow In the analogy with a mechanical system, the random variable q is our position vector, and an enlarged phase system is formed by artificially adding a corresponding impulsion vector p. Its total energy is separable $H(q,p) = U(q) + K(p)$. The Hamiltonian flow is then given by (4.4). The kinetic energy is chosen to be of the form

$$K(p) = \sum_{i=1}^{n_q} \frac{p_i^2}{2m_i}. \tag{4.7}$$

That is p is Gaussian with zero mean and a diagonal covariance matrix M, with diagonal elements $m_1, \ldots, m_{n_q}$ to be selected by the user (see below). Moreover, given the volume and Hamiltonian preservation properties of φ^t, the flow preserves also the probability measure in the phase space $\mathbb{R}^{2n_q}$, with density the joint one $\mathit{\Pi}(q,p)$

$$\mathit{\Pi}(q,p) = \frac{1}{Z_H}\exp(-H(q,p)) = \frac{1}{Z_U}\exp(-U(q))\frac{1}{Z_K}\exp(-K(p)) = \pi(q)\pi(p). \tag{4.8}$$

(4.8) emphasizes the fact that by construction, the position q and the impulsion p are independent. A sample of q can thus be obtained simply by marginalizing out p when sampling from $\pi(q,p)$. Hence, if $q(0) \sim \pi(q)$ and $p(0) \sim \mathcal{N}(0, M)$ then $q(t) \sim \pi(q)$.

Numerical integrator In practice, the Hamiltonian flow is simulated from the Verlet/leapfrog scheme (4.5). Given the form of $K(p)$, the full step in position reads more specifically

$$q(t+\varepsilon) = q(t) + \varepsilon M^{-1} p(t+\varepsilon/2). \tag{4.9}$$

Accept/reject rule Since the leapfrog scheme does not preserve the Hamiltonian value, an accept/reject rule is introduced to remove this bias. Consequently, the HMC algorithm is hybrid in the sense that it combines a Gibbs step and a Metropolis update:

1) Draw p from $\pi(p|q) = \pi(p)$ (by independence (4.8)). This step is easy to perform, since p is Gaussian.

Algorithm 4.1 Hybrid Monte Carlo

Inputs: M, ε, L, target density $\pi(q)$.

0. Initialize $q^{(0)}$ and set $k = 1$.
1. *Gibbs step:* draw $p^{(k-1)} \sim \mathcal{N}(0, M)$.
2. *Metropolis update:*
   a) draw $\varepsilon^{(k)}$ uniformly from $[\varepsilon - \delta, \varepsilon + \delta]$, $\delta > 0$ small;
   b) Perform L leapfrog steps with stepszie $\varepsilon^{(k)}$ according to (4.5) or Algorithm 4.2. Label the new state (q^*, p^*).
   c) Set $(q^{(k)}, p^{(k)}) = (q^*, p^*)$ with probability r computed from (4.10), or $(q^{(k)}, p^{(k)}) = (q^{(k-1)}, p^{(k-1)})$ with probability $1 - r$.
3. If convergence is reached, terminate, otherwise set $k \leftarrow k + 1$ and return to step 1.

Output: samples $q^{(k)}$ of $\pi(q)$.

2) From the current state (q, p), simulate L steps of the Hamiltonian dynamics using the leapfrog method (4.5) with stepsize ε to reach the proposal (q^*, p^*). Negate p^* and accept the new state $(q^*, -p^*)$ with probability

$$r = \min\left[1, \exp(H(q,p) - H(q^*, -p^*))\right]. \tag{4.10}$$

Negating p^* ensures time reversibility. It is however skipped in pratice, since $K(p^*) = K(-p^*)$ and p^* is then replaced at the first step of the next iteration. It is shown in [Liu08] that this hybrid procedure is valid, in the sense that it leaves the joint canonical distribution of (q, p) invariant (detailed balance condition).

We see also from (4.10) that the proposals are more likely to be accepted if they tend to decrease the total energy, which means with (4.6) that the algorithm does look for regions of high densities for q. (4.5) shows that the trajectory of q follows $\nabla K(p) \propto p$ and p is directed according to ∇U. HMC avoids thus the inefficient random walk behaviour of the Metropolis algorithm. However, the drawbacks of the method is that it has more parameters to be tuned than Metropolis. Indeed, the mass matrix M, the time step ε and the length L must be chosen carefully.

4.2.3 Practical improvements

The HMC algorithm is summarized in Algorithm 4.1. This standard procedure may be improved in several ways. For instance, [BPSSS11] suggests to improve the acceptance rate in high dimensions by generalizing Algorithm 4.1 to an infinite-dimensional Hilbert space. Here, we will focus on the ergodicity of the algorithm, the tuning of the step size ε and the mass matrix M, and the handling of constraints in q.

Ergodicity

The question of the ergodicity of the algorithm has been raised in [Mac89, Nea11]: the chain should asymptotically converge to the target distribution, and not be trapped in some subset of the

state space. In order to ensure it in practice, the solution is to randomly choose ε from some small interval around a given nominal value at the beginning of each Hamiltonian trajectory.

Tuning ε and M

Tuning ε is the crucial point when applying the HMC procedure (see e.g. [BPR+13]). Too large an ε would lead to a very small acceptance rate, whereas too small an ε would either waste computational time or generate a random walk behaviour [Nea11]. ε should not be larger than the typical length of the most constrained direction of q. Yet other directions might be much less constrained, resulting in a slow exploration of the latter. This can be alleviated by an appropriate choice of the diagonal elements of the mass matrix M: if q_i, $i = 1, \dots, n_q$ scales with s_i, set $m_i = 1/s_i^2$. (4.5) is then equivalent to a leapfrog scheme with stepsize $\varepsilon_i = \varepsilon s_i$ in each direction $i = 1, \dots, n_x$ and mass matrix identity. [Nea11] indicates that L and ε are not so intricated, and it is possible to tune them by trial-and-error.

Handling constraints

An advantage of HMC is that it can be easily adapted to handle constraints of the form $q_i \leq u_i$ or $q_i \geq l_i$ for any $i \in \{1, \dots, n_q\}$. The method, taken from [Nea11], can be named 'billiard HMC': if a leapfrog step is such that q_i violates the limit, then q_i bounces off the boundary, and p_i is negated.

The principle is the following. First of all, note that the leapfrog scheme (4.5) is a special case of the discretization by splitting scheme. In this more general setting, the Hamiltonian $H(q,p)$ is written as the sum of k terms

$$H(q,p) = H_1(q,p) + \cdots + H_k(q,p), \tag{4.11}$$

where each H_l, $l = 1, \dots, k$ is twice differentiable and $H_l(q,p) = H_{k-l+1}(q,p)$. The discretization by splitting is obtained by applying successively the dynamics (4.4) to each subsystem with energy H_l, $l = 1, \dots, k$. Applying the scheme with $k = 3$ and $H_1(q,p) = H_3(q,p) = U(q)/2$ and $H_2(q,p) = K(p)$ yields the leapfrog scheme (4.5).

Assume now that a constraint $q_i \leq u_i$ for some $i \in \{1, \dots, n_q\}$ holds. The idea is to propose a new splitting scheme close to the one leading to the leapfrog. The constraint is included in the potential energy:

$$U(q) = \lim_{r \to \infty} [U_*(q) + C_r(q_i, u_i)], \quad \text{where } C_r(q_i, u_i) := \begin{cases} 0 & \text{if } q_i \leq u_i \\ r^{r+1}(q_i - u_i)^r & \text{if } q_i > u_i \end{cases}. \tag{4.12}$$

U_* is the unconstrained potential, while the term $C_r(q_i, u_i)$ tends to infinity if the constraint is broken. In such case, U tends also to infinity, and the acceptance probability of this state is null. The Hamiltonian is then split in three terms according to

$$H(q,p) = \frac{1}{2}U_*(q) + [K(p) + C_r(q_i, u_i)] + \frac{1}{2}U_*(q). \tag{4.13}$$

The first and last terms yield the same half-step in p as the leapfrog scheme. If the full step for q leads to $q_i(t+\varepsilon) \leq u_i$ then necessarily $C_r = 0$ along the trajectory and the position update is also unchanged: (4.13) is equivalent to the leapfrog scheme. On the other hand, if the position update scheme leads to $q_i(t+\varepsilon) > u_i$, then the trajectory 'bounces off' the limit, and returns with opposite impulsion (see [Nea11]). This variation of the leapfrog scheme is described in [Nea11].

Algorithm 4.2 Leapfrog scheme with constrained state space

1. Half impulsion update: $p(t+\varepsilon/2) = p(\varepsilon) - \frac{\varepsilon}{2}\nabla U_*(q(t))$.
2. Position update : for each variable $i = 1, \dots, n_q$:
   (a) let $p_i' = p_i(t+\varepsilon/2)$;
   (b) let $q_i' = q_i(t) + \varepsilon p_i'/m_i$;
   (c) if q_i is constrained, repeat the following until q_i' satisfies all the constraints:
      i. if q_i has an upper constraint u_i and $q_i' > u_i$,
         let $q_i' = u_i - (q_i' - u_i)$ and $p_i' = -p_i'$;
      ii. if q_i has a lower constraint l_i and $q_i' < u_i$,
         let $q_i' = l_i + (l_i - q_i')$ and $p_i' = -p_i'$;
   (d) let $q_i(t+\varepsilon) = q_i'$ and $p_i(t+\varepsilon/2) = p_i'$
3. Half impulsion update: $p(t+\varepsilon) = p(\varepsilon/2) - \frac{\varepsilon}{2}\nabla U_*(q(t+\varepsilon))$.

4.3 Application to building identification

4.3.1 Principle

The idea of applying HMC sampling techniques to the problem of building identification was suggested in [Laj11], with application to a R1C1 network with saturated Proportional controller. We extend therefore its application to the higher-order R3C2 network. In this chapter, we set $\theta = \begin{pmatrix} C_r & C_s & z_f & z_o & z_i & K & z_\tau \end{pmatrix}^\intercal$, where K and $z_\tau := 1/\tau$ are the parameters of the controller. Moreover, we define $d := z_\tau \int (T_r - \vartheta_i)$, which is the integrator component of the PI controller.

The estimation of θ is treated as a Bayesian inversion problem [KS05, Sär13]. We assume thus that θ is a random variable with prior distribution $\pi(\theta)$. In our state-space framework, the full posterior distribution is, using Bayes' theorem and ignoring the normalizing term:

$$p(x_{0:N}, \theta | y_{1:\mathrm{T}}) \propto p(y_{1:\mathrm{T}} | x_{0:\mathrm{T}}, \theta) p(x_{0:\mathrm{T}} | \theta) \pi(\theta), \tag{4.14}$$

where the hidden state is $X_t := \begin{pmatrix} \vartheta_i^t & \vartheta_s^t & d^t \end{pmatrix}^\intercal$, with ϑ_s the hidden structure temperature of the R3C2 network, and the observation is $Y_t := Q_h^t$. Given the Markovian properties of the model, the potential U associated to the full posterior distribution $p(x_{0:\mathrm{T}}, \theta | y_{1:N})$ (4.14) is the sum of the following terms, where Q_{fg} represents the free heating gains, Q_s the solar gains and ϑ_o the outdoor temperature:

$$\begin{aligned} U_1(\theta, \vartheta_i, \vartheta_s, d) = \frac{1}{2\sigma_1^2} \sum_{t=0}^{\mathrm{T}-1} \{ & \Phi[K(T_r^{t+1} - \vartheta_i^t + d^t)] + Q_{fg}^{t+1} - C_r(\vartheta_i^{t+1} - \vartheta_i^t) - z_i(\vartheta_i^t - \vartheta_s^t) \\ & - z_o(\vartheta_i^t - \vartheta_o^{t+1}) \}^2, \end{aligned} \tag{4.15a}$$

$$U_2(\theta, \vartheta_i, \vartheta_s) = \frac{1}{2\sigma_2^2} \sum_{t=0}^{\mathrm{T}-1} \{ Q_s^{t+1} - C_s(\vartheta_s^{t+1} - \vartheta_s^t) - z_i(\vartheta_s^t - \vartheta_i^t) - z_o(\vartheta_s^t - \vartheta_o^{t+1}) \}^2, \tag{4.15b}$$

$$U_3(\theta, \vartheta_i, d) = \frac{1}{2\sigma_3^2} \sum_{t=0}^{\mathrm{T}-1} \left\{ -(d^{t+1} - d^t)) - z_\tau(\vartheta_i^t - T_r^{t+1}) \right\}^2, \tag{4.15c}$$

$$U_4(\theta, \vartheta_i, d) = \frac{1}{2\sigma_4^2} \sum_{t=1}^{\mathrm{T}} \left\{ Q_h^t - \Phi(K(T_r^t - \vartheta_i^t + d^t)) \right\}^2, \tag{4.15d}$$

$$U_5(\theta) = -\log(\pi(\theta)), \tag{4.15e}$$

where t denotes a discrete time index and Y_t, T_r^t, ϑ_o^t, Q_{fg}^t and Q_s^t are measured for $t = 1, \dots, \mathrm{T}$. The first two potentials in (4.15a)-(4.15b) correspond to a first-order (instead of exact) discretization of Kirchhoff's equations (2.1), with stepsize δ omitted for sake of clarity, and state noise process $v_t \sim \mathcal{N}\left(0, \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_2^2 \end{pmatrix}\right)$. Similarly, the third potential U_3 (4.15c) corresponds to a first-order discretization of the state equation in d, and U_4 (4.15d) to the observation equation.

In this work, we suggest to estimate jointly the parameters θ and the hidden states. Hence, the potentials U_1 to U_5 are seen as functions not only of θ, but also of the trajectories ϑ_i^t, ϑ_s^t and d^t, for $t = 0, \dots, \mathrm{T}$, and the gradient ∇U is a vector of dimension $3(\mathrm{T}+1)+7$, where T is the number of observations (the noise variances $\sigma_1^2, \dots, \sigma_4^2$ are not learned).

Moreover, we add a potential corresponding to the priors on the trajectories of the hidden states. In particular, the knowledge of the setpoint temperature, a supposedly good approximation of the indoor temperature, informs *a priori* accurately on ϑ_i. We choose therefore to penalize strongly the deviations of ϑ_i to the setpoint T_r with an *a priori* potential of the form $\sum_t (T_r^t - \vartheta_i^t)^4$. On the other hand, the priors on θ, ϑ_s and d are less accurate and choosen Gaussian, which corresponds to quadratic potentials. Ultimately, the potential U is thus the sum of the five potentials (4.15a)-(4.15e), and of the priors on the trajectories. We note that U is differentiable and ∇U is well-defined everywhere, because of the $\mathcal{C}^\infty$ smooth sigmoid saturation function Φ (4.2).

4.3.2 Numerical illustration

Validation procedure

This method is applied to solving the building identification with missing observations of the indoor temperature problem for several one-week datasets. The data are generated according to the procedure described in Chapter 2. In particular, meteorological records from January 2016 are used, as well as free internal gains corresponding to an office building. For each simulation, θ is initialized randomly with initial error between -100% and +25%. The Gaussian prior on θ has a mean which is also randomly drawn within $[-100\%, +25\%]$ of the true value, and a diagonal covariance with standard deviations set to 35% of the true values of the parameters. The initial series of the hidden states are drawn randomly around fixed values: for all $t \in \{0, \dots, N\}$, we set $\vartheta_i^{t,(0)} \sim \mathcal{N}(16, 0.02)$, $\vartheta_s^{t,(0)} \sim \mathcal{N}(10, 0.2)$ and $d^{t,(0)} \sim \mathcal{N}(3, 0.2)$. The prior on ϑ_i^t is defined by the setpoint temperature, whereas the priors on ϑ_s and d are their initial distributions $\mathcal{N}(16, 0.02)$ and $\mathcal{N}(10, 0.2)$, respectively.

The key parameters ε (step-size) and L (number of steps) are tuned by trial and error to $\varepsilon = 5.10^{-5}$ and $L = 150$. The diagonal coefficients of the mass matrix M are tuned according to the scale of the parameters θ, and the square roots of the scales of the hidden state. Indeed, the parameters θ have a strong influence on the whole time series and need therefore smaller impulsion compared to each hidden state x_t, which affects only one instant t [Laj11]. Finally, we add the constraint that $\theta \in (\mathbb{R}_+^*)^7$. We used $N_{iter} = 1000$ iterations, and the point estimates of the

Table 4.1: Estimated parameters with twice their standard errors for HMC, on 100 experiments. 5 outliers are not taken into account.

Parameter	True value	HMC
UA	0.354	0.338 ± 0.007
g	0.385	0.288 ± 0.106
τ_1	33.7	38.3 ± 14.3
τ_2	3.6	3.5 ± 0.9
K	2.79	1.33 ± 0.14
z_τ	0.02	0.019 ± 0.0085

parameters are the averages over the last 100 iterations. This procedure is repeated 100 times, with the same set of true parameters but different weather conditions.

Results

Table 4.1 shows the mean of the estimated physical parameters of the R3C2 model and of the parameters of the controller. The performances are encouraging, with a good accuracy on the heat loss coefficient UA and on the two time constants. The estimation of the solar transmission g is less accurate, although the order of magnitude remains correct. However, the average acceptance rate is 48% (see Figure 4.4 for the distribution), which means that slightly decreasing the stepsize ε would improve the accuracy of the leapfrog scheme. By way of comparison, [BPR+13] identifies the asymptotically optimal acceptance probability, which is 0.651.

For each of the 100 estimations of the parameters, we simulated the heating power and the indoor temperature with the setpoint and meteorological data contained in the training set. The averages of these simulated trajectories are plotted in Figure 4.3. The root mean square error (RMSE) of the heating power is 7% of the average power, whereas the estimated energy consumption over the week is 1.2% higher than the actual consumption. Hence, although the parameters of the regulation estimated in Table 4.1 are less accurate than those of the R3C2 network, the simulated trajectories are in agreement with the observations. Most of the divergence between the two heating loads occurs at saturation (see Figure 4.3(a)), which is consistent with the saturation function used to generate the data in Chapter 2 by the sigmoid function in Figure 4.2. For an increased accuracy, we suggest therefore to replace the sigmoid function by the convolution of the saturation function with a Gaussian kernel, as in [Laj11]. Similarly, the simulated indoor temperature is closed to the actual trajectory, with a normalized root mean square error less than 4% of the average indoor temperature.

In order to further validate the estimation, we used the Monte Carlo average of the estimated parameters on the 100 experiments to simulate the heating power and the indoor temperature under different weather conditions. Indeed, we use data from November 2015, whereas the parameters were learned with data from January 2016. The simulated and true trajectories are shown in Figure 4.5. The predicted power has an RMSE representing 20% of the average heating power, but the estimated total consumption over the entire simulation period is only 6% smaller than the true consumption. The predicted indoor temperature is in close agreement with the true measure, with an RMSE of 0.35°C, corresponding to a normalized RMSE of 1.9%.

These results accentuate thus the good performances of the estimation procedure suggested in this chapter. This is partly due to the inclusion of the informative prior based on the setpoint signal for the indoor temperature. Hence, the joint estimation of the hidden states and of the parameters

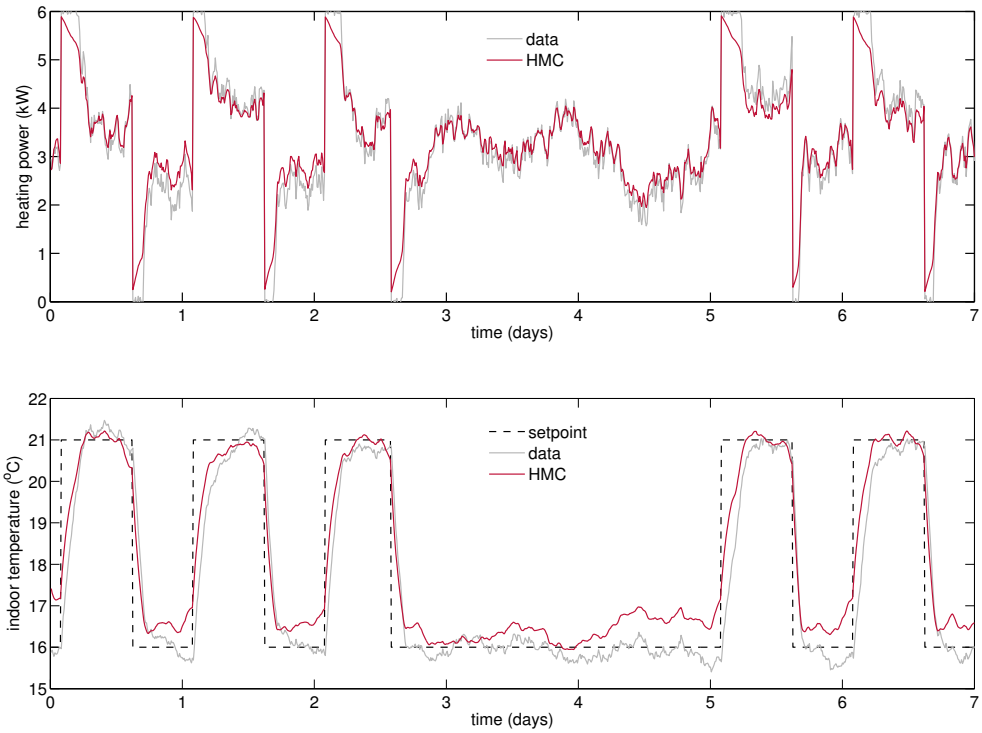

Figure 4.3: In red, average of (a) the heating power and (b) the indoor temperature simulated with the 100 estimations of the parameters on the training set (meteorological data from BESTLab, January 2016). The true data are in grey.

Figure 4.4: Histogram of the acceptance rates of the HMC sampler, 100 experiments.

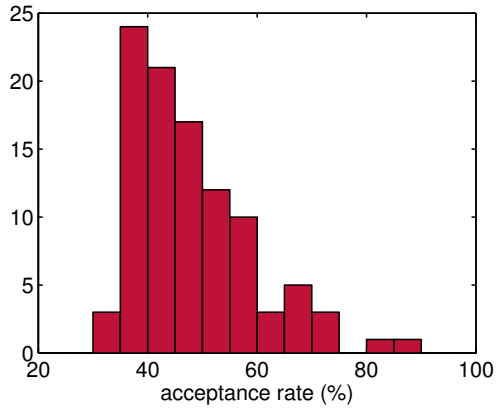

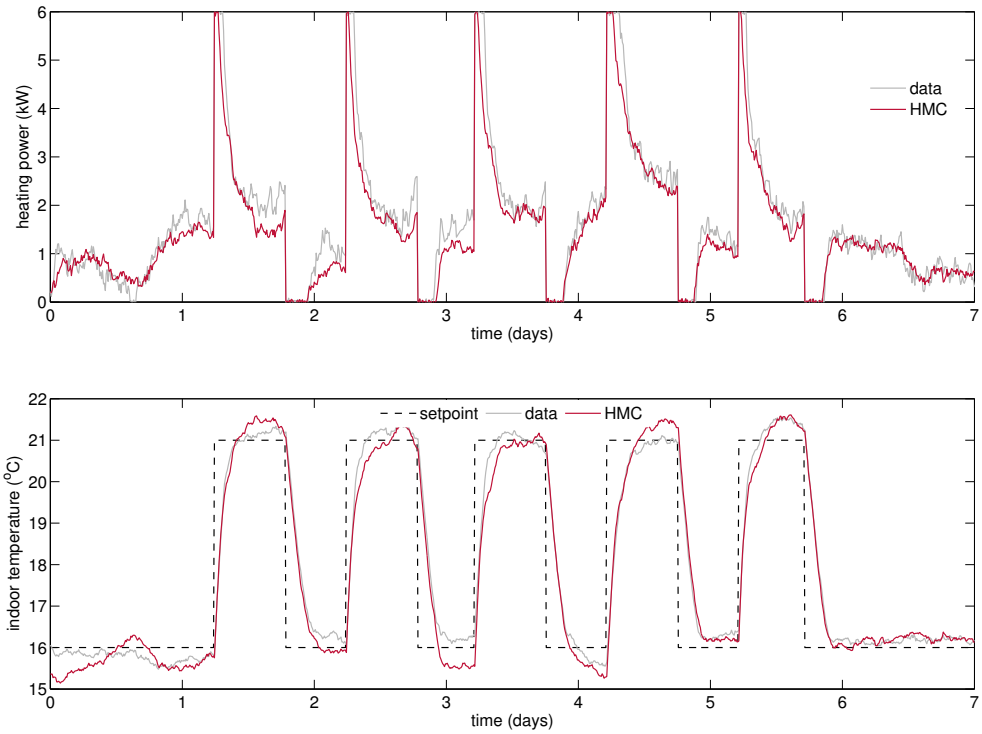

Figure 4.5: In red, (a) heating power and (b) indoor temperature simulated with the mean of the 100 estimations of the parameters on a test set with meteorological data from BESTLab, November 2015. The grey curves are generated with the true parameters.

of the R3C2 network and of the controller based on the HMC sampling technique is an efficient way of handling the missing observations of the indoor temperature and the nonlinearity of the controller. The main difficulty consists in tuning carefully the parameters of the HMC algorithm, in particular the stepsize ε and the number L of leapfrog steps.

4.4 Summary

In this chapter, we start investigating the application of building identification in a constrained context of missing or partially observed data. We have assumed that the output of the thermal building model, namely the indoor temperature, was not recorded. This corresponds to a non-intrusive approach of the energy utility, which does not have access *a priori* to a thermometer inside the building. Since the output of the model is not measured, we suggest a closed loop model whose new output is the regulated heating power. This model includes a nonlinear regulation, with the saturation of the command. We suggest to estimate jointly the hidden states and the parameters of this nonlinear system by means of the Hybrid Monte Carlo (HMC) algorithm. The numerical evaluation shows promising results on simulated data, with a good estimation of the physical constants of the R3C2 model. The main difficulty lies in the tuning of the key parameters of the HMC algorithm, which requires further examination.

In the next chapters, we continue the analysis of the practical constraints that limit the application of building identification for an energy utility. Chapter 5 is more particularly devoted to handling the uncertainty inherent to the measurement of the outdoor temperature.

Chapter 5

Bias reduction of the outdoor temperature

We address in this chapter the problem of the uncertainty of the measure of the outdoor temperature, a key input to any thermal building model. Indeed, the context of the smart homes and buildings prevents us from assuming that the measured outdoor temperature is the "true" noise-free input of the building system. Instead of having access to the accurate measure provided by a controlled climate station, an energy utility may only have access in practice to a measurement from a connected object, such as a connected weather station. This measurement lacks accuracy, as further exemplified in Section 5.1. Hence, the aim of this chapter is to adapt the estimation procedure of the open loop R3C2 model described in Chapter 2, in order to account for this uncertainty. The other inputs (solar radiation, heating flux) and the indoor temperature are supposedly correctly measured in this chapter.

5.1 Motivation

To the best of our knowledge, the outdoor temperature is seen as a known, correctly measured, input of the thermal model in the rest of the building identification literature. This means that in the literature, the identification algorithms are tested either with an average outdoor temperature in order to generate a simulated dataset (see e.g. [MCPF10, HGP12]), or with a measure from an on-site climate station (see e.g. [BM11, FABG14]). In both cases, the measured or statistical outdoor temperature is regarded as the noise-free input of the system, and neither the availability or the quality of the data are questioned.

Yet, it is unlikely, at a more global scale and for any building, to have access to a scientific climate station nearby, while an actual measurement, and not a statistical signal, is required. Consequently, the two aforementioned solutions meet with practical limitations, and the building microclimate remains unavailable. An answer to the issue of the data availability may come from the emerging context of the Internet of Things (see Introduction). One may thus have access to *a* measurement of the outdoor temperature, whether it is directly *via* a connected weather station for instance, or *via* an open access database such as the Weathermap. However, compared to the fully controlled climate station, the quality of this measurement suffers two drawbacks:

1. the temperature sensor is less accurate than a scientific probe (for instance the Netatmo weather station has an accuracy of $\pm 0.3°C$, whereas the sensors from the EDF R&D laboratory BESTLab have an accuracy of $\pm 0.1°C$);

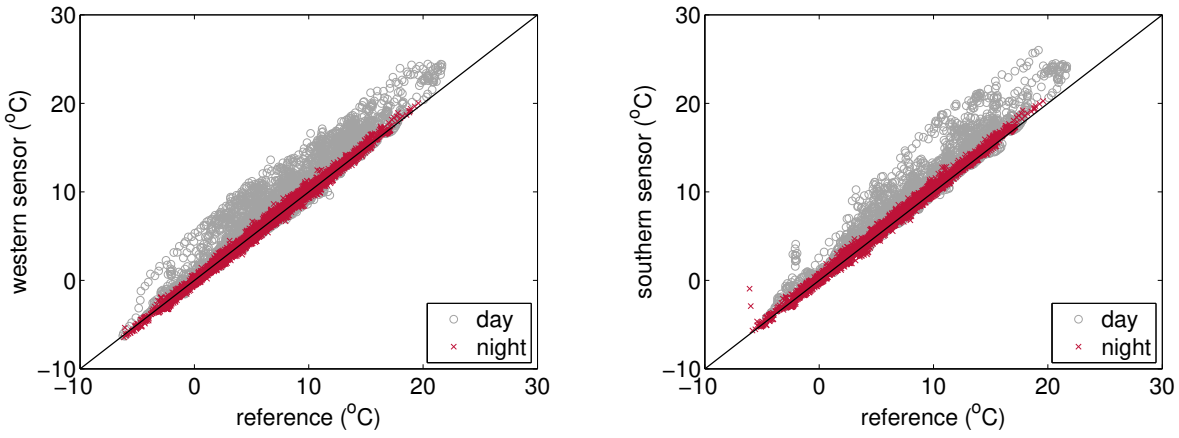

(a) (b)

Figure 5.1: Measurements of the dry-bulb outdoor temperature with a sensor with a (a) western or (b) southern orientation, compared to a reference temperature measured at the same location (data from November 2015 to February 2016).

2. an energy utility does not know the precise location of the sensor, which is up to the user. Hence, the sensor may either be well placed or exposed to the sun, or under cover,... or may not be located at the exact same location as that of the building (for instance if we use a temperature from an open access database). In comparison, the temperature sensor of a climate station is necessarily protected from the solar radiation.

In order to illustrate the second point, we conducted a simple experiment at EDF R&D laboratory BESTLab: we measured a reference dry-bulb air temperature (sensor with 0.15°C accuracy, at a controlled climate station a few meters aways from BESTLab) on one hand, and two other dry-bulb air temperatures, from sensors located in the immediate vicinity of BESTLab's cells, that is in particular with a southern and a western orientation respectively (accuracy 0.1°C). A brief description of the characteristics of these sensors is given in Table 2.2 in Chapter 2. Figure 5.1 shows that there is indeed a bias in the measurements from the two unprotected sensors. Interestingly, this discrepancy between the two measurements occurs during the whole range of winter outdoor temperatures, from about −5°C to 20°C. Moreover, it appears that this bias may be up to 8°C during day time whereas it is much closer to zero at night. This confirms the influence of the solar radiation on an on-site but exposed temperature sensor. Hence, if the building model is identified according to the method described in Chapter 3 with the southern sensor instead of the reference sensor, the gap between the indoor and outdoor temperature would be less important than it actually is and one might expect for instance the heat loss coefficient UA to be overestimated to a certain extent.

A typical screenshot of the Weathermap by Netatmo, which maps the outdoor temperatures measured by the personal weather stations of Netatmo's clients, as in Figure 5.2 confirms the disparity of the measures in a restricted area, due to an average accuracy (0.3°C) and probably to mislocated sensors: the average temperature lies probably between 13°C and 14°C, but two adjacent sensors might give measurements as different as 13°C and 17°C or 15°C and 20°C.

In this chapter, we assume thus that the observed outdoor temperature is possibly biased, the noise-free outdoor temperature being actually hidden. This suggests that the outdoor temperature

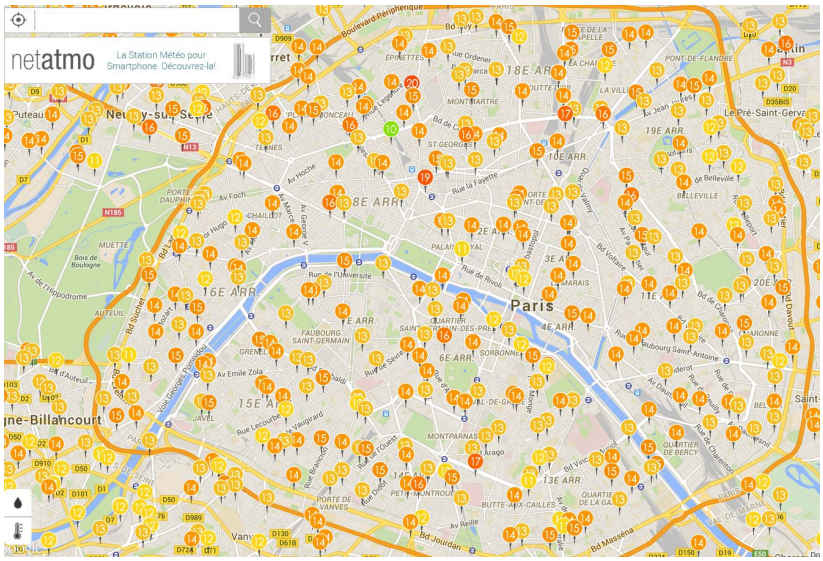

Figure 5.2: Screenshot of the measurements of the outdoor temperatures from the Weathermap by Netatmo in Paris.

may be seen as a stochastic process rather than a deterministic input, as it is only an approximation of the true temperature. Accordingly, the building system belongs here to the so-called class of *errors-in-variables* dynamic models [Söd07].

With these elements in mind, our purpose is to propose a new method for identifying dynamic thermal building models, that takes into account the uncertainty inherent to the outdoor temperature records, and investigate whether this uncertainty significantly affects the estimation of the physical parameters of the model. Assuming that the other inputs and the output of the R3C2 model are correctly measured, we show how to include a state-space representation of the outdoor temperature in the reference framework. The estimation procedure in this constrained context is based on the EM algorithm, similarly to Chapter 3.

5.2 State space representation

In order to account for the possible bias of the measurement of the outdoor temperature, we decide to model this measurement as a state-space system. Two modeling steps are therefore required in addition to the open loop R3C2 building model from Chapter 2 that explains the indoor temperature in terms of weather conditions and heating flux. We must first specify the dynamics of the hidden outdoor temperature, at time-step 10 minutes. Secondly, the bias model is contained in an observation equation. The global model is thus also a state space system, with noise-free inputs and the noisy outputs summarized in Table 5.1 (compare to Table 2.1).

5.2.1 Dynamics and observation of the outdoor temperature

First of all, at time-step $\delta = 10$ minutes and over a span of a few days only, we suppose that the actual outdoor temperature ϑ_o evolves according to a one-dimensional random walk. At discrete time $t := t\delta$, we have

$$\vartheta_{o,t} = \vartheta_{o,t-1} + \varepsilon_t, \tag{5.1}$$

Table 5.1: Inputs and outputs of the R3C2 representation with biased outdoor temperature.

Inputs	Outputs
solar gains (Q_s) heating flux (Q_h) internal gains (Q_{fg})	indoor temperature (T_i) outdoor temperature (T_o)

where ε_t is a Gaussian white noise process with constant variance σ_ε^2. This is the simplest dynamic model possible, which assumes that the current value $\vartheta_{o,t}$ of the outdoor temperature is in the neighborhood of the past value $\vartheta_{o,t-1}$, this neighborhood being characterized by σ_ε. In particular, we do not include any periodic component in the model, because such component is already included in the observation T_o.

The observed outdoor temperature T_o is a biased version of the hidden state ϑ_o, where the bias remains to be specified. In accordance with our dataset (BESTLab experiment), let us consider the case of a temperature sensor close to the building but not necessarily protected from the solar radiations. An example of the corresponding series is given in Figure 5.3. It illustrates in particular the fact that the difference between the two sensors does not depend on the actual outdoor temperature: high values of the bias may occur even for cold climates. This represents typically the situation where an energy utility has access to the measurements of the end-user's temperature sensor but does not know its precise location around the building. We suggest to represent the bias as another hidden state, denoted ϑ_b, with a random walk behaviour:

$$\vartheta_{b,t} = \vartheta_{b,t-1} + \epsilon_t, \tag{5.2a}$$

$$T_{o,t} = \vartheta_{o,t} + I_t \vartheta_{b,t} + \eta_t, \tag{5.2b}$$

where $T_{o,t}$ is the noisy observation of the "true" outdoor temperature $\vartheta_{o,t}$, ϵ_t a Gaussian white noise process with time-invariant variance σ_ϵ^2, η_t a Gaussian observation white noise process with time-invariant variance σ_η^2 and independent of ϵ_t and I_t is an indicator function, set to 0 at night and 1 during day time in order to account for the effect of the solar radiation, in accordance with our remark subsequent to Figure 5.1.

We note finally that the model (5.1)-(5.2) of the outdoor temperature does not introduce any additional parameter to the R3C2 network, keeping thus the dimensionality of the estimation problem to a low level. In particular, this model assumes that the switching times of δ_t are known.

5.2.2 Global model

The global model is a time-varying errors-in-variables linear state-space representation of the building. We choose to augment the initial state-space representation of the R3C2 model (2.1a)-(2.1b)

$$Q_r = C_r \frac{d\vartheta_i}{dt} + \frac{1}{R_i}(\vartheta_i - \vartheta_s) + \frac{1}{R_f}(\vartheta_i - \vartheta_o), \tag{5.3a}$$

$$Q_s = C_s \frac{d\vartheta_s}{dt} + \frac{1}{R_i}(\vartheta_s - \vartheta_i) + \frac{1}{R_o}(\vartheta_s - \vartheta_o), \tag{5.3b}$$

with the model given by eqs (5.1)-(5.2). Observe that ϑ_i, ϑ_s and ϑ_o are respectively the hidden indoor, structure and outdoor temperatures, whereas $Q_r := Q_h + Q_{fg}$ is the total heating power,

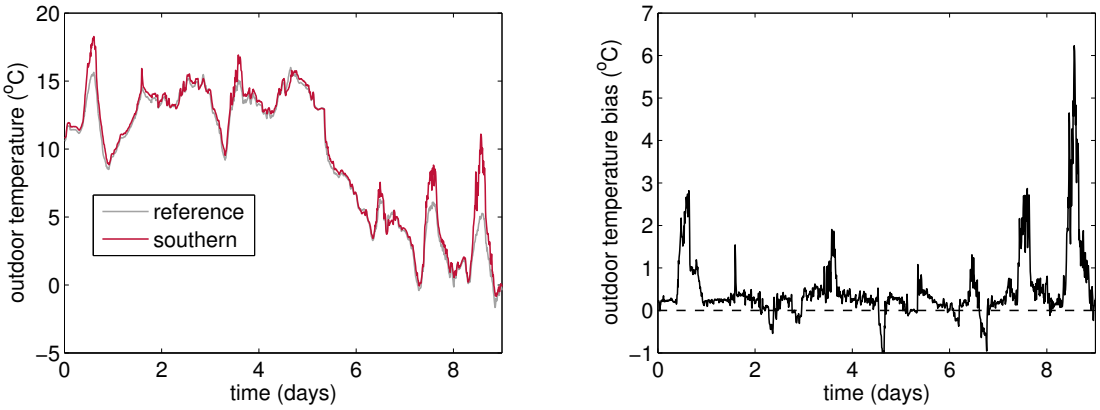

(a) (b)

Figure 5.3: (a) Example of a nine-day record of the outdoor temperature using a reference sensor under cover (ϑ_o, light grey) and a sensor with southern orientation and exposed to the sun (T_o, red). The data come from the BESTLab measurement campaign of November 2015. (b) Difference ϑ_b between the measures of T_o and ϑ_o during the same period of time.

wiht notations consistent with Table 5.1. This latter model is appended to the R3C2 network as follows. Let $U := \begin{pmatrix} Q_r & Q_s \end{pmatrix}^\intercal$ denote the measured input vector, $X := \begin{pmatrix} \vartheta_i & \vartheta_s & \vartheta_o & \vartheta_b \end{pmatrix}^\intercal$ the hidden state vector and $Y := \begin{pmatrix} T_i & T_o \end{pmatrix}^\intercal$ the noisy observation of $\begin{pmatrix} \vartheta_i & \vartheta_o \end{pmatrix}^\intercal$. The state equation is deduced from the discretization of (5.3) and from the random walk models (5.1) and (5.2a). At time $t := t\delta$, where δ is the sampling time, we have

$$X_t = A_\delta(\theta) X_{t-1} + B_\delta(\theta) U_t + V_t, \tag{5.4}$$

$$Y_t = C_t X_t + W_t, \tag{5.5}$$

where V_t is a model Gaussian white noise process with time-invariant covariance matrix Q_δ, W_t is an observation Gaussian white noise with time-invariant covariance matrix R_δ uncorrelated with V, $A_\delta(\theta) = \exp(\delta A(\theta))$, $B_\delta(\theta) \simeq \delta B(\theta)$ and

$$A(\theta) = \begin{pmatrix} -z_r(z_f + z_i) & z_r z_i & z_r z_f & 0 \\ z_s z_i & -z_s(z_o + z_i) & z_s z_o & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad B(\theta) = \begin{pmatrix} z_r & 0 \\ 0 & z_s \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \tag{5.6}$$

$$C_t = \begin{cases} M_1 \text{ if t is a diurnal index} \\ M_2 \text{ otherwise} \end{cases}, \quad M_1 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{pmatrix}, M_2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}. \tag{5.7}$$

5.3 Estimation of the parameters

5.3.1 EM for a time-varying linear state-space

We choose the EM algorithm in order to estimate the parameters of the augmented model, under the assumption that the initial state is normal. See [SS82] and Chapter 3 for a more detailed description of the algorithm in the case of linear Gaussian time-invariant state space models, and [WHHF16] for another approach of errors-in-variables models estimated with EM. The main difference with the reference case in Chapter 3 is that the model is now a time-varying system, to

which EM and the Kalman smoother involved in the procedure adapt easily. This is prensented e.g. [DMS14, Chapter 2]. Hence, the estimation procedure of the global model is essentially the same as the one described in Chapter 3. We merely focus on the estimation of the inverse of the resistors and capacitors, and are not interested in estimating the initial state, nor the noise covariances. The closed-form solution of the Maximization step is discarded, because the state matrices are structurally defined by θ. Instead, we use thus the Matlab routine `fminunc` with the trust-region algorithm, which requires a gradient computation of the EM cost. In order to keep the problem unconstrained, the five parameters of the R3C2 network are expressed as $\theta = \exp(\gamma)$ and the optimization is carried out with respect to γ.

5.3.2 EM for a switching linear state-space

In this section, we suggest to adapt the estimation procedure to a weaker version of the observation equation (5.2b) of the outdoor temperature. Indeed, it is assumed in (5.2) that the bias is necessarily null at night, with $I_t = 0$ at night. This hypothesis may be relaxed in the following manner. We suppose that the system has two possible configurations, or states, depending on the value of I_t in (5.2b). These two configurations are the states of a nonstationary *independent* process defined by the time-varying *a priori* probabilities

$$\pi_j(t) := \mathrm{P}(C_t = M_j),\ j = 1, 2, \tag{5.8}$$

with M_1 and M_2 given in (5.6). Hence, the system has a dynamic, stochastic switching behaviour, instead of the deterministic switching of the initial model.

Algorithm 5.1, reproduced from [SS91], details how to adapt the EM algorithm to this case. In particular, the Kalman filter necessary to solve the Expectation step is modified to account for the multiple possible states of the system. Standard notations, such as $x_k^l := \mathbb{E}[x_k|y_1, \ldots, y_l]$ and $P_k^l := \mathrm{cov}(x_k|y_1, \ldots, y_l)$ are used. It is based on the computation of the filtered probabilities

$$\pi_j(t|t) := \mathrm{P}(C_t = M_j|Y_1, \ldots, Y_t). \tag{5.18}$$

One obtains

$$\pi_j(t|t) = \frac{\pi_j(t) f_j(t|t-1)}{\pi_1(t) f_1(t|t-1) + \pi_2(t) f_2(t|t-1)}, \tag{5.19}$$

where $f_j(t|t-1)$ is the conditional density of Y_t given $C_t = M_j$ and the past observations $y_1, \ldots, y_{t-1}$. $f_j(t|t-1)$ is a mixture of normal distributions which has an exponentially increasing computational cost. Consequently, $f_j(t|t-1)$ is approximated by a single normal density, with mean vector $M_j x_t^{t-1}$ and covariance matrix $M_j P_t^{t-1} M_j^\intercal + R$ (see [SS91] and the corrections therein; this density is optimal in the sense of the Kullback-Leibler distance).

If the covariance matrix R_δ is estimated, Algorithm 5.1 is a "pseudo-EM" algorithm, since solving for R_δ involves the smoothed probabilities $\pi_j(t|\mathrm{T})$, where T is the total number of measurements, which are very expensive to compute. In such case, they are replaced by $\pi_j(t|t)$, hence the name "pseudo-EM". However, in this application and the numerical experiments in the next section, we choose not to estimate R_δ and the only smoothed probability that is required is $\pi_j(\mathrm{T}|\mathrm{T})$ in (5.16). This one can be computed exactly by the filtering equations (5.10).

Algorithm 5.1 EM for switching linear state-space

Inputs: θ_0, *a priori* probabilities $\pi_j(t)$, $j = 1, 2$, $t = 1, \ldots, \mathrm{T}$.

0. Initialize $\theta = \theta_0$ and set $k = 1$.

1. *E-step:* compute

$$\mathcal{Q}(\theta, \theta_k) := \mathbb{E}_{\theta_k} \left[\ln L_\theta(X, Y) | Y = y_1, \ldots, y_n \right] \tag{5.9}$$

based on (a) the filtering scheme

$$x_t^{t-1} = A_\delta x_{t-1}^{t-1} + B_\delta u_t, \tag{5.10}$$

$$P_t^{t-1} = A_\delta P_{t-1|t-1} A_\delta^\intercal + Q_\delta, \tag{5.11}$$

$$K_{tj} = P_{t|t-1} M_j^\intercal (M_j P_t^{t-1} M_j^\intercal + R_\delta)^{-1}, \tag{5.12}$$

$$x_t^t = x_t^{t-1} + \sum_{j=1}^{2} \pi_j(t|t) K_{tj} (y_t - M_j x_t^{t-1}), \tag{5.13}$$

$$P_t^t = \sum_{j=1}^{2} \pi_j(t|t) (I - K_{tj} M_j) P_t^{t-1}, \tag{5.14}$$

where the filtered probabilities are, with $f_j(t|t-1) := p(Y_t | C_t = M_j, Y_1, \ldots, Y_{j-1})$,

$$\pi_j(t|t) = \frac{\pi_j(t) f_j(t|t-1)}{\pi_1(t) f_1(t|t-1) + \pi_2(t) f_2(t|t-1)}, \; j = 1, 2, \tag{5.15}$$

(b) the standard Kalman smoother, and (c) the lag-one covariance smoother initialized with

$$P_{\mathrm{T},\mathrm{T}-1}^{\mathrm{T}} = \sum_{j=1}^{2} \pi_j(\mathrm{T}|\mathrm{T}) (I - K_{\mathrm{T}j} M_j) A_\delta P_{\mathrm{T}-1}^{\mathrm{T}-1}. \tag{5.16}$$

2. *M-step:* compute θ_{k+1} :

$$\theta_{k+1} = \arg\max_\theta \mathcal{Q}(\theta, \theta_k). \tag{5.17}$$

3. If convergence is reached, terminate, otherwise set $k \leftarrow k + 1$ and return to step 1.

Output: maximum likelihood estimator $\widehat{\theta}$.

5.4 Numerical illustration

5.4.1 Validation procedure

In order to avoid the convergence of the EM algorithm 5.1 towards a local minimum, $M = 50$ random initializations are drawn uniformly in $[0, 1]$ for every parameter (all true values are smaller than 1), except for z_s which is drawn between 0 and z_r. We obtain thus M estimators, and select the one with the greatest log-likelihood. Moreover, we require that the sub-state matrix defined corresponding to the system (5.3a)-(5.3b) is well-conditioned, since its inverse appears both in the discretization and for the resolution of the M-step. We also check that the Kalman filter is stable, which is needed for the asymptotic study of the estimates. Hence, if the eigenvalues of A_δ are outside the unit circle, we reinitialize the procedure. This is the same procedure as the one described in Chapter 3 for the estimation of the R3C2 model alone. Besides, for a given dataset, three models are estimated:

1. M0 is the naive estimation of the R3C2 network uniquely, with the recorded outdoor temperature as a deterministic input; the estimation procedure is that of Chapter 3;
2. M1 is the R3C2 network combined with eqs. (5.1)-(5.2), estimated with the standard EM algorithm (deterministic switching times, see Section 5.3.1)
3. M2 is the R3C2 network combined with eqs. (5.1)-(5.2), estimated with Algorithm 5.1 (probabilistic switching).

In the latter case, the probabilities $\pi_j(t)$ are set to the uninformative prior $\pi_1(t) = \pi_2(t) = 0.5$, for all time instant t. All three models are estimated successively with three outdoor temperatures: a reference outdoor temperature as well as the measurements of two sensors with a western and southern orientation, respectively (see Chapter 2). The length of the learning data set is two weeks. This procedure is iterated 25 times, each simulation with its own set of true parameters.

5.4.2 Smoothed outdoor temperature

We check first that the smoothed outdoor temperature $\widehat{\vartheta}_o$, estimated once Algorithm 5.1 has converged, is consistent with the observation T_o. In average over the N simulations, the Root Mean Square Error (RMSE) is significantly reduced for the southern sensor, from 1.2°C to 0.7°C. This error is more slightly reduced for the western sensor, from 0.6°C to 0.5°C. The bias generated on the true data is also reasonably low (0.4°C). These results tend to indicate that the chosen model (5.1)-(5.2) is relevant. This is confirmed graphically by inspecting the histograms of the relative errors, in Figure 5.4. In particular, the extreme values have been removed, and the errors are much more centered.

5.4.3 Estimation of the R3C2 model

The results are summarized in Table 5.2 and Figure 5.5. For a given parameter θ, the computed error is $(\theta^\star - \widehat{\theta})/\theta^\star$, where $\theta^\star$ and $\widehat{\theta}$ denote respectively the true and estimated value. We observe first that the heat loss coefficient UA is robustly estimated regardless of the estimation method, M1 and M2 yielding a slight improvement in average (see Table 5.2). The estimated UA is most often larger than the corresponding true value (see Figure 5.5(a)). This is consistent with the physical interpretation of UA. Indeed, UA is estimated with warmer temperatures than the true excitation of the building. Since the input corresponding to the heating load is unchanged, this means that the identified building is less efficient, and thus UA increases.

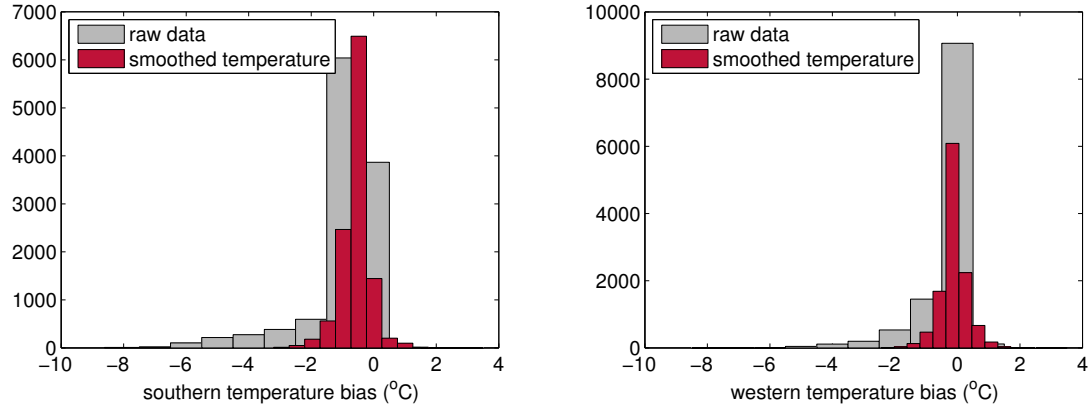

(a) (b)

Figure 5.4: Histograms of the temperature difference between the reference and (a) southern (b) western sensors. The darker bars correspond to the raw data, the lighter to the smoothed temperature estimated by Algorithm 5.1 and averaged over the 25 data sets.

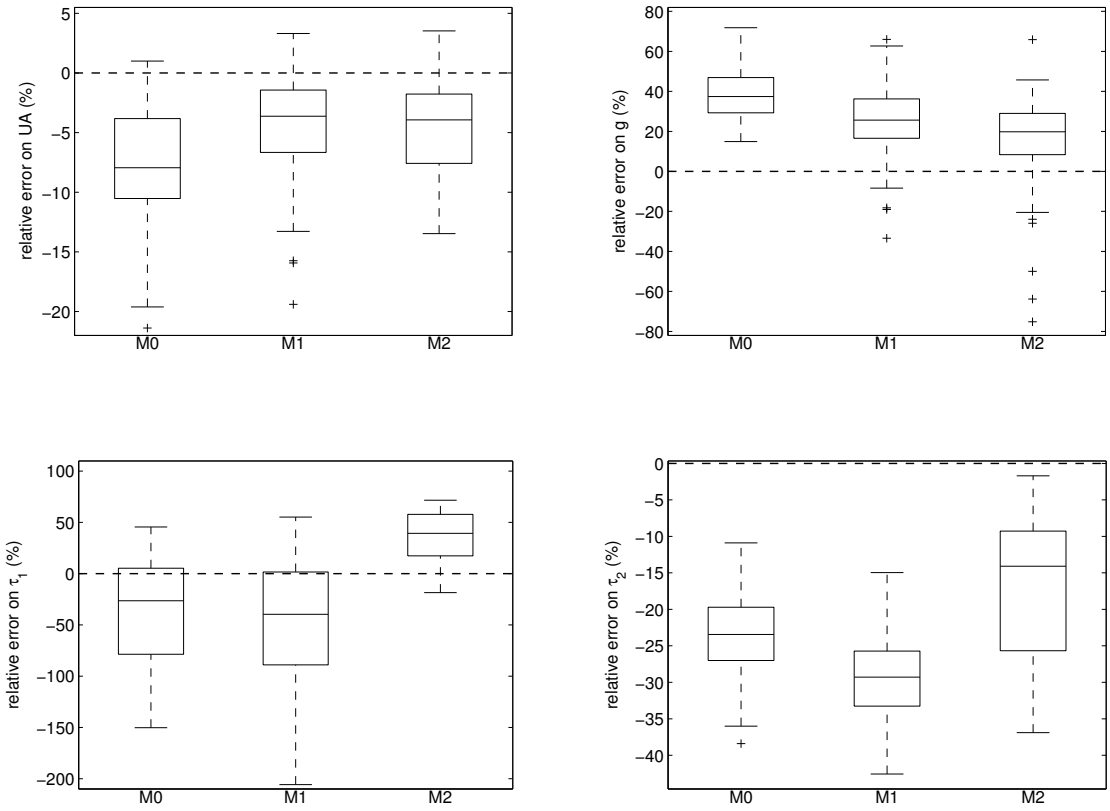

Figure 5.5:
Boxplots of the 25 estimation errors of the heat loss coefficient UA (top left corner), the solar transmittance g (top right corner), the slow (bottom left corner) and fast (bottom right corner) time constants, against the estimated model.

Table 5.2: Quadratic mean of the estimation errors (%) of the physical parameters of the R3C2 network for the three models M0, M1, M2. Average over 25 data sets.

Parameter	**M0**	**M1**	**M2**
UA	9.4	6.4	6.0
g	40.2	31.6	26.8
τ_1	109	100	44.3
τ_2	24.4	29.8	19.1

The estimation of the solar transmittance g illustrates the benefits of the errors-in-variables models M1 and M2 over the noise-free approach M0. Hence, these models are capable of reproducing the static thermal behaviour of the building more precisely than the noise-free approach M0. The dynamic behaviour is more difficult to estimate, which is in agreement with the numerical study in Chapter 3. This is partly due to the non-informative temperature setpoint used to generate the data which switches only twice a day, similarly to the real exploitation of a building. It appears nevertheless that the model M2, with stochastic switching, is much more efficient than the other two. In particular, the slower time constant is more accurately estimated, with a lower variance, whereas the faster time constant is rather correctly estimated in average, but the values are more dispersed.

5.5 Summary

In the context of smart buildings, it is reasonable to assume that an energy utility has access to a measurement of the outdoor temperature. However, we have questioned in this chapter the standard assumption of a noise-free measurement of this quantity. In practice, this measurement is indeed possibly biased, because the location of the sensor is unknown and up to the end-user: the sensor may for instance be left exposed to the solar radiation. Assuming that every other input and output of the building model is well measured, this chapter aimed at handling this uncertainty on the measurment of the outdoor temperature. To address this issue, we suggested to represent the possible outdoor temperature bias with an empirical discrete linear state-space model, without adding new parameters. A key feature of this model is that the temperature bias is accounted for by term which switches either deterministically or stochastically. When appending this model to the R3C2 model, this forms a global time-varying linear state-space model, which may be learned by the Expectation-Maximization algorithm with Kalman smoothing.

The numerical evaluation on simulated data shows the ability of the method to reduce the bias on the outdoor temperature. Moreover, the suggested algorithms improve the naive estimation of a R3C2 model with no bias correction for the outdoor temperature. This chapter shows therefore that the R3C2 model, and more generally and equivalent electrical network, may easily be transformed into an errors-in-variables representation that accounts for the uncertainty.

Chapter 6

Learning the dynamics of the unobserved solar gains

In the Introduction to this thesis, we took the example of the solar flux to motivate our work and emphasize the practical constraints that degrade the observation of the building. The solar radiation Q_s being measured by an expensive specific sensor, the pyranometer, it seems therefore unreasonable to assume that such a measurement would be available for any building. For instance, the pyranometer is not included in the new generation of connected weather stations. Besides, the naive estimation of the open-loop R3C2 network with $Q_s \equiv 0$ is not advocated, because Q_s is a non-negligible gain. It is all the more true in the emerging context of the smart and low-energy buildings, which aim at maximizing such free gains for improved energy efficiency performances [CDRR12]. We suggest therefore two ways of overcoming this difficulty:

- either estimating the solar radiation prior to the identification of the R3C2 network, typically from other meteorological data, and including then this estimation and its uncertainty in the identification process;
- or modelling the solar radiation as a hidden quantity estimated jointly with the other states and parameters of the R3C2 network.

The first approach is very challenging, because most statistical models of the solar radiation are based on supervised learning algorithms. That is, a dataset with records of the solar radiation must be available at every location of interest, which is precisely not our case. For instance, 33 out of the 35 references cited in the review [TLH+17, Table 9] of solar irradiance prediction models require a measurement of the direct or global solar irradiance as an input, while the other two require a cloud cover record. In Appendix A, we suggest a new method, which constructs a statistical spatial model of the cloud cover, and is based on open access records of this quantity at given, fixed, locations (for example, in France). The solar radiation at any location can in turn be estimated from the cloud cover index by means of a simple physical model. This method is promising because it requires no ground instrumentation, but still lacks accuracy and the studied standard spatial models are largely perfectible.

On the other hand, the second approach is similar to the idea developed in Chapter 5, and it is the focus of this chapter. The main difference with Chapter 5 is that we do not have access to a biased measure of the quantity of interest: the system is not an errors-in-variables model. Nevertheless, Chapter 5 suggests that there is a link between the solar radiation, the reference air

Table 6.1: Inputs and outputs of the R3C2 representation with undirectly observed solar gains.

Inputs	Outputs
internal gains heating flux clear sky radiation outdoor temperature	indoor temperature outdoor temperature bias

temperature and a biased measure of the air temperature from a sensor exposed to the sun, the reference temperature being estimated from the other two. This leads naturally to the following question: can we estimate the solar radiation from the measurements, respectively biased and unbiased, of the outdoor temperature? The underlying idea is that larger differences between these two measurements are caused by larger values of the solar radiation. We will investigate the question by assuming throughout the chapter that these two measurements of the outdoor temperature, as well as the indoor temperature and the internal gains, are measured at the location of interest. Although it requires an additional temperature sensor compared to the standard setting of Chapter 3, the associated cost is negligible compared to the pyranometer's one, since such sensors are widespread and much more affordable now, especially in the new context of the smart buildings and the Internet of Things. This limits thus the practical constraints.

The rest of this chapter is organized as follows. Section 6.1 describes the chosen state-space representation of the solar flux. In particular, this model accounts for the variability of the weather conditions (cloud movements, etc.) by allowing for occasional high variance increments of the hidden state. The magnitude of the high values and the times at which they occur are unknown and estimated with an ℓ_1-regularized maximum likelihood approach, and this model is then included into the R3C2 framework. The resulting model is nonlinear with a linear substructure, which calls for a new smoothing scheme within the EM algorithm. An adapted solution, called Rao-Blackwellised particle smoothing is therefore presented in Section 6.2. Finally, we conclude the chapter with a numerical evaluation of the model in Section 6.3.

6.1 Nonlinear modelling

The modelling step is two-fold. First of all, the open loop thermal building model is the R3C2 network described in Chapter 2. The second step, which we focus on in this section, entails a physical representation of the solar radiation, written under state space form. This model is then appended to the R3C2 network, forming a global state space model whose inputs and outputs are summarized in Table 6.1. This is to be compared to the standard setting described in Table 2.1.

6.1.1 Solar radiation model

Rather than modelling directly the dynamics of the solar flux, we suggest first to use a physical model. Indeed, there exist simple physical models, depending only on time and geographical coordinates, that estimate the solar radiation. These models are accurate enough provided that they are corrected by a measure of the cloud cover. In this work, we consider Perrin de Brichambaut's model in [PdBV82, RW99]:

$$G_{h,clr} = \kappa G_0 (\cos \theta_z)^{1.15} \tag{6.1}$$

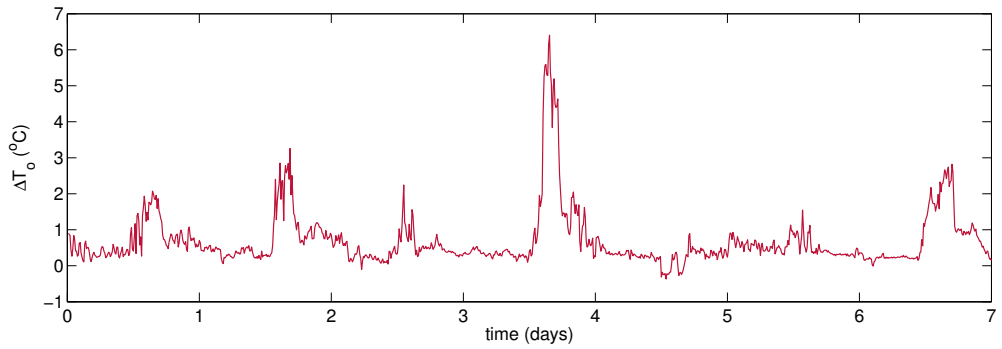

Figure 6.1: A one-week record of ΔT_o, the difference of outdoor temperatures measured respectively by a sensor exposed to the sun and second sensor under cover (BESTLab, November 2015).

where $G_{h,clr}$ is the global horizontal solar flux under clear sky conditions, κ is a constant set in the literature to $\kappa = 0.81$ [RW99], θ_z is the azimuth angle (that depends only on the latitude, longitude and solar time) and G_0 the extraterrestrial solar constant (theoretical radiation at the top of the atmosphere, depends on the day number). [DB13, Chapter 1] gives the formulas to compute these quantities. In [RW99], it is shown that this model is a first reasonable approximation, with errors decreasing when the Linke turbidity decreases. This means that the model works better for a dry and clean atmosphere. This model is corrected by a measurement $N \in [0, 1]$ of the cloud cover according to

$$G_{h,cld} = G_{h,clr}(1 - aN^b). \tag{6.2}$$

If the sky is completely clear ($N = 0$), 100% of the initial model $G_{h,clr}$ is kept, whereas a fraction $(1 - a)$ remains even if the sky is totally overcast ($N = 1$). In the sequel, we presuppose that a and b are known, with $a = 0.67$ and $b = 1.08$. Those values were obtained from the kriging model of the cloud cover detailed in Appendix A.

Hence, we consider that the latent quantity to be estimated is the cloud cover index N, instead of the solar flux. The cloud cover model is formulated as a discrete-time state-space model, with observation and dynamic equations to specify. First of all, we cannot assume that the cloud cover is directly measured at the location of interest. Instead, we suppose that a good approximation is given by the difference ΔT_o between two measures of the outdoor temperature, namely a measure from a biased sensor, in the sense that it is exposed to the sun, and a second measure from a reference sensor, under cover. The latter sensor is the usual sensor measuring the outdoor air temperature. This requires therefore two distinct outdoor temperature sensors. A one-week record of this outdoor temperature bias is plotted in Figure 6.1. The underlying idea is that smaller values of ΔT_o are caused by higher values of the cloud cover since in such case the sensors under cover and exposed to the sun give similar outdoor temperatures. This echoes back to the model in Chapter 5, where the reference outdoor temperature is reconstructed from a biased measure as well as the solar radiation. The relation between the cloud cover and ΔT_o is modelled as follows:

$$N = \phi(\Delta T_o), \qquad \text{where } \phi : x \in \mathbb{R} \mapsto 1 - \frac{1}{1 + \exp(-\alpha(x - x_0))} \in (0, 1), \tag{6.3}$$

with $x_0 = (x_1 + x_2)/2$, $\alpha = 4/(x_2 - x_1)$, where x_1 and x_2 are some parameters to estimate. In this model, clearest skies correspond to $\Delta T_o > x_2$ whereas overcast skies on the other hand correspond

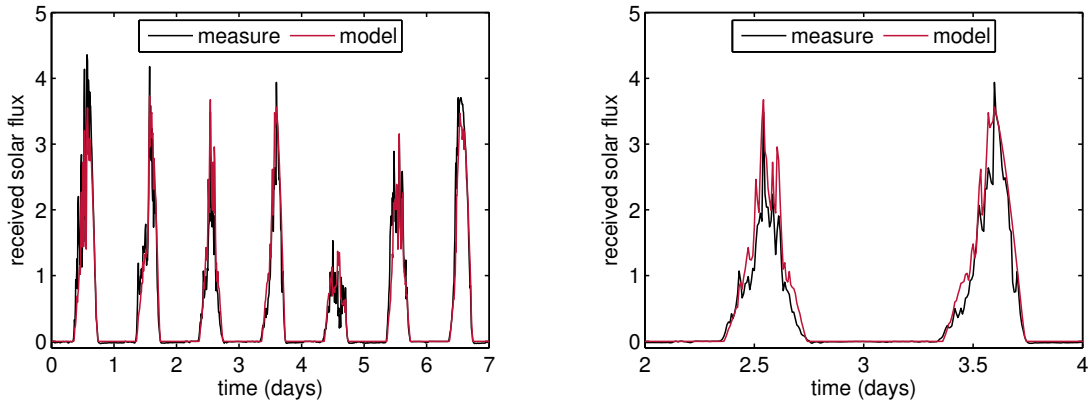

(a) (b)

Figure 6.2: (a) Estimated solar flux from the measured outdoor temperature bias ΔT_o and model (6.2)-(6.3). (b) Two-day zoom.

to $\Delta T_o < x_1$. (6.3) is inspired by the estimation of the cloud cover from the clearness index data (the ratio of the received global solar radiation during a given period to the extraterrestrial global horizontal radiation for this period - we do not have access to it) in [FXM+09]. In this sense, the bias ΔT_o may be seen as a surrogate clearness index. It remains to select the thresholds x_1 and x_2 (in °C). We used the data from BESTLab to minimize the mean square error of the predicted solar flux (using (6.2)-(6.3)). A cross-validation procedure with a one-week BESTLab record data from November 2015 yields for instance $x_1 = 0.44\,^{\circ}\mathrm{C}$ and $x_2 = 1.77\,^{\circ}\mathrm{C}$. An example in Figure 6.2 shows that this model provides a sensible guess of the received solar radiation, able to catch the periods of high and low intensities.

We derive now the discrete-time state-space representation of the cloud cover. Let us assume first that ΔT_o is the noisy observation of a hidden state η:

$$\Delta T_{o,t} = \eta_t + w_{\eta,t}, \tag{6.4}$$

where w_η is a Gaussian white noise with constant variance herafter denoted r_η^2. The hidden cloud cover is therefore given by $N_t = \phi(\eta_t)$, with ϕ as in (6.3), and the estimated solar radiation is $Q_s = G_{h,clr}(1 - a\phi(\eta)^b)$.

The last step consists in specifying the dynamics of η. At a time-step of 10 minutes, we may suppose that this dynamics is accurately represented by a random walk. However, the high variability of the weather (cloud movements, etc) produces some jumps in the series, as exemplified by Figure 6.1. Hence, we choose to represent the dynamics of η by a discrete-time random walk with time-varying variance:

$$\eta_t = \eta_{t-1} + (q_\eta + \sigma_t)v_{\eta,t} \tag{6.5}$$

where v_η is a Gaussian white noise with unitary variance. The constant q_η is the minimum standard deviation of the additive noise, while σ_t is either null, or positive whenever a change in the variability occurs. We suppose that most σ_t's are null, whereas a few positive values account for the jumps in the dynamics of η. Accordingly, the series is a random-walk with increments' variances all equal

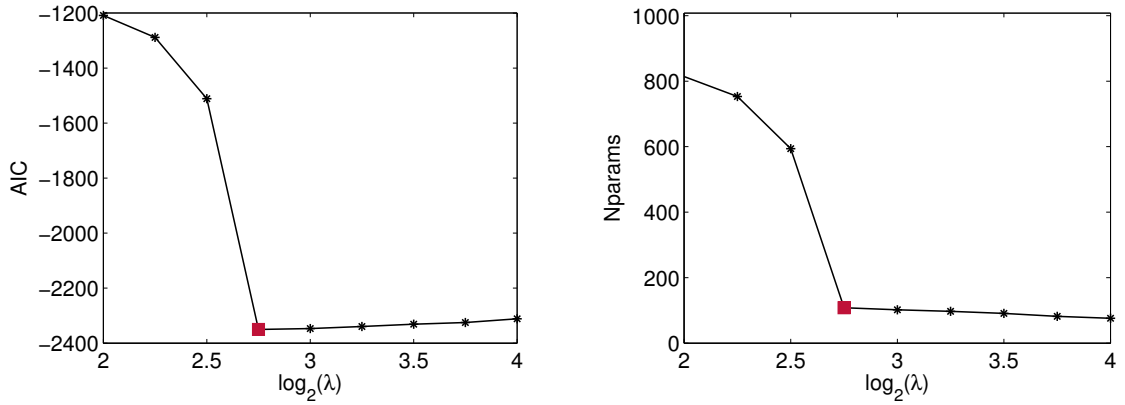

(a) (b)

Figure 6.3: Example of a grid search for the best penalty coefficient λ, with (a) the AIC and (b) the number of parameters $n(\widehat{\beta})$ displayed. The red square denotes the selected coefficient $\lambda = 2^{2.75} \simeq 6.7$.

to the minimal variance q_η^2 except for a few of them which are higher, at unknown times and with unknown values. Estimating the sequence of σ_t's is treated as a variable selection problem, similarly to the approach in [LLH08] (which deals with additive jumps in the mean). The parameters $\beta := \{q_\eta, r_\eta, \sigma_1, \ldots, \sigma_{\mathrm{T}}\}$ are found by minimizing a ℓ_1-penalized log-likelihood criterion:

$$\widehat{\beta}(\lambda) := \arg\min_{\beta} \left\{ -2 \log L_\beta(\Delta T_{o,1}, \ldots, \Delta T_{o,\mathrm{T}}) + \lambda \sum_{t=1}^{\mathrm{T}} |\sigma_t| \right\}, \tag{6.6}$$

where L_β, the likelihood of the linear Gaussian state space model (6.4)-(6.5), is a by-product of the Kalman filter equations. For a given λ, there exist already many optimization algorithms solving the problem (6.6). In the sequel, we have used Mark Schmidt's toolbox for Matlab, with algorithms detailed in [SFR07, Sch10]. λ is a trade-off coefficient, balancing the predictive power of the model and the total number of parameters, with larger values of λ penalizing more strongly the non-zero σ_t's. λ is selected by testing different values and retaining the one minimizing the Akaike Information Criterion (AIC):

$$AIC := 2n(\widehat{\beta}) - 2\log(L_{\widehat{\beta}}), \tag{6.7}$$

where $n(\widehat{\beta})$ is the number of non-zero parameters in $\widehat{\beta}$ and $L_{\widehat{\beta}}$ is the maximized value of the likelihood function $L_\beta(\Delta T_{o,1}, \ldots, \Delta T_{o,\mathrm{T}})$. More precisely, λ is generated as a power of 2 during a two-step grid search, where the first grid is $2^{-6}, 2^{-5}, \ldots, 2^5, 2^6$ and the second grid is centered around the best node of the former grid, with a step-size of $1/4$. In practice, we choose fixed values $q_\eta = 0.15$ and $r_\eta = 0.15$, and estimate only the σ_t's. An example is provided in Figures 6.3 and 6.4, with selected penalty coefficient $\lambda = 2^{2.75} \simeq 6.7$. In this example, there are 89 out of 1008 (8.8%) identified high variance increments such that $\sigma_t > 0$. Figures 6.3(b) confirms that smaller values of λ result in a very large number of parameters $n(\widehat{\beta})$. On the other hand, larger values of λ do not yield a significant reduction of $n(\widehat{\beta})$. Note that the problem (6.6) is unconstrained, and in particular the estimated σ_t's could be negative. Yet, by fixing in practice q_η to a sufficiently small value, the estimated σ_t's are always positive (in Figure 6.4 as well as in any of the other simulations in this chapter).

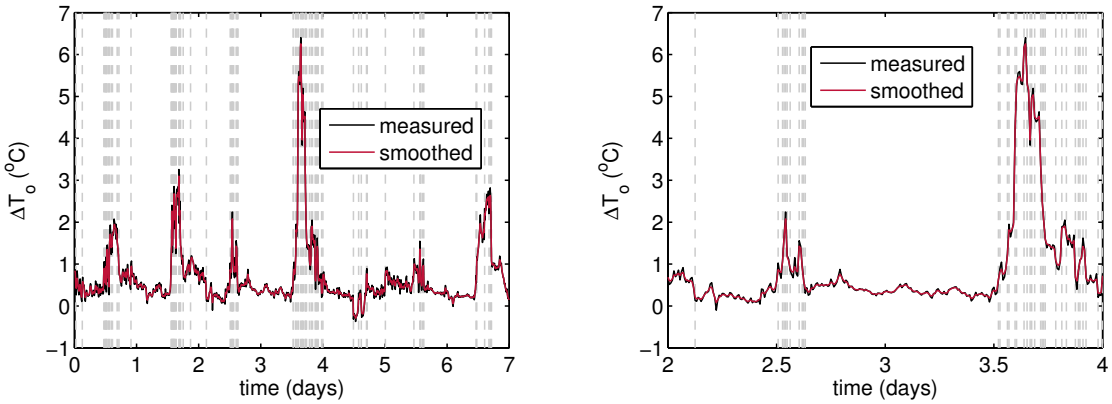

Figure 6.4: (a) One-week and (b) two-day zoom of the smoothed dynamics of the temperature bias process. The vertical lines locate the estimated changing points of the variance. Data from November 2015, BESTLab.

6.1.2 Global state-space model

The building thermal model is made of two blocks, namely the building dynamics given by the R3C2 network presented in Chapter 2, and the solar radiation model given by (6.4)-(6.5). Let $U := \begin{pmatrix}\vartheta_o & Q_r & G_{h,clr}\end{pmatrix}^\intercal$ be the observed input vector, $Y := \begin{pmatrix}T_i & \Delta T_o\end{pmatrix}^\intercal$ the noisy observation of $\begin{pmatrix}\vartheta_i & \eta\end{pmatrix}$, and $X := \begin{pmatrix}\vartheta_i & \vartheta_s & \eta\end{pmatrix}^\intercal$ the hidden state vector. Moreover, we split the state vector into $\xi := \eta$ and $\alpha := \begin{pmatrix}\vartheta_i & \vartheta_s\end{pmatrix}^\intercal$. Indeed, while the global model in X is nonlinear because of the solar radiation model, we have that, conditionally on ξ, the dynamics of α is linear: conditionally on the solar radiation model, the thermal building model, that is the R3C2 network, is a linear model. This conditional linearity will be exploited in the estimation process.

The modelling hypotheses result thus in a discrete-time conditionally linear Gaussian state-space (CLGSS) model, with for $t = 1, \dots, \mathrm{T}$:

$$\xi_t = f_\xi(\xi_{t-1}) + v_{\xi,t}, \tag{6.8a}$$
$$\alpha_t = f_\alpha(\xi_{t-1}, u_t, \theta) + A_\alpha(\theta)\alpha_{t-1} + v_{\alpha,t}, \tag{6.8b}$$
$$Y_t = h(\xi_t) + C\alpha_t + w_t, \tag{6.8c}$$

where V_ξ, V_α and W are pairwise-independent Gaussian white noises with respective covariances $(q_\eta + \sigma_t)^2$ (time-varying), Q_α and R (time independent). Besides, $\theta = \begin{pmatrix}z_r & z_s & z_f & z_o & z_i\end{pmatrix}^\intercal$ denotes the parameter vector of the R3C2 network, and

$$f_\xi : x \in \mathbb{R} \mapsto x \in \mathbb{R}, \tag{6.9a}$$

$$f_\alpha : (x, u, \theta) \in \mathbb{R} \times \mathbb{R}^3 \times \mathbb{R}^5 \mapsto \begin{pmatrix}\theta(1)\theta(3)u(1) + \theta(1)u(2) \\ \theta(2)\theta(4)u(1) + \theta(2)u(3)(1 - a\phi(x)^b)\end{pmatrix} \in \mathbb{R}^2, \tag{6.9b}$$

$$h : x \in \mathbb{R} \mapsto \begin{pmatrix}0 \\ x\end{pmatrix} \in \mathbb{R}^2, \tag{6.9c}$$

$$A_\alpha(\theta) = \begin{bmatrix} 1 - z_r(z_f + z_i) & z_r z_i \\ z_s z_i & 1 - z_s(z_o + z_i) \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \tag{6.9d}$$

where f_α and A_α are obtained from a first-order discretization of Kirchhoff's equations (2.1) with stepsize $\delta = 1$ omitted for sake of clarity. Finally, the initial state x_1 is assumed Gaussian with

$$x_1 \sim \mathcal{N}\left(\begin{bmatrix} \xi_{1|0} & 0 \\ 0 & \alpha_{1|0} \end{bmatrix}, \Sigma_1 \right). \tag{6.10}$$

$\xi_{1|0}$, $\alpha_{1|0}$ and Σ_1 are unknown but not learned in the sequel.

For notational convenience, an alternative notation of the state equation (6.8a)-(6.8b) is

$$x_t = f(\xi_{t-1}, u_t, \theta) + A(\theta)\alpha_{t-1} + v_t, \tag{6.11}$$

with $v_t \sim \mathcal{N}(0, Q_t)$, where (with $\sigma_{\xi,t} := q_\eta + \sigma_t$ from (6.5))

$$Q_t := \begin{bmatrix} \sigma^2_{\xi,t} & 0 \\ 0 & Q_\alpha \end{bmatrix}, \quad f(\xi_{t-1}, u_t, \theta) := \begin{bmatrix} f_\xi(\xi_{t-1}) \\ f_\alpha(\xi_{t-1}, u_t, \theta) \end{bmatrix}, \quad A(\theta) := \begin{bmatrix} 0 \\ A_\alpha(\theta) \end{bmatrix}. \tag{6.12}$$

In this work, we suppose that the parameters of the solar flux model, that is mainly the function ϕ and the high variance increments σ_t, are learned beforehand, such that the global model is parameterized by θ only.

The model (6.8) is a special case of the so-called class of mixed linear/nonlinear state space models, whose inference via a particle filter based version of the EM algorithm has been studied in some recent publications, such as [LS10a]. We reproduce this method in the next section. From a statistical point of view, the key point of the model (6.8) is that unlike the more general framework of [LS10a], the dynamics of the nonlinear state ξ in (6.8a) does not depend on the conditionally linear state α.

6.2 Rao-Blackwellised particle smoother EM algorithm

6.2.1 Problem formulation

The parameters θ of the R3C2 network are estimated via a maximum likelihood procedure, relying on the EM algorithm. EM was previously introduced in Algorithm 3.1. It necessitates to compute the quantity

$$\mathcal{Q}(\theta, \theta') = \mathbb{E}_{\theta'} \left[L_\theta(X_{1:\mathrm{T}}, Y_{1:\mathrm{T}}) | Y_{1:\mathrm{T}} \right]. \tag{6.13}$$

Given the Markovian properties of the model, we have

$$L_\theta(X_{1:\mathrm{T}}, Y_{1:\mathrm{T}}) = \log p_\theta(Y_{1:\mathrm{T}}|X_{1:\mathrm{T}}) + \log p_\theta(X_{1:\mathrm{T}}) \tag{6.14}$$

$$= \log p_\theta(X_1) + \sum_{t=1}^{\mathrm{T}-1} \log p_\theta(X_{t+1}|X_t) + \sum_{t=1}^{\mathrm{T}} \log p_\theta(Y_t|X_t). \tag{6.15}$$

Hence,

$$\mathcal{Q}(\theta, \theta') = \mathcal{I}_1(\theta, \theta') + \mathcal{I}_2(\theta, \theta') + \mathcal{I}_3(\theta, \theta'), \tag{6.16a}$$

where

$$\mathcal{I}_1(\theta, \theta') = \mathbb{E}\left[\log p_\theta(X_1) \middle| Y_{1:\mathrm{T}}\right], \tag{6.16b}$$

$$\mathcal{I}_2(\theta, \theta') = \sum_{t=1}^{\mathrm{T}-1} \mathbb{E}\left[\log p_\theta(X_{t+1}|X_t) \middle| Y_{1:\mathrm{T}}\right], \tag{6.16c}$$

$$\mathcal{I}_3(\theta, \theta') = \sum_{t=1}^{\mathrm{T}} \mathbb{E}\left[\log p_\theta(Y_t|X_t) \middle| Y_{1:\mathrm{T}}\right]. \tag{6.16d}$$

The crucial step is therefore to compute the expectations in (6.16). This requires in particular the computation of the smoothing densities, such as $p_\theta(X_t|Y_{1:\mathrm{T}})$. This is achieved in several possible ways, depending on the model structure:

- in the case of a linear Gaussian state-space, a closed-formed solution exists and is given by the Kalman smoother, see Chapter 3 and Algorithm 3.2;
- in the nonlinear case however, the analytical computation is infeasible, and numerical approximations are needed. An in-depth investigation of this task is proposed in [DMS14]. A standard solution is provided by the class of Sequential Monte Carlo (SMC) methods, which computes Monte Carlo based empirical averages of the expectations in (6.16).

SMC methods approximate a sequence of target distributions by a set of weighted particles that are updated recursively, hence the names *particle filters* and *particle smoothing*. The idea has been introduced first by [GSS93], and has been extensively studied since. See e.g. [DGA00, GDW04, CGM07, DJ09] for Monte Carlo smoothing based on particle filters, and [ADST04, KDSM09, WSN08] for its use for parameter identification of a nonlinear state-space model. More specifically, a maximum likelihood estimation of a general nonlinear system based on the EM algorithm and particle filter based approximations is described in [SWN11].

Here, the conditional linearity of the model (6.8) can be exploited in order to improve the particle filtering and smoothing approximations. The idea is that the computations involving the conditional linear state α can be carried out analytically, thereby reducing the dimension of the nonlinear state space. This filtering technique is presented in Section 6.2.2, while its inclusion within the EM algorithm is explained in Section 6.2.3.

6.2.2 Rao-Blackwellised particle filtering and smoothing

Background on particle filtering and smoothing

This brief overview of the particle filtering and smoothing methods is based on [LBS+16]. For comprehensive tutorials, see e.g. [GDW04, DJ09, Gus10]. Standard references include [DGA00, CGM07, DGM+11]. The principle is the following. Given a general nonlinear state-space model of the form

$$x_{t+1} \sim p(x_{t+1}|x_t), \tag{6.17a}$$

$$y_t \sim p(y_t|x_t), \tag{6.17b}$$

with arbitrary probability density function p, latent state $x_t \in \mathbb{R}^{n_x}$ and observation $y_t \in \mathbb{R}^{n_y}$, a particle filter (PF) approximates the joint filtering distribution $p(x_{1:t}|y_{1:t})$ by a set of M_1 weighted

particles $\{x_{1:t}^j, w_t^j\}_{j=1}^{M_1}$. Consequently, the joint filtering distribution is represented by the point-mass approximation

$$\widehat{p}^{M_1}(x_{1:t}|y_{1:t}) := \sum_{j=1}^{M_1} w_t^j \delta_{x_{1:t}^j}(x_{1:t}), \tag{6.18}$$

where $\delta_x(\,\cdot\,)$ is the Dirac distribution at x. In order to compute this weighted system, the target distribution is factorized using conditional probabilities and the Markovian properties of the model (6.17) into

$$p(x_{1:t}|y_{1:t}) \propto p(y_t|x_{1:t}, y_{1:t-1})p(x_{1:t}|y_{1:t-1}) = p(y_t|x_t)p(x_t|x_{t-1})p(x_{1:t-1}|y_{1:t-1}). \tag{6.19}$$

Since we cannot sample directly from this target, SMC assumes instead that M_1 particles are drawn from a known instrumental (importance) distribution $q(x_{1:t}|y_{1:t})$. Assuming that the importance distribution factorizes into $q(x_{1:t}|y_{1:t}) = q(x_{1:t-1}|y_{1:t-1})q(x_t|x_{1:t-1}, y_t)$, the weights w_t^j and the particles $x_{1:t}^j$, $j = 1, \dots, M_1$, are then generated sequentially. $x_{1:t-1}^j$ is sampled from the previous distribution $\widehat{p}^{M_1}(x_{1:t-1}|y_{1:t-1})$ (resampling step), while x_t^j is sampled from the proposal $q(x_t|x_{1:t-1}, y_t)$. The weights w_t^j are computed such as to account for the discrepancy between the proposal and target distributions:

$$w_t^j := \frac{p(x_{1:t}^j|y_{1:t})}{q(x_{1:t}^j|y_{1:t})} \propto w_{t-1}^j \frac{p(y_t|x_t^j)p(x_t^j|x_{t-1}^j)}{q(x_t^j|x_{1:t-1}^j, y_t)}, \tag{6.20}$$

where we have used (6.19). Due to the systematic resampling step, we have at time t, $w_{t-1}^j = 1/M_1$, meaning that it can be discarded in (6.20). The PF is thus a sequential importance sampling scheme, with resampling. The simplest choice for q is $q(x_t^j|x_{1:t-1}^j, y_t) := p(x_t^j|x_{t-1}^j)$, which leads to the so-called bootstrap filter [GSS93]. It is a non-optimal choice, because the information from y_t is lost, yet it is commonly used in practice since all the densities are easy to compute and the weights simplify to $w_t^j = p(y_t|x_t^j)$. Finally, note that a representation of the marginal filtering density $p(x_t|y_{1:t})$ is obtained straightforwardly by marginalization of (6.18), that is by discarding the past trajectories $x_{1:t-1}^j$.

The smoothing step approximates the distributions $p(x_{1:t}|y_{1:\mathrm{T}})$ for all $t = 1, \dots, \mathrm{T}$. The first idea would be to marginalize the joint filtering distribution $p(x_{1:\mathrm{T}}|y_{1:\mathrm{T}})$. However, this methods suffers from path degeneracy: $p(x_{1:t}|y_{1:\mathrm{T}})$ would be approximated by a single particle for $t \ll \mathrm{T}$, due to the resampling step in the PF [DJ09]. Several strategies have been designed in order to mitigate this effect, see e.g. [LS13]. One such strategy is the forward filtering/backward simulation (FFBSi) algorithm reported in [GDW04], which generates states successively in the reverse-time direction, conditioning upon future states. It is based on the factorization

$$p(x_{1:\mathrm{T}}|y_{1:\mathrm{T}}) = p(x_{\mathrm{T}}|y_{1:\mathrm{T}}) \prod_{t=1}^{\mathrm{T}-1} p(x_t|x_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) \tag{6.21}$$

of the joint smoothing density. The Markovian properties of the model (6.17) and Bayes' theorem yield

$$p(x_t|x_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = p(x_t|x_{t+1}, y_{1:t}) \propto p(x_{t+1}|x_t)p(x_t|y_{1:t}). \tag{6.22}$$

Hence, using the particle approximation $\{w_t^j, x_t^j\}_{j=1}^{M_1}$ of the marginal filtering density obtained from the forward filtering pass (6.18), we obtain the representation

$$\widehat{p}^{M_1}(x_t|x_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = \sum_{j=1}^{M_1} w_{t|t+1}^j \delta_{x_t^j}(x_t), \tag{6.23}$$

with smoothing weights

$$w_{t|t+1}^j(x_{t+1}) \propto p(x_{t+1}|x_t^j) w_t^j. \tag{6.24}$$

The backward simulation consists thus, given a sample $\widetilde{x}_{t+1:\mathrm{T}}$ of $p(x_{t+1:\mathrm{T}}|y_{1:\mathrm{T}})$, into sampling $\widetilde{x}_t$ from $p(x_t|\widetilde{x}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$ using (6.23) and (6.24). The scheme is initialized with $\widetilde{x}_\mathrm{T}$ sampled from the point-mass marginal filtering distribution $p(x_\mathrm{T}|y_{1:\mathrm{T}})$. Repeating the scheme M_2 times, we obtain a collection of backward trajectories $\{\widetilde{x}_{1:\mathrm{T}}^i\}_{i=1}^{M_2}$ and a point-mass approximation of the joint smoothing distribution, according to

$$\widetilde{p}^{M_2}(x_{1:\mathrm{T}}|y_{1:\mathrm{T}}) := \frac{1}{M_2} \sum_{i=1}^{M_2} \delta_{\widetilde{x}_{1:\mathrm{T}}^i}(x_{1:\mathrm{T}}). \tag{6.25}$$

Note that we do not need to have $M_1 = M_2$. [LS13] discusses the design of these parameters, advocating $M_1 > M_2$, that is more computational effort on the forward pass.

Rao-Blackwellisation

The general FFBSi algorithm may be applied to our nonlinear model in (6.8a)-(6.8c). However, the tractable substructure in this model may be exploited to improve the performances and obtain better estimates than the PF. Indeed, the linear substructure may be treated analytically, instead of resorting to a particle approximation. This variance-reduction technique is called *Rao-Blackwellisation*, see [DGA00, CGM07] and F. Lindsten's PhD Thesis [Lin11] for an extensive treatment of the topic. [SGN05] derives the Rao-Blackwell particle filter (RBPF) for the class of mixed linear/nonlinear models that includes the model (6.8).

Regarding the filtering pass, the Rao-Blackwellisation exploits the factorization

$$p(\xi_{1:t}, \alpha_t|y_{1:t}) = p(\alpha_t|\xi_{1:t}, y_{1:t}) p(\xi_{1:t}|y_{1:t}), \tag{6.26}$$

with, given the model in (6.8a)-(6.8c),

$$p(\alpha_t|\xi_{1:t}, y_{1:t}) = \mathcal{N}(\alpha_t; \bar{\alpha}_{t|t}, P_{t|t}) \tag{6.27}$$

for some mean and covariance functions $\bar{\alpha}_{t|t}(\xi_{1:t})$ and $P_{t|t}(\xi_{1:t})$ and where $\mathcal{N}(\,\cdot\,; \mu, \Sigma)$ denotes generically the Gaussian density with mean vector μ and covariance matrix Σ. Hence, the PF may be applied to estimate the nonlinear state joint filtering distribution $p(\xi_{1:t}|y_{1:t})$ by a weighted system $\{w_t^j, \xi_{1:t}^j\}_{j=1}^{M_1}$, while the moments $\bar{\alpha}_{t|t}(\xi_{1:t}^j)$ and $P_{t|t}(\xi_{1:t}^j)$ of the conditionally linear state are estimated from a Kalman filter run for each particle $\xi_{1:t}^j$. The approximated joint filtering distribution is therefore

$$\widehat{p}^{M_1}(\xi_{1:t}, \alpha_t|y_{1:t}) = \sum_{j=1}^{M_1} w_t^j \mathcal{N}(\alpha_t; \bar{\alpha}_{t|t}^j, P_{t|t}^j) \delta_{\xi_{1:t}^j}(\xi_{1:t}). \tag{6.28}$$

The mean and covariance of the conditionally linear state are obtained analytically by a straightforward application of the Kalman filter, since the linear state does not bring information on the nonlinear state.

Lemma 6.2.1. *For the CLGSS model given by (6.8), the conditional density functions for $\alpha_{t|t-1}$ and $\alpha_{t|t}$ are*

$$p(\alpha_t|y_{1:t-1},\xi_{1:t}) = \mathcal{N}(\bar{\alpha}_{t|t-1}, P_{t|t-1}), \tag{6.29a}$$
$$p(\alpha_t|y_{1:t},\xi_{1:t}) = \mathcal{N}(\bar{\alpha}_{t|t}, P_{t|t}), \tag{6.29b}$$

where

$$\bar{\alpha}_{t|t-1} = f_\alpha(\xi_{t-1}, u_t, \theta), A_\alpha(\theta)\bar{\alpha}_{t-1|t-1}, \tag{6.30a}$$
$$P_{t|t-1} = A_\alpha P_{t-1|t-1} A_\alpha^\intercal + Q_\alpha, \tag{6.30b}$$

and

$$\bar{\alpha}_{t|t} = \bar{\alpha}_{t|t-1} + K_t(y_t - h(\xi_t) - C\bar{\alpha}_{t|t-1}), \tag{6.31a}$$
$$P_{t|t} = (I - K_t C)P_{t|t-1}, \tag{6.31b}$$
$$S_t = CP_{t|t-1}C^\intercal + R, \tag{6.31c}$$
$$K_t = P_{t|t-1}C^\intercal S_t^{-1}. \tag{6.31d}$$

Proof. See [SGN05]. □

The particle filtering step, targeting $p(\xi_{1:t}|y_{1:t})$ is similar to the general scheme previously described. The main difference lies into the fact that now $p(y_t|\xi_{1:t}, y_{1:t-1}) \neq p(y_t|\xi_t)$, that is the Markovian property is lost. As a consequence, instead of (6.19), the target distribution is factorized into

$$p(\xi_{1:t}|y_{1:t}) \propto p(y_t|\xi_{1:t}, y_{1:t-1})p(\xi_{1:t}|y_{1:t-1}) = p(y_t|\xi_{1:t}, y_{1:t-1})p(\xi_t|\xi_{1:t-1}, y_{1:t-1})p(\xi_{1:t-1}|y_{1:t-1}). \tag{6.32}$$

If the importance density is chosen such that it factorizes into

$$q(\xi_{1:t}|y_{1:t}) = q(\xi_t|\xi_{1:t-1}, y_{1:t})q(\xi_{1:t-1}|y_{1:t-1}), \tag{6.34}$$

then the Rao-Blackwellised weights are

$$w_t^j \propto w_{t-1}^j \frac{p(y_t|\xi_{1:t}^j, y_{1:t-1})p(\xi_t^j|\xi_{1:t-1}^j, y_{1:t-1})}{q(\xi_t^j|\xi_{1:t-1}^j, y_{1:t})}. \tag{6.35}$$

With a proposal density $q(\xi_t|\xi_{1:t-1}, y_{1:t}) := p(\xi_t|\xi_{1:t-1}, y_{1:t-1})$ and a systematic resampling scheme, (6.35) shrinks down to $w_t^j \propto p(y_t|\xi_{1:t}^j, y_{1:t-1})$. The task is therefore to sample from the density $p(\xi_t|\xi_{1:t-1}, y_{1:t-1})$ and to compute $p(y_t|\xi_{1:t}^j, y_{1:t-1})$. These densities are normal and given by Lemma 6.2.2.

Lemma 6.2.2. *For the CLGSS model given by (6.8), the densities $p(\xi_t|\xi_{1:t-1}, y_{1:t-1})$ and $p(y_t|\xi_{1:t}^j, y_{1:t-1})$ are given by*

$$p(\xi_t|\xi_{1:t-1}, y_{1:t-1}) = \mathcal{N}\left(f_\xi(\xi_{t-1}), \sigma_{\xi,t-1}^2\right), \tag{6.36}$$
$$p(y_t|\xi_{1:t}^j, y_{1:t-1}) = \mathcal{N}\left(h(\xi_t) + C\bar{\alpha}_{t|t-1}, CP_{t|t-1}C^\intercal + R\right). \tag{6.37}$$

Proof. See [SGN05]. □

Algorithm 6.1 Rao-Blackwellised particle filter

Inputs: $M_1, \bar{\alpha}_0, \bar{P}_0$, initial distribution $p_{\xi_0}(\xi_0)$.

- Initialize M_1 particles $\xi^j_{0|-1} \sim p_{\xi_0}(\xi_0)$, set $\{\alpha^j_{0|-1}, P^j_{0|-1}\} = \{\bar{\alpha}_0, \bar{P}_0\}$, $j = 1, \ldots, M_1$.
- Forward pass: for $t = 1, \ldots, \mathrm{T}$,
  1. Compute the importance weights
  $$w^j_t \propto p(y_t|\xi^j_{1:t}, y_{1:t-1}), \quad j = 1, \ldots, M_1 \tag{6.33}$$
  via Lemma 6.2.2 and normalize.
  2. Measurement update of the linear state: use Lemma 6.2.1 to compute $\bar{\alpha}^j_{t|t}$ and $\bar{P}^j_{t|t}$ such that $p(\alpha_t|y_{1:t}, \xi_{1:t}) = \mathcal{N}(\bar{\alpha}^j_{t|t}, \hat{P}^j_{t|t})$, for $j = 1, \ldots, M_1$.
  3. Time update (prediction) of the nonlinear state: for $j = 1, \ldots, M_1$, use Lemma 6.2.2 to sample $\xi^j_{t+1|t} \sim p(\xi_{t+1}|y_{1:t}, \xi^j_{1:t})$ and set $\xi^j_{1:t+1} := \{\xi^j_{1:t}, \xi^j_{t+1|t}\}$.
  4. Time update (prediction) of the linear state: use Lemma 6.2.1 to compute $\bar{\alpha}^j_{t+1|t}$ and $P^j_{t+1|t}$ such that $p(\alpha_{t+1}|y_{1:t}, \xi_{1:t}) = \mathcal{N}(\bar{\alpha}^j_{t+1|t}, P^j_{t+1|t})$, for $j = 1, \ldots, M_1$.

Output: a particle representation $\{\xi^j_t, w^j_t, \bar{\alpha}^j_t, \bar{P}^j_t\}_{j=1}^{M_1}$ of the joint filtering distribution.

Finally, the RBPF consists thus into sampling from the nonlinear state ξ and propagating the sufficient statistics of the conditionally linear state α. The main steps are sketched in Algorithm 6.1.

Let us focus now on the more challenging smoothing task. Similarly to the general nonlinear case, the Rao-Blackwellised particle smoother (RBPS) is a FFBSi algorithm, with forward filtering given by the RBPF. It remains to construct a backward simulator, which samples only the nonlinear state ξ. The target density is

$$p(\xi_{1:\mathrm{T}}|y_{1:\mathrm{T}}) = p(\xi_{1:t}|\xi_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})p(\xi_{t+1:\mathrm{T}}|y_{1:\mathrm{T}}). \tag{6.38}$$

If a backward simulator has been run from time T to time $t+1$, resulting in a trajectory $\widetilde{\xi}_{t+1:\mathrm{T}}$, and sufficient statistics $\widetilde{\alpha}_{t+1|\mathrm{T}}$, $\widetilde{P}_{t+1|\mathrm{T}}$ of the linear state, the strategy is two-fold:

1. Sampling: append a new state $\xi^{j_0}_t$ to $\widetilde{\xi}_{t+1:\mathrm{T}}$, where $\xi^{j_0}_t$ is drawn from the collection of particles $\{\xi^j_{1:t}\}_{j=1}^{M_1}$ obtained from the RBPF, $\xi^{j_0}_{1:t-1}$ being discarded. The updated trajectory is then $\widetilde{\xi}_{t:\mathrm{T}} := \{\xi^{j_0}_t, \widetilde{\xi}_{t+1:\mathrm{T}}\}$.

2. Smoothing: update the expressions $\widetilde{\alpha}_{t|\mathrm{T}}$ and $\widetilde{P}_{t|\mathrm{T}}$ of the sufficient statistics of the smoothed linear state α_t.

This scheme is then iterated backward in time until $t = 1$, while it is initialized at $t = \mathrm{T}$ by resampling from the filtering distribution at time $t = \mathrm{T}$, that is $\{\widetilde{\xi}_{\mathrm{T}}, \widetilde{\alpha}_{\mathrm{T}}, \widetilde{P}_{\mathrm{T}}\} = \{\xi^j_{\mathrm{T}}, \alpha^j_{\mathrm{T}|\mathrm{T}}, P^j_{\mathrm{T}|\mathrm{T}}\}$ with probability w^j_{T}.

The difficulty is to compute the backward sampling probabilities that generate $\xi_t^{j_0}$. Such computations are given in [LBS+16]. Here, we opt for a second option, described in [LS10a], according to which it is actually easier to sample from the joint distribution

$$p(\xi_{1:t}, \alpha_{t+1}|\xi_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = p(\xi_{1:t}|\alpha_{t+1}, \xi_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})p(\alpha_{t+1}|\xi_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}). \tag{6.39}$$

The idea is to sample $\tilde{\alpha}_{t+1}$ from the marginal density $p(\alpha_{t+1}|\widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$, and then sample $\widetilde{\xi}_{1:t}$ from the conditional density $p(\xi_{1:t}|\widetilde{\alpha}_{t+1}, \widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$, discard $\widetilde{\xi}_{1:t-1}$ and set $\widetilde{\xi}_t = \xi_t^{j_0}$. The marginal $p(\alpha_{t+1}|\widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$ is normal and simply given by:

$$p(\alpha_{t+1}|\widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = \mathcal{N}\left(\alpha_{t+1}; \widetilde{\alpha}_{t+1|\mathrm{T}}, \widetilde{P}_{t+1|\mathrm{T}}\right) \tag{6.40}$$

The first factor $p(\xi_{1:t}|\alpha_{t+1}, \xi_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$ of (6.39) is approximated using the particle representation $\{w_t^j, \xi_{1:t}^j\}_{j=1}^{M_1}$ obtained from the forward filtering pass. The resulting point-mass approximation

$$p(\xi_{1:t}|\widetilde{\alpha}_{t+1}, \widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = \sum_{j=1}^{M_1} w_{t|\mathrm{T}}^j \delta_{\xi_{1:t}^j}(\xi_{1:t}), \tag{6.41}$$

is given in Lemma 6.2.3. Hence, we have $\xi_{1:t} = \xi_{1:t}^j$ with probability $w_{t|\mathrm{T}}^j$, for $j = 1, \ldots, M_1$.

Lemma 6.2.3. *For the CLGSS given by (6.8), the density $p(\xi_{1:t}|\widetilde{\alpha}_{t+1}, \widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}})$ is approximated by*

$$p(\xi_{1:t}|\widetilde{\alpha}_{t+1}, \widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = \sum_{j=1}^{M_1} w_{t|\mathrm{T}}^j \delta_{\xi_{1:t}^j}(\xi_{1:t}), \tag{6.42}$$

where

$$w_{t|\mathrm{T}}^j(\widetilde{\xi}_{t+1}) \propto w_t^j p(\widetilde{\xi}_{t+1}, \widetilde{\alpha}_{t+1}|\xi_{1:t}^j, y_{1:t}), \tag{6.43}$$

and

$$p(\xi_{t+1}, \alpha_{t+1}|\xi_{1:t}^j, y_{1:t}) = \mathcal{N}\left(\begin{pmatrix}\xi_{t+1}\\ \alpha_{t+1}\end{pmatrix}; f(\xi_t^j) + A\bar{\alpha}_{t|t}^j, AP_{t|t}^j A^\intercal + Q\right). \tag{6.44}$$

Proof. See [LS10b]. □

Once the sampling step is completed, and a new particle $\widetilde{\xi}_t$ is appended to the nonlinear state trajectory $\widetilde{\xi}_{t+1:\mathrm{T}}$, it remains to update the sufficient statistics of the linear state, i.e. to compute the mean $\widetilde{\alpha}_{t|\mathrm{T}}$ and the covariance $\widetilde{P}_{t|\mathrm{T}}$. The cross-covariance

$$M_{t|\mathrm{T}} := \mathrm{cov}\left\{\alpha_t \alpha_{t+1}^\intercal \middle| \widetilde{\xi}_{t:\mathrm{T}}, y_{1:\mathrm{T}}\right\} \tag{6.45}$$

is also needed in order to compute $\mathcal{I}_2(\theta, \theta')$ in (6.16c). The updating equations are provided by Lemma 6.2.4.

Algorithm 6.2 Rao-Blackwellised particle smoother

Inputs: $M_1, \bar{\alpha}_0; , \bar{P}_0, p_{\xi_0}(\xi_0), M_2$

- Forward pass: run a forward pass of the RBPF Algorithm 6.1, and store the particles, the weights and the sufficient statistics for the linear state $\{\xi_t^j, w_t^j, \bar{\alpha}_{t|t}^j, P_{t|t}^j\}_{j=1}^{M_1}$
- Initialize: Resample M_2 particles $\{\widetilde{\xi}_{\mathrm{T}}^i, \widetilde{\alpha}_{\mathrm{T}|\mathrm{T}}^i, \widetilde{P}_{\mathrm{T}|\mathrm{T}}^i\}_{i=1}^{M_2}$ at time $t = \mathrm{T}$, that is with probabilities w_{T}^j, $j = 1, \ldots, M_1$.
- Backward simulation: for $t = \mathrm{T}-1, \ldots, 1$, for $i = 1, \ldots, M_2$
  1. *Sampling:* for each trajectory $\widetilde{\xi}_{t+1:\mathrm{T}}^i$ and associated sufficient statistics $\widetilde{\alpha}_{t+1|\mathrm{T}}^i, \widetilde{P}_{t+1|\mathrm{T}}^i$, $i = 1, \ldots, M_2$,
     (a) sample $\tilde{\alpha}_{t+1}^i \sim p(\alpha_{t+1}|\widetilde{\xi}_{t+1:N}^i, y_{1:N})$ (6.40);
     (b) for $j = 1, \ldots, M_1$, compute $w_{t|\mathrm{T}}^j(\widetilde{\xi}_{t+1}^i)$ using Lemma 6.2.3, and set $\widetilde{\xi}_t^i = \xi_t^j$ with probability $w_{t|\mathrm{T}}^j(\widetilde{\xi}_{t+1}^i)$.
  2. *Smoothing:* update the sufficient statistics according to Lemma 6.2.4.

Output: sample trajectories $\{\widetilde{\xi}_{t|\mathrm{T}}^i, \widetilde{\alpha}_{t|\mathrm{T}}^i, \widetilde{P}_{t|\mathrm{T}}^i\}_{i=1}^{M_2}$ approximating the smoothing distribution.

Lemma 6.2.4. *For the CLGSS given by (6.8) and under the assumption that*

$$p(\alpha_{t+1}|\xi_t^j, \widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) \approx p(\alpha_{t+1}|\widetilde{\xi}_{t+1:\mathrm{T}}, y_{1:\mathrm{T}}) = \mathcal{N}\left(\alpha_{t+1}; \widetilde{\alpha}_{t+1|\mathrm{T}}, \widetilde{P}_{t+1|\mathrm{T}}\right), \tag{6.46}$$

the sufficient statistics and the cross-covariance for α_t are

$$\widetilde{\alpha}_{t|\mathrm{T}} = \Sigma_{t|t}^{j_0} W_\alpha \widetilde{\alpha}_{t+1|\mathrm{T}} + c_{t|t}^{j_0}(\widetilde{\xi}_{t+1}), \tag{6.47a}$$
$$\widetilde{P}_{t|\mathrm{T}} = \Sigma_{t|t}^{j_0} + M_{t|\mathrm{T}}(W_\alpha)^\intercal \Sigma_{t|t}^{j_0}, \tag{6.47b}$$
$$M_{t|\mathrm{T}} = \Sigma_{t|t}^{j_0} W_\alpha \widetilde{P}_{t+1|\mathrm{T}}, \tag{6.47c}$$
$$\tag{6.47d}$$

where j_0 refers to the index of the forward filtering particle that corresponds to the appended state $\widetilde{\xi}_t$ and

$$\begin{bmatrix} W_\xi & W_\alpha \end{bmatrix} := A^\intercal Q^{-1}, \tag{6.48a}$$
$$\Sigma_{t|t}^j := P_{t|t}^j - P_{t|t}^j A^\intercal \left(Q + A P_{t|t}^j A^\intercal\right) A P_{t|t}^j, \tag{6.48b}$$
$$c_{t|t}^j(\xi_{t+1}) := \Sigma_{t|t}^j \left(W_\xi(\xi_{t+1} - f_\xi(\xi_t^j)) - W_\alpha f_\alpha(\xi_t^j)) + (P_{t|t}^j)^{-1} \bar{\alpha}_{t|t}^j\right). \tag{6.48c}$$

Proof. See [LS10b]. □

The main steps of the RBPS are sketched in Algorithm 6.2. They are repeated M_2 times, in order to obtain M_2 smoothed trajectories.

6.2.3 Algorithm

Computing the integrals in the Expectation step

Assuming that the RBPS of Algorithm 6.2 has been run, yielding a particle approximation $\{\widetilde{\xi}^i_{t|\mathrm{T}}, \widetilde{\alpha}^i_{t|\mathrm{T}}, \widetilde{P}^i_{t|\mathrm{T}}\}_{i=1}^{M_2}$ of the smoothing densities $p(x_{1:t}|y_{1:\mathrm{T}})$, it is now possible to evaluate the $\mathcal{Q}-$function in (6.16). We reproduce here the computations of [LS10a]. Let us start with the term $\mathcal{I}_2(\theta,\theta')$ in (6.16c). We have

$$\begin{aligned}-2\log p_\theta(X_{t+1}|X_t) &= -2\log p_{v,\theta}(x_{t+1}-f(\xi_t)-A\alpha_t)\\ &= \log|Q| + (x_{t+1}-f(\xi_t)-A\alpha_t)^\intercal Q^{-1}(x_{t+1}-f(\xi_t)-A\alpha_t)\\ &= \log|Q| + \mathrm{tr}\left\{Q^{-1}\ell_2(\xi_{t:t+1},\alpha_{t:t+1})\right\}, \end{aligned} \tag{6.49}$$

where we have ignored the constant term in the normal density, and defined $\ell_2(\xi_{t:t+1},\alpha_{t:t+1}) := (x_{t+1}-f(\xi_t)-A\alpha_t)(x_{t+1}-f(\xi_t)-A\alpha_t)^\intercal$. It follows that

$$\mathcal{I}_2(\theta,\theta') = -\frac{1}{2}\sum_{t=1}^{\mathrm{T}-1}\mathbb{E}_{\theta'}\left[\log|Q| + \mathrm{tr}\left\{Q^{-1}\ell_2(\xi_{t:t+1},\alpha_{t:t+1})\right\}\middle| Y_{1:\mathrm{T}}\right]. \tag{6.50}$$

Using the RBPS approximation of the smooting densities, the particle-based estimation of the integral $\mathcal{I}_2(\theta,\theta')$ is

$$\mathcal{I}_2(\theta,\theta') \simeq \widehat{\mathcal{I}}_2(\theta,\theta') := -\frac{1}{2M_2}\sum_{t=1}^{\mathrm{T}-1}\sum_{j=1}^{M_2}\left(\log|Q| + \mathrm{tr}\left\{Q^{-1}\widehat{\ell}^j_{2,t}\right\}\right), \tag{6.51}$$

where $\widehat{\ell}^j_{2,t}$ is an expectation over the $\alpha-$variable only:

$$\widehat{\ell}^j_{2,t} := \mathbb{E}_{\theta'}\left[\ell_2(\widetilde{\xi}^j_{t:t+1},\alpha_{t:t+1})\middle|\widetilde{\xi}^j_{t:t+1}, Y_{1:\mathrm{T}}\right] \tag{6.52}$$

Since

$$\mathbb{E}[\alpha_t\alpha_t^\intercal|\widetilde{\xi}^j_{t:t+1}, Y_{1:\mathrm{T}}] = \widetilde{\alpha}^j_{t|\mathrm{T}}\widetilde{\alpha}^{j\intercal}_{t|\mathrm{T}} + \widetilde{P}^j_{t|\mathrm{T}}, \tag{6.53a}$$

$$\mathbb{E}[\alpha_t\alpha_{t+1}^\intercal|\widetilde{\xi}^j_{t:t+1}, Y_{1:\mathrm{T}}] = \widetilde{\alpha}^j_{t|\mathrm{T}}\widetilde{\alpha}^{j\intercal}_{t+1|\mathrm{T}} + M^j_{t|\mathrm{T}}, \tag{6.53b}$$

with $\widetilde{\alpha}^j_{t|\mathrm{T}}$, $\widetilde{P}^j_{t|\mathrm{T}}$ and $M^j_{t|\mathrm{T}}$ given in Lemma 6.2.4, straightforward calculations yield, with $\widetilde{x}^j_t := \begin{pmatrix}\widetilde{\xi}^j_t & \widetilde{\alpha}^{j\intercal}_{t|\mathrm{T}}\end{pmatrix}^\intercal$,

$$\widehat{\ell}^j_{2,t} = (\widetilde{x}^j_{t+1} - f(\widetilde{\xi}^j_t) - A\widetilde{\alpha}^j_{t|\mathrm{T}})(\widetilde{x}^j_{t+1} - f(\xi^j_t) - A\widetilde{\alpha}^j_{t|\mathrm{T}})^\intercal + \tag{6.54}$$

$$A\widetilde{P}^j_{t|\mathrm{T}}A^\intercal - \begin{bmatrix}0 & AM^j_{t|\mathrm{T}}\end{bmatrix} - \begin{bmatrix}0 & M^{j\intercal}_{t|\mathrm{T}}A^\intercal\end{bmatrix}^\intercal + \begin{bmatrix}0 & 0\\ 0 & \widetilde{P}_{t+1|\mathrm{T}}\end{bmatrix}, \tag{6.55}$$

where the 0's are matrices of appropriate dimensions.

Similar particle-based approximations are derived for the integrals $\mathcal{I}_1(\theta,\theta')$ and $\mathcal{I}_3(\theta,\theta')$ respectively:

$$\mathcal{I}_1(\theta,\theta') \simeq \widehat{\mathcal{I}}_1(\theta,\theta') := -\frac{1}{2M_2}\sum_{j=1}^{M_2}\left(\log|\Sigma_1| + \mathrm{tr}\left\{\Sigma_1^{-1}\widehat{\ell}^j_{1,t}\right\}\right) \tag{6.56}$$

and

$$\mathcal{I}_3(\theta,\theta') \simeq \widehat{\mathcal{I}}_3(\theta,\theta') := -\frac{1}{2M_2}\sum_{t=1}^{N}\sum_{j=1}^{M_2}\left(\log|R| + \mathrm{tr}\left\{R^{-1}\widehat{\ell}^{j}_{3,t}\right\}\right), \tag{6.57}$$

where

$$\widehat{\ell}^{j}_{1,t} := \mathbb{E}_{\theta'}\left[\left(x_1 - x_{1|0}\right)\left(x_1 - x_{1|0}\right)^{\intercal}\middle|\, \widetilde{\xi}^{j}_{1}, Y_{1:\mathrm{T}}\right] \tag{6.58}$$

$$= \left(\widetilde{x}_{1|\mathrm{T}} - x_{1|0}\right)\left(\widetilde{x}_{1|\mathrm{T}} - x_{1|0}\right)^{\intercal}, \tag{6.59}$$

with, following (6.10), $x_{1|0} := \begin{bmatrix}\xi_{1|0} & \alpha_{1|0}\end{bmatrix}^{\intercal}$, and

$$\widehat{\ell}^{j}_{3,t} := \mathbb{E}_{\theta'}\left[\left(y_t - h(\widetilde{\xi}^{j}_{t}) - C\alpha_t\right)\left(y_t - h(\widetilde{\xi}^{j}_{t}) - C\alpha_t\right)^{\intercal}\middle|\, \widetilde{\xi}^{j}_{t}, Y_{1:\mathrm{T}}\right] \tag{6.60}$$

$$= \left(y_t - h(\widetilde{\xi}^{j}_{t}) - C\widetilde{\alpha}^{j}_{t|N}\right)\left(y_t - h(\widetilde{\xi}^{j}_{t}) - C\widetilde{\alpha}^{j}_{t|\mathrm{T}}\right)^{\intercal} + C\widetilde{P}_{t|\mathrm{T}}C^{\intercal}. \tag{6.61}$$

Finally, the particle-based approximation of the $\mathcal{Q}-$function (6.16) is

$$\widehat{\mathcal{Q}}(\theta,\theta') = \widehat{\mathcal{I}}_1(\theta,\theta') + \widehat{\mathcal{I}}_2(\theta,\theta') + \widehat{\mathcal{I}}_3(\theta,\theta'), \tag{6.62}$$

where $\widehat{\mathcal{I}}_1(\theta,\theta')$, $\widehat{\mathcal{I}}_2(\theta,\theta')$ and $\widehat{\mathcal{I}}_3(\theta,\theta')$ are respectively given by (6.56), (6.51) and (6.57). This concludes the Expectation step of the EM algorithm.

Initialization

In order to facilitate the convergence of the algorithm, we use as initial value a naive estimation of the open-loop R3C2 network, with solar radiation $Qs \equiv 0$ as a deterministic input. This first estimator is computed using Algorithm 3.1 described in Chapter 3. Hence, the three main steps of the overall estimation procedure, sketched in Algorithm 6.3, are:

1. learning of the high variance increments with a ℓ_1- penalized criterion;
2. naive estimation of θ with $Q_s \equiv 0$;
3. estimation of θ starting from the naive estimator and using the RBPS-EM algorithm.

6.3 Numerical illustration

6.3.1 Simulated data

Validation procedure

The method described in this chapter is tested on several one-week datasets. The data are generated according to the procedure described in Chapter 2. In particular, meteorological records from November 2015 to January 2016 are used, as well as free internal gains corresponding to an office building. We use 200 iterations of Algorithm 6.3, with $M_1 = 30$ and $M_2 = 10$ particles for the forward filtering and backward smoothing, respectively. The number of particles is limited by the computing time: it typically requires around an hour and a half to run the 200 iterations with $M_1 = 30$ and $M_2 = 10$ on a laptop with an IntelCore i7 2.80Ghz processor - while the

Algorithm 6.3 EM for R3C2 with undirectly observed solar flux

Inputs: M_1, M_2, observed series $\{\Delta T_o(t),\, \vartheta_o(t),\, T_i(t),\, Q_r(t)\}_{t=1}^{\mathrm{T}}$.

1. Estimate the series of variances of the dynamics of the nonlinear state (6.6).
2. Initialize the parameters to θ_0 and set $k = 0$, where θ_0 is the standard open-loop estimator of the R3C2 network with obtained by setting $Q_s \equiv 0$ in Algorithm 3.1.
3. *E-step:*
   (a) smoothing: run the RBPS Algorithm 6.2 under the hypothetised model $\theta = \theta_k$ and store the particles and the sufficient statistics for the linear state;
   (b) approximate the $\mathcal{Q}-$function, using
   $$\widehat{\mathcal{Q}}(\theta, \theta_k) = \widehat{\mathcal{I}}_1(\theta, \theta_k) + \widehat{\mathcal{I}}_2(\theta, \theta_k) + \widehat{\mathcal{I}}_3(\theta, \theta_k), \tag{6.63}$$
   with $\widehat{\mathcal{I}}_1(\theta, \theta_k)$, $\widehat{\mathcal{I}}_2(\theta, \theta_k)$ and $\widehat{\mathcal{I}}_3(\theta, \theta_k)$ respectively given by (6.56), (6.51) and (6.57).
4. *M-step:* compute θ_{k+1}:
   $$\theta_{k+1} = \arg\max_{\theta} \widehat{\mathcal{Q}}(\theta, \theta_k). \tag{6.64}$$
5. If convergence is reached, terminate, otherwise set $k \leftarrow k + 1$ and return to step 3.

Outputs: estimated R3C2 parameters $\widehat{\theta}$, smoothed trajectories of the solar flux.

initial estimation θ_0 is obtained in around one minute and a half with the same machine. The Maximization step is carried out with Matlab's routine `fminunc`, based on the BFGS quasi-Newton method. Moreover, we choose fixed values $q_\eta = 0.15$ and $r_\eta = 0.15$, and estimate only the σ_t's at the first step of the algorithm. This procedure is repeated 85 times, with the same set of true parameters but different weather conditions. In order to evaluate the benefits of our approach, we compare the estimated $\widehat{\theta}$ obtained by the end of Algorithm 6.3 to the initial value θ_0 as well as an ARX model of order 2, both being representative of standard approaches of the literature with unobserved solar radiation.

Smoothed solar flux

We check first that the smoothed solar flux, estimated once Algorithm 6.3 has converged, is consistent with the (actually hidden) observations. The accurate smoothed trajectory of the hidden state η as well as the location of the high variance increments, shown in Figure 6.5 (top line) validates the state-space model (6.4)-(6.5). This smoothed trajectory is computed as the average of the $M_2 = 10$ backward trajectories of the RBPS. However, the estimation of the solar flux (Figure 6.5, bottom line) is less accurate, with a slight overestimation and an average root mean square error of 0.29 over all the particle trajectories. Nevertheless, the estimated dynamic trajectory reproduces the periods of low and high intensities of the solar radiation. This confirms that higher values of the outdoor temperature bias ΔT_o are indeed caused by higher values of the solar flux.

Moreover, this error on the solar flux is the combination of several errors: the estimation errors as well as the modelling errors in (6.1), (6.2) and (6.3). In particular, the physical model (6.1) of the solar radiation is among the simplest in the literature. Its advantage is that it is based on the geographical and time coordinates of the building location only, which are easily obtained data. Hence, despite its simplicity, the model still provides a good approximation of the solar gains entering into the building, and supports our approach.

The comparison to standard algorithms and models developed in the literature is made difficult due to the dissimilarities in the experimental conditions. Indeed, as highlighted in the introduction to this chapter, most of the literature is devoted to the problem of predicting future values of the solar radiation based on past measurements, while we do not measure the solar flux at all. However, the problem addressed in [OS14] is somewhat similar to our context. This paper aims at estimating the hourly solar radiation during long-term missing gaps in the measured series, with the application to building performance simulation as a motivation. Their results vary with the length of the missing gaps. For a gap of five days, the best technique is the singular spectrum analysis, achieving normalized root mean squared errors (NRMSE) varying in the range 30.83%-70.35% between November and February in Oklahoma City North (Oklahoma, warm and humid climate), where the NRMSE is the RMSE normalized by the mean of the observations, and is furthermore averaged on a monthly basis. For a gap of ten days however, the NRMSE varies between 49.98% and 78.91%, where their best algorithm is the temperature-based approach.

On the other hand, the NRMSE obtained with the algorithm suggested in this chapter, and averaged over the all the particle trajectories at time-step 10 minutes is 67.56%. When aggregating the series at the hourly time-step used in [OS14], the averaged NRMSE is 56.43%. Note that when only the diurnal data are taken into account in the computation, the average NRMSE is then 35.14%. Although the regions and climate differ, this value compares rather well with [OS14], which further validates the model.

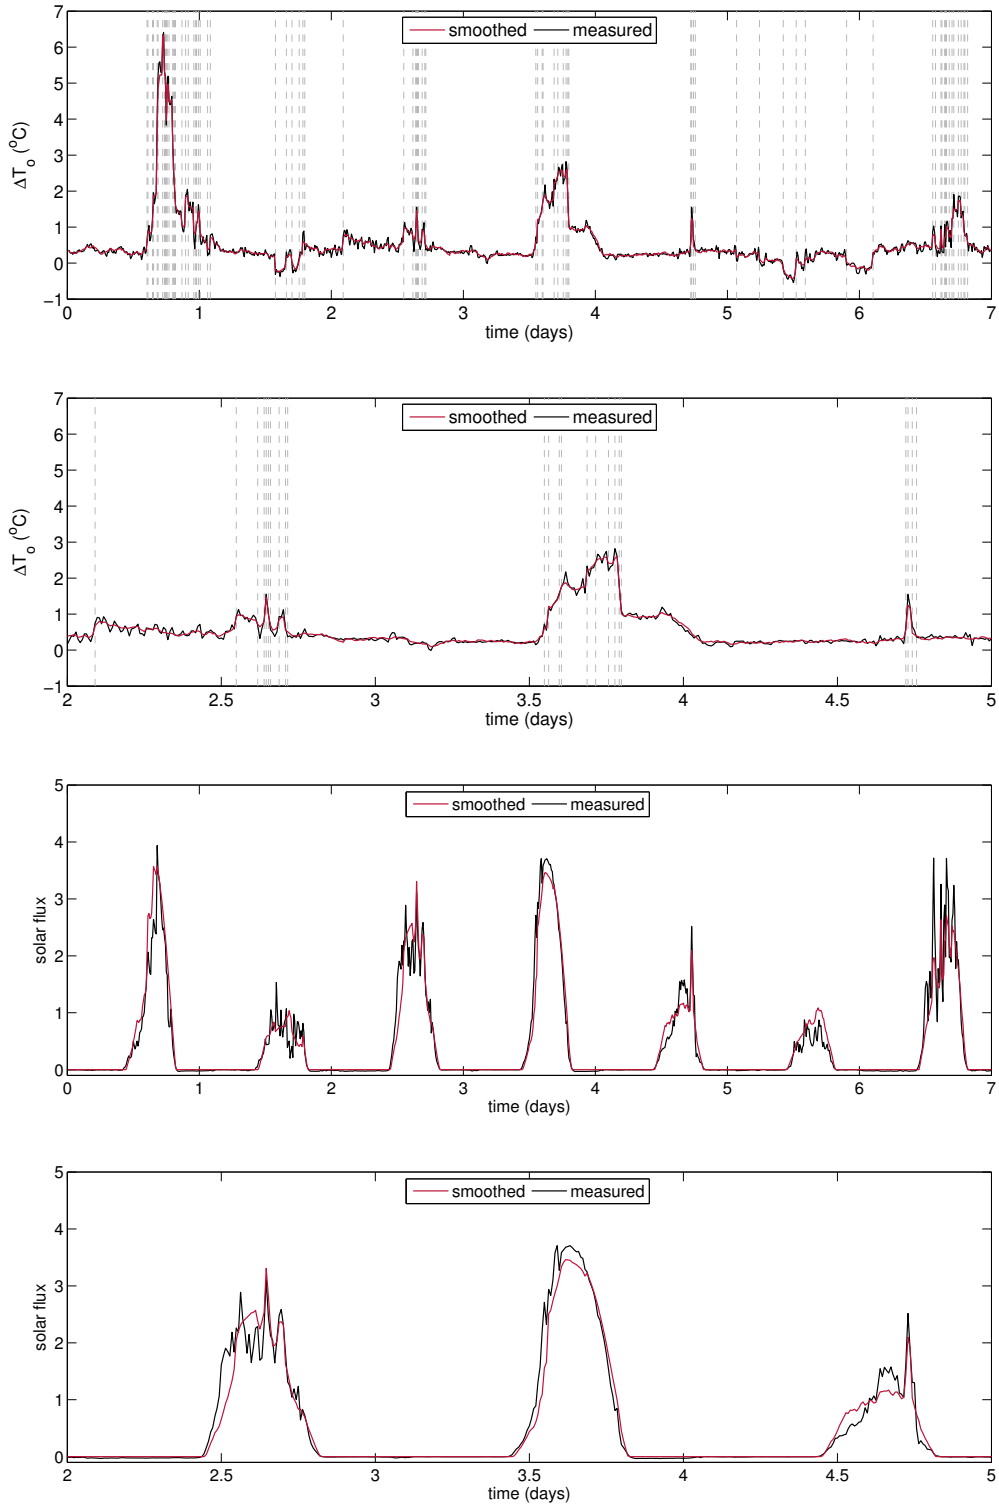

Figure 6.5: From top to bottom, examples of smoothed trajectories of ΔT_o, three-day zoom, smoothed trajectories of the solar flux, three-day zoom. The vertical lines locate the estimated high variance increments.

Table 6.2: Estimated parameters with twice their standard errors for the naive method (θ_0), RBPS-EM (Algorithm 6.3, $\widehat{\theta}$) and an ARX model of order 2. 85 experiments, of which 5 outliers are not taken into account for the naive estimator, 6 for RBPS-EM.

	True value	**Naive estimator**	**RBPS-EM (Alg. 6.3)**	**ARX**
UA	0.370	0.354 ± 0.083	0.367 ± 0.021	0.326 ± 0.025
g	0.400	0.376 ± 0.338	0.390 ± 0.162	-
τ_1	41.4	37.9 ± 26.3	33.0 ± 17.6	5.8 ± 1.3
τ_2	3.6	4.5 ± 1.6	3.8 ± 0.4	0.20 ± 0.15

Accuracy of the estimates

The estimations of the four physical parameters of the R3C2 network (heat loss coefficient UA, solar transmission g, slow and fast time constants τ_1 and τ_2) are summarized in Table 6.2 and Figure 6.6. Figure 6.7 shows that the convergence of the EM algorithm is reached by the end of the 200 iterations. We observe first that $\widehat{\theta}$, the output of Algorithm 6.3, yields a significant reduction of the variance of the estimation of the four physical parameters, compared to the naive estimation θ_0. The estimations of UA and of the fast time constant τ_2 are thus much more accurate. This emphasizes the importance of including the Rao-Blackwellisation within the particle filters and smoother, in accordance with the experimental results in [LS10a]. Besides, our method reduces also the bias on g, consistently with its physical interpretation as the proportion of the solar flux transmitted inside the building. This illustrates the benefits of learning the dynamics of the solar flux. Finally, the slow time constant τ_1 is underestimated, although $\widehat{\theta}$ provides the correct order of magnitude. This bias may be partly explained by the length of the training dataset, i.e. seven days only whereas the numerical simulations in Chapter 3 show that adding a few more days significantly improves the estimation.

Furthermore we simulated the indoor temperature and heating power on a test dataset, with the 85 estimations of the R3C2 network. The temperature is simulated in closed loop, with parameters of the regulation set to their true values. The simulated trajectories, and the one obtained with the true parameters, are displayed on Figure 6.8. The predicted power in Figure 6.8(a) has a normalized RMSE ranging between 4.9% and 17.8%, with an average of 7.9%, while the absolute error on the energy consumption lies between 0.2% and 13.8%, with an average of 3.8%. The temperature trajectories are in closer agreement, even for the less accurate parameter estimations. The RMSE ranges between 0.16°C and 0.82°C, with an average of 0.26°C, corresponding to a normalized RMSE's of 0.9%, 4.8% and 1.4%, respectively.

The results of Algorithm 6.3 are also compared in Table 6.2 to a state-of-the-art identification model, an ARX model. Such a model has been used for instance in [JMA08, HGP12] and in [MCPF10] for benchmarking. For a fair comparison with the R3C2 network, we choose a model of order 2, such that two time constants can be computed from the poles of the model. The auto-regressive output is the indoor temperature, while the exogenous inputs are the outdoor temperature ϑ_o, the heating power Q_r, the outdoor temperature bias ΔT_o and the clear-sky radiation Q_{clr}. Hence, the model is identified from the same dataset as $\widehat{\theta}$. The orders with respect to the exogenous inputs are set to 2, as in [MCPF10, HGP12]. The parameters of the resulting ARX(2,2,2,2,2) model are estimated by minimization of the one-step ahead prediction error, using ordinary least squares. In addition to the two time constants, a heat transfer coefficient can also be computed from the identified ARX representation, see e.g. [JMA08]. However, the solar transmission g cannot be

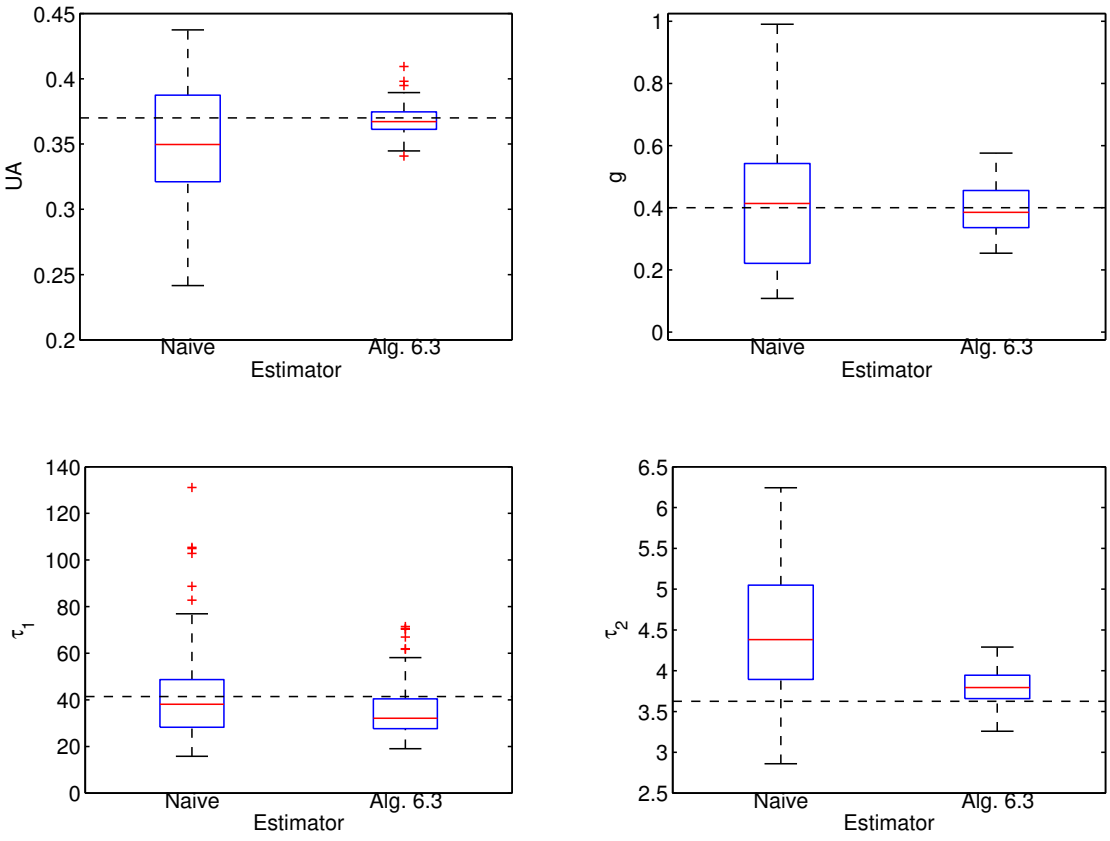

Figure 6.6: From left to right, top to bottom: boxplots on 85 experiments of the estimated UA, solar transmission, slow and fast time constants, with true values as horizontal dashed lines. Left boxes of each plot are for the naive estimator, right for Algorithm 6.3.

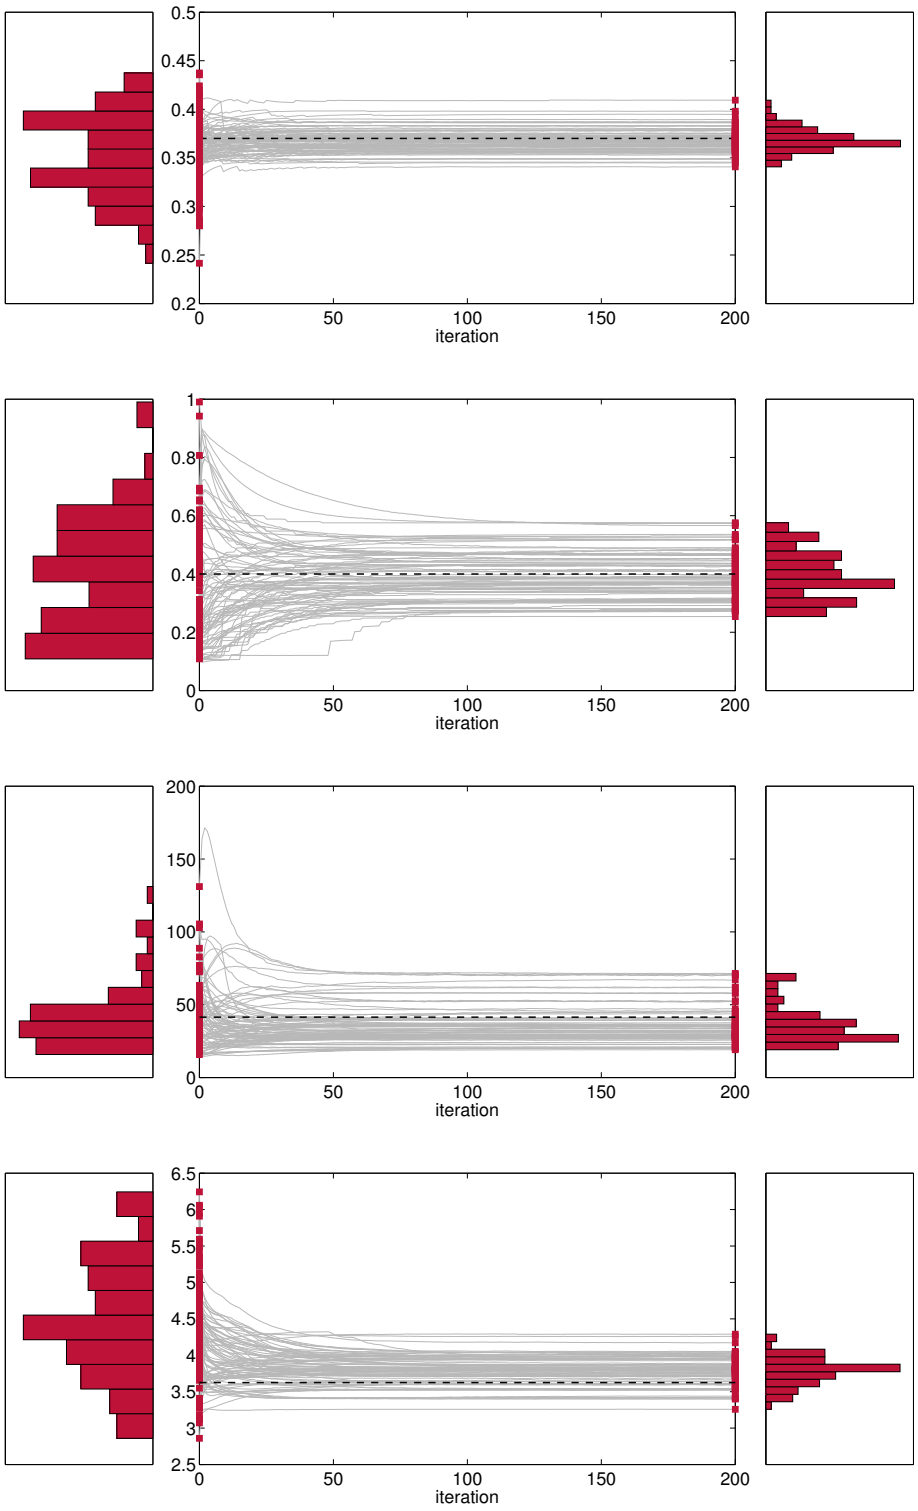

Figure 6.7: From top to bottom, estimation of the heat loss coefficient UA, the solar transmission g, the slow and fast time constants τ_1 and τ_2, as a fonction of the number of iterations of the RBPS-EM algorithm. Each line represents one of the 85 estimations. The true values are the horizontal black dashed lines. The histograms are the distributions of θ_0 (naive estimator, left) and $\widehat{\theta}$ (RBPS-EM, right).

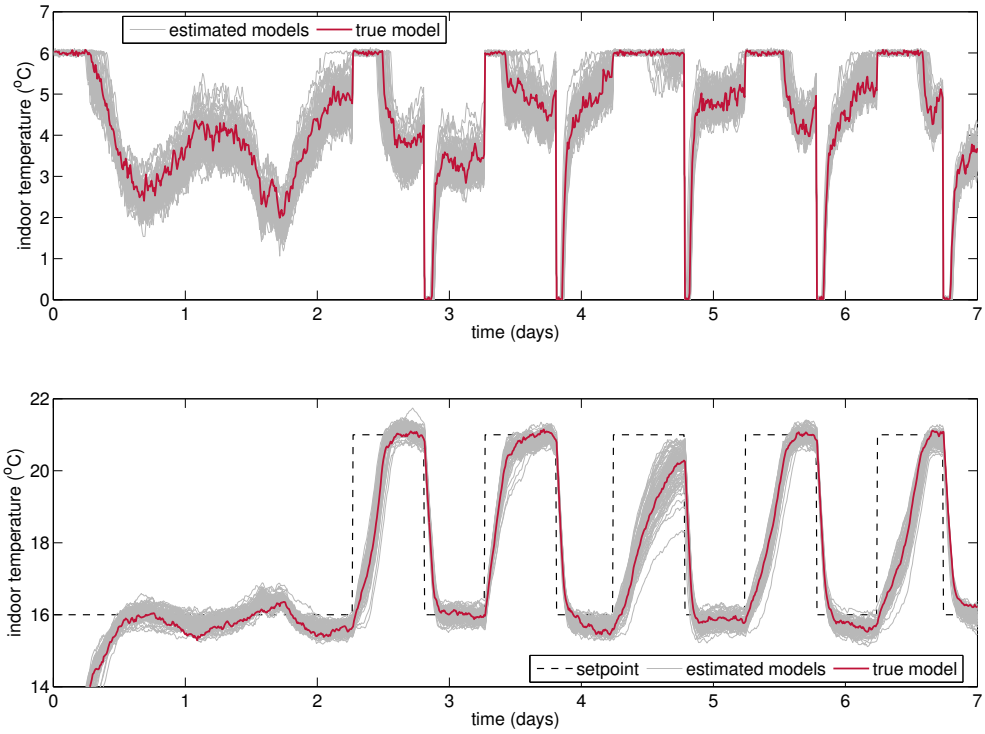

Figure 6.8: Closed loop simulation of the heating power (top) and the indoor temperature (bottom) with the true (red) and the 85 estimated (grey) parameters, on a test set with meteorological data from BESTLab and true regulation parameters.

computed, since the solar radiation is not measured.

Table 6.2 shows that the estimator based on Algorithm 6.3 largely outperforms the ARX model. Indeed, the ARX model can approximate the UA coefficient, but largely underestimates the time constants of the building. This confirms the relevancy of reconstructing the dynamics of the solar flux from the observation of the outdoor temperature bias, rather than using a black box model such as an ARX. These numerical experiments support thus the idea that a simple way to reduce the implementation cost of building identification algorithms consists in replacing the pyranometer by a cheap outdoor temperature sensor exposed to the sun. In the next section, we analyze the performance of Algorithm 6.3 on real data from the BESTLab experiment described in Section 2.2.1.

6.3.2 BESTLab data

Algorithm 6.3 is also tested with BESTLab data, with parameters tuned identically to those for the simulated dataset: 200 iterations, $M_1 = 30$ particles for the forward filtering step, $M_2 = 10$ for the backward smoothing, 7 days of data. The sheltered temperature sensor is located at climate station ETNA, whereas the sun-exposed sensor is alongside the southern or western cell (see Table 2.2). We repeated 35 independent estimations for the southern and western cells, each of them with a random excerpt of BESTLab data.

The convergence of the four physical parameters is shown in Figure 6.9 for the southern cell. More experiments are necessary before drawing unequivocal conclusions, yet the findings agree

Table 6.3: Estimated parameters with twice their standard errors for the naive method (θ_0) and RBPS-EM (Algorithm 6.3), for BESTLab's cells, with 35 experiments for each cell.

	Southern cell		**Western cell**	
	Naive estimator	RBPS-EM	Naive estimator	RBPS-EM
U	0.847 ± 0.182	0.986 ± 0.049	0.920 ± 0.252	1.016 ± 0.066
g	0.738 ± 0.406	0.607 ± 0.200	0.682 ± 0.443	0.596 ± 0.270
τ_1	75.9 ± 46.8	43.4 ± 13.3	71.4 ± 37.5	47.9 ± 17.4
τ_2	13.5 ± 9.8	7.5 ± 5.8	14.6 ± 8.2	9.6 ± 6.9

already with those of the simulated dataset. In particular, the naive estimator give plausible but much widespread estimates. On the contrary, the RBPS-EM estimations converge also to plausible values, but within narrower intervals (see also Table 6.3). For instance, the 35 estimated heat loss coefficients U converge to values within the bounds of the error bars in Figure 3.8. These results are all the more promising given that there are only 7 days of data in the learning sets.

The smoothed trajectories of the solar flux are not shown, since the environmental data also are identical to those used in the previous section: we obtain plots similar to those already displayed in Figure 6.5. Finally, the same conclusions can be drawn from Table 6.3 for the identification of the western cell, with sun-exposed temperature sensor alongside this cell. This illustrates the fact that the method can be used regardless of the orientation of the glazed wall, and not only for southern facades. The advantage of our method is therefore that it is very local, corresponding as much as possible to the microclimate of the building.

6.4 Summary

In this chapter, we have proposed a new model and learning algorithm such that it is not necessary to measure directly the solar radiation. The motivation was that local solar radiation measurements are most of the time missing for an energy utility. On the contrary, our approach is based on the availability of two measurements of the outdoor temperature, respectively with a sheltered and an exposed sensor. These measurements are easy to obtain for utilities. Hence, the main hypothesis is that the difference between the two measurments, called outdoor temperature bias, replaces the observation of the solar flux. The identification procedure is adapted consequently, by appending an *ad hoc* "meteorological" model to the R3C2 building model. This meteorological model links the solar radiation to the outdoor temperature bias by means of a simplified physical solar flux model, a cloud cover correction, and an *ad hoc* mapping of the outdoor temperature bias to the cloud cover index. This forms globally a conditionally linear Gaussian state space model, whose identification may be performed by a state-of-the-art algorithm, the Expectation-Maximization with Rao-Blackwellised Particle Smoothing.

The numerical evaluation of the model, both on simulated and real data, shows that it is possible to retrieve the received solar flux, at a time-step of 10 minutes, with a reasonable accuracy. Besides, the algorithm allows a good estimation of the physical constants of the building model, with a reduced bias and variance. Its limiting aspect is the computational time, although the code written in this work did not optimize it. It is in particular possible to execute the code on parallel workers for the time-consuming particle filtering.

Because thermometers are much more affordable and easily deployed than pyranometers, we believe that this chapter supports the idea that the microclimate of any given building may be

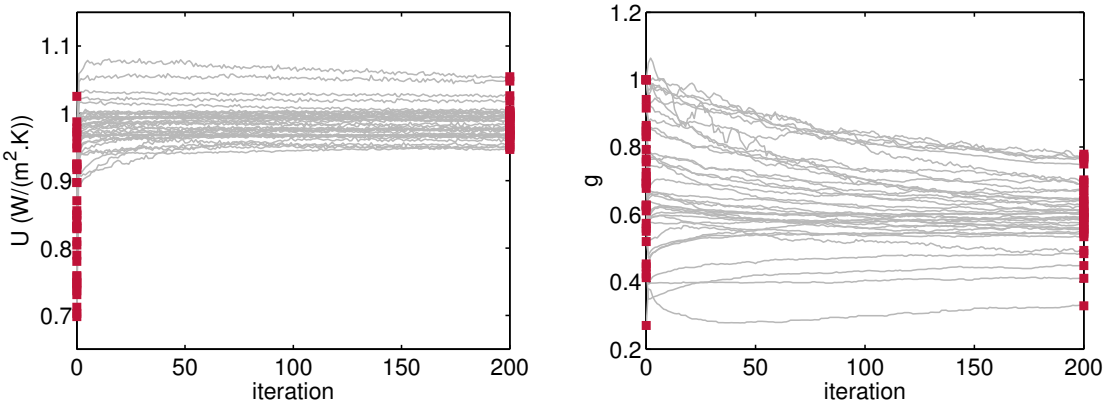

(a) U (b) g

(c) τ_1 (d) τ_2

Figure 6.9: Estimation of (a) the heat loss coefficient U, (b) the solar transmittance g, (c) the slow time constant τ_1, (d) the fast time constant τ_2, for BESTLab Southern cell, as a fonction of the iterations of the RBPS-EM algorithm. Each line represents one of the 35 estimations. Left squares denote the naive estimation θ_0 (iteration = 0), right squares (iteration = 200) $\widehat{\theta}$ obtained with RBPS-EM.

accessed in this convoluted manner, while not deteriorating the process of building identification.

Chapter 7

Uncertainty on the heating load curve

This chapter is devoted to the last input of the R3C2 building model, namely the heating power. The deployment of the smart electricity and gas meters - for instance, 35 millions of them are to be installed in France by 2021 for the electricity, 11 millions by 2022 for the gas, whereas a complete rollout is expected by the end of 2020 for the United Kingdom - makes it reasonable to assume that an energy utility has access to the measurement of the overall energy consumption of a given building, at a dynamic time-step typically less than an hour. However, the open-loop identification of the R3C2 model described in Chapter 2 requires more especially the extraction from this total load curve of the consumption attributable only to the specific loads that contribute to heating. Equivalently, this means that the measured total load curve contains non-heating loads that must be discarded, i.e. substracted from the initial measurement. Hence, the quantity of interest remains hidden. Nevertheless, the total load curve constitutes at least an upper bound. Besides, the metabolic heat gains due to the occupancy of the building are not measured at all, whereas they should also be taken into account for the identification.

In this context, the problem of building identification seems difficult to solve without any additional information about the building. As previously stated in the Introduction to this thesis, the idea developed in this chapter consists thus into considering the more specific case of a connected building, that is to say a building equipped with connected objects. Moreover, the main assumption is that the energy utility has an access to the information supplied by these connected objects. Our purpose is therefore to use these information to help reduce the uncertainty on the actual heating power, which we call the *useful load*. Note that how the utility obtains these information in practice goes beyond the scope of the thesis, and is not discussed furthermore.

Consequently, the first task is to characterize more precisely the connected objects of the building. In the following Section 7.1, we establish a typology and review the information provided by the connected objects. This typology helps us formulate a set of work hypotheses, consistent with the current technology. Then, the resulting mathematical model is described in Section 7.2. Before evaluating numerically the model in Section 7.4, a new data generating process is introduced in Section 7.3. Indeed, the former dataset used in Chapters 3 to 6 contains only heating loads, which is unfit for this chapter. Section 7.3 presents thus a method for generating stochastic domestic load curves, with appliance specific consumptions including heating and non-heating loads.

7.1 Characterization of the connected objects

7.1.1 Typology

In this section, we review the loads that contribute most to the total electric energy consumption, in a domestic building. Our purpose is to characterize the information that *smart* or *connected* versions of these appliances bring - or would be likely to bring in a near future - for improving the observation of the thermal behaviour of the building. The loads are categorized depending on whether or not they contribute to heating.

Heating loads

Heaters The heating system represents most of the load curve during winter times. For instance, heating represented an average of 61% of the total consumption during an experiment carried out internally at EDF R&D over 22 dwellings. Moreover, an electrical heater requires a maximum power that ranges typically between 500 W and 2000 W. This explains why in the absence of any information and as a first approximation, the measured aggregated load curve may be seen as contributing entirely to heating, and is thus used as input of the R3C2 building model.

In France, about a third of the housing market is dominated by the electrical heaters (CEREN 2014), which we focus on from now on. A typical connected solution already on the market provides the daily consumption of the heaters, together with a measurement of the indoor temperature. However, it is not clear whether this daily consumption is actually measured, or if it is estimated by the heater. This underlying uncertainty of the information provided by the connected object is shared by many other devices. From an internal test at EDF R&D, it seems nevertheless that these consumptions are accurate. Unfortunately, the daily consumption does not say much of the fine-grained consumption that would be representative of the power regularly injected during the day. Hence, this fine-grained consumption remains hidden, although it is necessary for identifying a dynamic thermal model with a 10-minute resolution.

Refrigerator, oven Although couterintuitive at first sight perhaps, we consider that both the refrigerator and the oven contribute to heating. As a first approximation, we suppose that the heat gains of these two appliances amount to their respective electrical consumptions. Typical one-day load curves for the refrigerator and the oven are shown in Figure 7.1, top left and center plots respectively. These data come from an experiment performed at EDF R&D, in an actual dwelling. As can be seen from these plots, both appliances represent non-negligible heat gains in this approximation: the refrigerator through a cyclic component including a peak and a low plateau, the oven through occasional high peak values.

The development of connected refrigerators and ovens aims mainly at providing remote monitoring and control to the user. For instance, a typical solution allows one to control the internal temperature and functioning mode of both appliances, via an application. Nevertheless, the actual temperature is not necessarily returned to the user. One additional feature is a feedback on the daily consumption of the appliance. However, as of today, this feedback takes the form of discrete levels: either low, medium or high. Furthermore, internal tests at EDF R&D showed that these consumptions were inaccurate, thereby indicating that the technology is not ready yet.

Lighting The weight of lighting in the total load curve is admittedly low, with power ratings ranging from 10W for a LED lamp to 100W and more for incandescent lamps. However, lighting represents a use with strong repetability, typically several hours everyday during the heating season. A typical one-day measurement of the lighting units in a given dwelling is shown in Figure 7.1,

top right plot. We can consider that lighting contributes to heating, whatever the underlying technology. Moreover, we may assume that the heat generated by the lamp amounts to its total electrical consumption.

There are already some connected lights on the market. For instance, a standard solution typically integrates a wireless lighting system and smart controls via an application. Another way of having a connected lighting system is to install a (wireless) relay and integrate it in a home automation system (see e.g. EnOcean products, similarly to our experiment at BESTLab). However, both solutions do not make available the electrical consumption of the lamps, which is the quantity of interest.

Computers, TV apparatus,... We assume that various electronic appliances such as a television, a computer, etc. contribute to heating. It seems unlikely however to have access to any information related to their consumption or use.

Occupancy Finally, the last main source of heat inside a building is linked to the occupancy itself, with the metabolic heat gains emitted by each human being. These gains depend on the ongoing activity of the person, but we may first assume a constant gain of 50 W per person on average. This free heat gain depends then solely on the number of occupants of the building. Yet, estimating this number is a difficult task, and much research effort is devoted to this problem. The generic approach consists in choosing a set of on-site sensors, selecting the relevant features, and then training a classifier from real experiments, before evaluating and validating its performances. The task can either be the estimation of the number of people in a given room, or the separation between two states of the room, namely unoccupied and occupied. As an example, an entire line of work consists in infering the occupancy from a set of connected environmental sensors. For instance, [AAP+16] suggests to combine several sensors such as CO_2 concentration sensors, motion detection and power consumption, in order to robustly infer the number of occupants. Most methods in the literature have in common a supervised learning approach for solving this problem. This means that the real occupation of the room has to be recorded for some time, either via video recording, questionnaires, etc.

Non-heating loads

Hot water tank The consumption of the hot water tanks is a significant part of the aggregated load curve, but it does not represent a heat gain for the building, since in the end, the water is drained away. It is therefore necessary to have an estimation of this consumption, in order to substract it from the main load curve.

The first connected versions of the domestic hot water tanks on the market provide the user with a feedback on the daily consumption: the fine-grained consumption is still unavailable. The context of smart electric hot water tanks (EHWT) is studied more thoroughly in N. Beeker-Adda's thesis [Bee16]. In the latter reference, it is suggested that a "smart EHWT" should embed a flow meter and means for measuring and recording the injected power. However, such appliances are not available yet on the market.

As of today, it seems therefore difficult to specify the information a smart EHWT would bring. However, at a dynamic time step, the shape of the EHWT load curve is very specific: it is a rectangle of constant height equal to the nominal power. A typical load curve from individual housing is given in Figure 7.1, bottom left plot. This plot highlights the fact that the EHWT represents a non-negligible electric consumption, both in magnitude and frequency, which strenghtens the need for estimating it. Moreover, given the strong repetability of the draining patterns, which may in

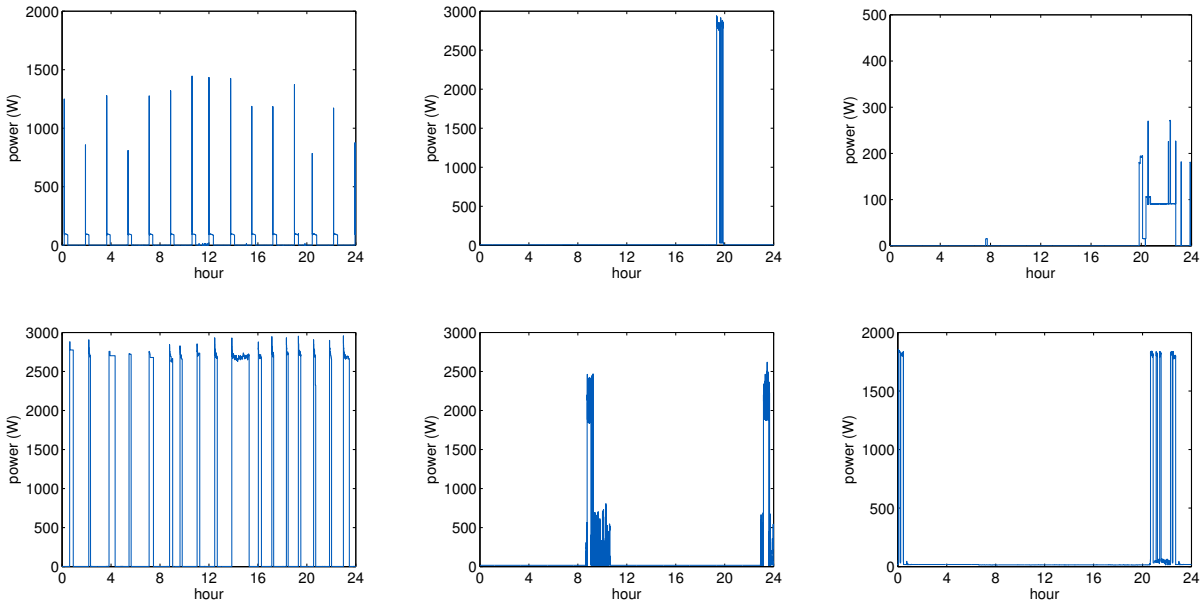

Figure 7.1: Typical domestic non-heating load curves at sampling time one second. From left to right, top to bottom: refrigerator, oven, lighting, electric hot water tank, washing machine, dish washer.

particular be synchronised with the electricity price signal, some algorithms have been designed internally at EDF R&D to successfully identify these consumptions, both in residential and office buildings. At sampling time 10 minutes, the HWT consumption presents more spikes, but some algorithms are also emerging, see e.g. [BdSACM16]. These algorithms should be even more efficient for those houses using both gas and electricity as energy vectors, since in such cases the gas meter measures the consumptions related to space heating and domestic hot water heating only, whereas the remaining loads are measured by the electricity meter.

We might therefore assume that even if no smart EHWT exists yet, an accurate estimation of the consumption is available in practice.

Washing machine, dish washer Washing machines and dish washers do not represent a heat gain in buildings, since the water is drained away, similarly to the hot water tank. Moreover, as shown in Figure 7.1, bottom center and right plots, these loads are also non-negligible, with nominal power of magnitude greater than 1 kW. Their use is also regular, with at least one use per week (depending on the family size). Connected washing machines and dish washers are emerging on the market. They essentially provide monitoring functions and remote control. Yet, no fine-grained consumption seems to be available at the moment. Furthermore, unlike the hot water tank, the electric consumption of the washing machine has different levels, as shown in Figure 7.1 (bottom center plot), corresponding to several functioning cycles. Extracting its consumption from the total load curve is therefore more challenging, and not discussed here. The random occurrence of starting times of the dish washer makes it also difficult to extract the related consumption.

Hybrid loads

Mechanical ventilation The ventilation in buildings is either natural (opening windows) or mechanical (driven by fans). We are interested here in the latter form. There are three types of mechanical ventilation: supply-only (not discussed here) / extract-only / supply and extract (balanced). In the extract-only case, the fan draws the air from the indoor spaces, whereas fresh air enters naturally through the openings in the building envelope. The extracted air represents a thermal loss that is accounted for separately in the model (not in the load curve but with the infiltrations instead), which means that it is not equivalent to a cooling system. We may thus consider that the measured consumption of the fan does not contribute to heating and has to be substracted from the main load curve. On the other hand, a balanced ventilation system uses a supply fan and an exhaust fan, both fans moving similar volumes of air, i.e. having the same power demand. The system includes moreover the possibility of heat recovery, with an exchanger (the extracted indoor air heats the supplied air). In that case, we may assume that the supplied air is at the same temperature as that of the indoor air. Consequently, the consumption of the extract fan needs to be substracted from the total load curve, but that of the supply fan may be considered as a heat gain, and kept in the total load curve. Therefore, half the consumption of the balanced ventilation system should be substracted from the total load curve.

Regarding the impact of the associated consumption, the power ratings range typically from 10W to less than 100W, depending on the speed of the fan. The ventilation system is a permanent use, but its load curve shows variability due to the variations of the speed of the fan. However, no "smart ventilation" system, able to report its consumption or functioning times for instance, seems to be available at the moment. Besides, from a practical point of view, we note that whenever the building is a block of flats, the consumption of the mechanical ventilation is not included in the counter of the flat, but in a counter common to the block instead.

Other connected objects

For the sake of completeness, let us mention three other objects that may bring information about the thermal state and energy consumption of the building.

- The recent development of **connected plugs** makes it possible for one to measure the specific consumption attached to a given plug. Although not generalizable to each and every plug of a building, because of the cost and the related privacy issues, it may be possible to monitor at least a specific appliance, such as a radiator for instance.
- There exist also connected **window/door contacts** such that it is possible to monitor the state (open vs closed) of any window or door inside the building. This opens the possibility to study time-varying models, instead of the time-invariant R3C2 network. See for instance the model suggested in [SBPW16].
- Finally, the **smartphone** is the most widely spread connected object. It may provide information about the position of the user, and serve as an interface between the different connected appliances of a building.

7.1.2 Main hypotheses

The hypotheses for taking into account the connected objects into the building identification framework are built from the above typology. First of all, we consider a building whose energy vector is exclusively gas and/or electricity, heating oil being therefore excluded. Secondly, the typology

shows that the current state of development of the technology does not allow one to assume that the non-heating loads are precisely known at each time, at a non-prohibitive cost. However, a more reasonable hypothesis would be to have access to a *timeline* of events. More particularly, we have seen that most non-heating loads are *recurring and irregular* loads, rather than permanent base loads. Hence, these loads may be first characterized by the time instants at which they are switched ON and OFF. The main hypothesis consists thus in assuming that the energy utility has access to this timeline of events. The global set of hypotheses used in this chapter is formulated below.

Observation with connected objects - main hypotheses

(H1) The aggregated total load curve Q is available at a dynamic resolution.

(H2) A timeline of events, specifying at which instants a non-heating load is either switched ON or OFF, is available.

(H3) The specific consumption of the hot water tank is either measured or accurately estimated.

(H4) The occupancy schedule and the metabolic heat gains are unknown.

(H5) The indoor and outdoor temperatures T_i[1]and ϑ_o and the solar gains Q_s are accurately measured.

(H6) The setpoint temperature T_r is also known.

(H7) The maximum power Q_h^{max} of the installed heating system is known.

7.2 Estimation procedure

7.2.1 Notations

Given the aforementioned hypotheses, we derive now an estimation procedure of the R3C2 network. Let us introduce some notations first. The measured total load curve, denoted Q, is at every moment the sum of the contributions of the useful heating loads Q_u and of the non-heating loads Q_{nh}. The useful load Q_u is the sum of two contributions, namely the regulated heating system, denoted Q_h, and the uncontrolled free heating gains, denoted Q_{fg}. Finally, the resulting heating power Q_r is the sum of the useful load Q_u and of the metabolic heat gains Q_m. Consistently with the notation used in the rest of this thesis, Q_r is thus the input of the R3C2 model. Hence, we observe Q while the quantity of interest is Q_r. Assuming that Q is measured between the time instants t_0 and t_f, $t_0 < t_f$, we have the relationships:

$$\forall t,\, t_0 \leq t \leq t_f, \quad \begin{cases} Q(t) &= Q_u(t) + Q_{nh}(t), \\ Q_u(t) &= Q_h(t) + Q_{fg}(t), \\ Q_r(t) &= Q_u(t) + Q_m(t). \end{cases} \tag{7.1}$$

[1]Remainder: T_i is the noisy observation of the model indoor temperature ϑ_i, ϑ_o is the model outdoor temperature. See Figure 2.1.

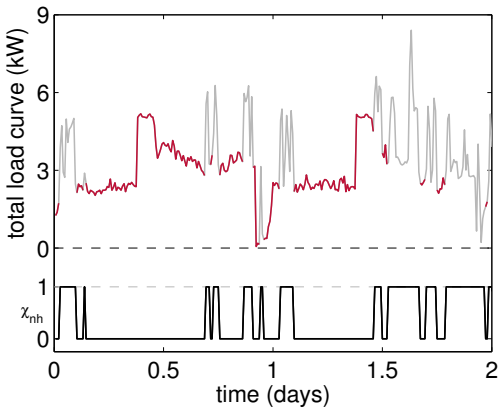

Figure 7.2: Measured total load curve and timeline of non-heating events. The non-heating events appear in light grey.

Let then χ_{nh} be the indicator function of the non-heating events:

$$\forall t,\, t_0 \leq t \leq t_f, \quad \chi_{nh}(t) = \begin{cases} 0 & \text{if } Q_{nh}(t) = 0 \\ 1 & \text{if } Q_{nh}(t) > 0 \end{cases} . \tag{7.2}$$

χ_{nh} is thus piecewise constant on the time interval $[t_0,\, t_f]$, as examplified in Figure 7.2 on simulated data - see Section 7.3 for the description of the data generating process. Considering that the continuous interval $[t_0,\, t_f]$ is discretized in T instants $\{t_i\}_{i=1}^{\mathrm{T}}$, with $t_1 := t_0$ and $t_{\mathrm{T}} := t_f$, $\{t_i\}_{i=1}^{\mathrm{T}}$ is the disjoint union $I_0 \sqcup I_1$ of two sets such that

$$I_0 := \{t_i, i = 1, \ldots, \mathrm{T} | \chi_{nh}(t_i) = 0\} = \bigsqcup_{j=1}^{n_0} I_{0j}, \tag{7.3}$$

$$I_1 := \{t_i, i = 1, \ldots, \mathrm{T} | \chi_{nh}(t_i) = 1\} = \bigsqcup_{j=1}^{n_1} I_{1j}, \tag{7.4}$$

with $n_0 + n_1 = \mathrm{T}$ and where each I_{0j} (respectively I_{1j}) is a set of consecutive instants for which the measured load curve contributes entirely (respectively not entirely) to heating. More particularly, we write for all $1 \leq j \leq n_0$, $I_{0j} = \{t_k\}_{k=k_j}^{k_j+l_j-1}$, where k_j (resp. l_j) is the first instant (resp. total number of consecutive instants) included in I_{0j}.

7.2.2 Overview of the algorithm

The estimation procedure is as follows. The first step consists in estimating the parameters θ of the R3C2 network from the data for which the measured load curve contributes totally to heating. This means that the EM algorithm described in Chapter 3 is adapted to the data corresponding to the time instants included in I_0. Let $\widehat{\theta}_1$ denote this estimator. In parallel with the estimation of $\widehat{\theta}_1$, we suggest as a second step to estimate the parameters η of the regulation model of the heating power Q_h. n_0 estimators $\widehat{\eta}^{(j)}$, $j = 1, \ldots, n_0$ are build, where each $\widehat{\eta}^{(j)}$ is computed exclusively from the measurements during the time interval I_{0j}, respectively. This step results thus in a set of estimators of the regulation sub-system. There are actually $K \leq n_0$ estimators $\widehat{\eta}^{(k)}$ in this set, because time invervals I_{0j} with small amounts of data are not taken into account.

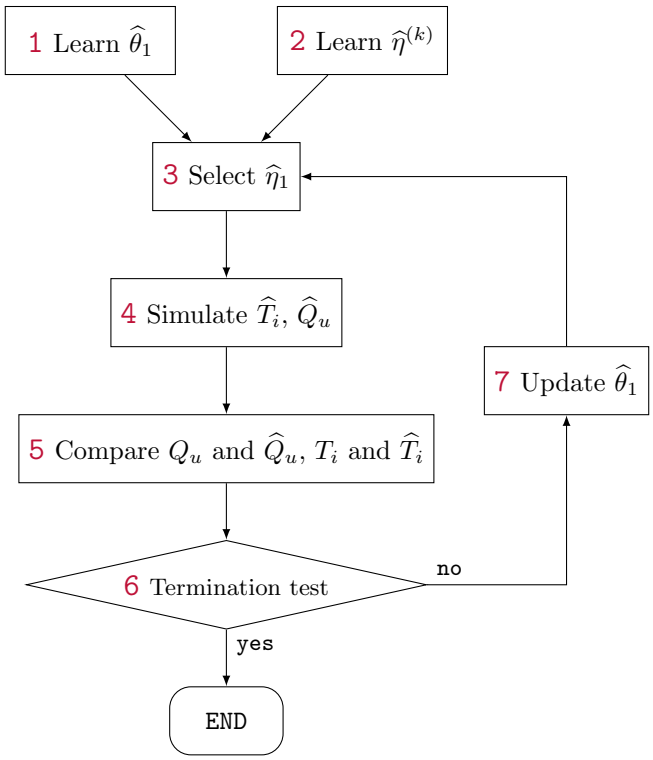

Figure 7.3: Flowchart of the learning procedure with an aggregated load curve.

The third step involves the selection, in some sense, of the "best" estimator $\widehat{\eta}_1$ amongst the $\widehat{\eta}^{(k)}$, $k = 1, \ldots, K$. This selection procedure depends also on the estimator $\widehat{\theta}_1$. Indeed, the indoor temperature being by hypothesis accurately measured at all times, the candidates $\widehat{\eta}^{(k)}$ are tested by simulation of the indoor temperature of the building and comparison to the actual measurements. This simulation depends thus on the estimation $\widehat{\theta}_1$ of the R3C2 network.

Then, the indoor temperature and heating power are simulated in closed loop from the parameters $(\widehat{\theta}_1, \widehat{\eta}_1)$. We compare the simulated trajectories of the indoor temperature and useful load with the observed ones. If this comparison satisfies a certain test, the algorithm ends with $(\widehat{\theta}_1, \widehat{\eta}_1)$ as the estimators of the building model. Otherwise, we suggest to update $\widehat{\theta}_1$, and then $\widehat{\eta}_1$, and repeat until satisfaction of the termination criterion.

A flowchart recapitulating this algorithm is given in Figure 7.3, whereas the different steps are described in detail in the next section.

7.2.3 Detailed description

Step 1/7: first identification of the R3C2 network

For each time interval I_{0j}, $j = 1, \ldots, n_0$ during which it is known that the measured load curve contributes entirely to heating, the inputs (heating power, solar gains, outdoor temperature) and output (indoor temperature) of the building model are supposed to be deterministic, i.e. known with certainty. This means in particular that the effect of the metabolic heat gains is neglected. Hence, we are in the standard case of a linear Gaussian state space model, and we may apply straightforwardly the EM equations, as in Chapter 3. More specifically, if $x_{I_{0j}} := \{x_k | k = k_j, \ldots, k_j + l_j - 1\}$, and $y_{I_{0j}} := \{y_k | k = k_j, \ldots, k_j + l_j - 1\}$ denote respectively the hidden states and observed output during

the time interval I_{0j}, the complete data log-likelihood reads

$$\begin{aligned}-2\ln L_{j,\theta}(x_{I_{0j}}, y_{I_{0j}}) &= \ln|\Sigma_{0j}| + (x_{0j}-\mu_{0j})^\intercal \Sigma_{0j}^{-1}(x_{0j}-\mu_{0j}) \\ &+ l_j \ln|Q| + \sum_{t=k_j}^{k_j+l_j-1} (x_t - Ax_{t-1} - Bu_t)^\intercal Q^{-1}(x_t - Ax_{t-1} - Bu_t) \\ &+ l_j \ln|R| + \sum_{t=k_j}^{k_j+l_j-1} (y_t - Cx_t)^\intercal R^{-1}(y_t - Cx_t), \end{aligned} \tag{7.5}$$

and the associated cost is $\mathbb{E}\left[-2\ln L_{j,\theta}(x_{I_{0j}}, y_{I_{0j}}) \middle| y_{I_{0j}}\right]$. See Algorithm 3.2 for solving this E-step. Adding the n_0 contributions, the estimator $\hat{\theta}_1$ of the R3C2 network is found by minimizing the total cost $J_1(\theta)$, where

$$J_1(\theta) := \sum_{j=1}^{n_0} \mathbb{E}\left[-2\ln L_{j,\theta}(x_{I_{0j}}, y_{I_{0j}}) \middle| y_{I_{0j}}\right]. \tag{7.6}$$

Estimating the parameters θ is thus very similar to the standard case, except that we need to apply the Kalman smoother n_0 times (with data of length l_j instead of T). The minimization step can be achieved with a standard unconstrained optimization solver, such as `fminunc` from Matlab.

Step 2/7: build a set of PI estimators

We consider a regulation model of the indoor temperature that takes the form of a Proportional-Integral (PI) controller with saturation at 0 and Q_h^{max} and anti-windup strategy, where Q_h^{max} is the limiting power of the heating system. In the following, it is assumed that Q_h^{max} is known, as well as the anti-windup gain. It remains thus to estimate the gain z_K and the time constant τ of the PI controller. To do so, we consider each interval I_{0j}, $j = 1, \ldots, n_0$ separately. Eliminating those intervals that contain less than 2 hours of data, the remaining intervals are re-labeled I_{0k}, $k = 1, \ldots, K_0$ for some $K_0 \leq n_0$. For the time instants included in each I_{0k}, the measured load curve contains only heating loads. However, the proportion of free gains, i.e. uncontrolled heating power, is not known. As a first approximation, we assume that the measured load curve is the regulated heating power. The estimation of $z_K^{(k)}$ and $\tau^{(k)}$ on each interval I_{0k}, $k = 1, \ldots, K$ uses a simple least squares formulation. If Q_t denotes the load curve at time $t \in I_{0k}$ and $\varepsilon_t := T_r^t - T_i^t$ the difference between the setpoint and measured temperatures, we have *during the linear regime* of the controller:

$$Q_t = Q_{t-1} + z_K^{(k)}\left(\varepsilon_t - \varepsilon_{t-1} + \frac{1}{\tau^{(k)}}\varepsilon_t\right). \tag{7.7}$$

In practice, we estimate $z_K^{(k)}$ and $z_\tau^{(k)} := 1/\tau^{(k)}$ by ordinary least squares, using the data measured during I_{0k} and such that they verify $\alpha < Q_t < (1-\alpha)Q_h^{max}$, $t \in I_{0k}$, for some small $\alpha > 0$, in order to make sure that the controller is not saturated. For instance, we may use $\alpha = 5\%$ or $\alpha = 10\%$. Again, we restrict ourselves to the intervals I_{0k} that contain sufficiently many data. Hence, we obtain a set of $K \leq K_0 \leq n_0$ estimators $\widehat{\eta}^{(k)} := (z_K^{(k)}, z_\tau^{(k)})$ of the PI controller.

Step 3/7: EWA aggregation strategy

In order to select the best estimator from the set of the estimated $\widehat{\eta}^{(k)}$, $k = 1, \ldots, K$, we suggest here to apply an aggregation strategy of the expert advice, where the experts are $\widehat{\eta}^{(k)}$, $k = 1, \ldots, K$.

Indeed, each $\widehat{\eta}^{(k)}$ is biased, with bias depending in particular on the unknown proportion of free heat gains in the load curve measured at the time instants in I_{0k}. By using an aggregation strategy, the aim is to approach the performances of the (unknown) best, least unbiased, estimator. An introduction to the context of prediction with expert advice is provided in the reference textbook [CBL06]. Moreover, we refer also to P. Gaillard's thesis, [Gai15], which contains several contributions to online robust aggregation and applications to electricity load forecasting. Briefly, the formalism is the following. The purpose is to predict any arbitrary time series $y_t \in \mathbb{R}^d$, $t = 1, \dots, \mathrm{T}$, based on available expert forecasts. At each time t, the statistician has access to the past observations $y_1, \dots, y_{t-1}$, as well as to the current and past predictions $\{x_{k,s}\}_{s=1}^{t}$ of the K experts. Note that a key feature of this framework is that it does not rely on any stochastic assumption on the time series y_t and $x_{k,t}$: they are arbitrary. A typical aggregation strategy solves the prediction problem by computing a weighted average $\widehat{y}_t := \sum_{k=1}^{K} w_{k,t} x_{k,t}$ of the expert forecasts. Then, the statistician observes the actual y_t, which enables the sequential computation of the cumulative loss, or its average $\widehat{L}_{\mathrm{T}} := \frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \ell(\widehat{y}_t, y_t)$, with $\ell : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ the chosen nonnegative loss function. Aggregation strategies aim at minimizing this cumulative loss. One such strategy is the Exponentially Weighted Average (EWA), presented in [LW94, CBL06, Gai15], for which the weights are computed according to

$$w_{k,t} := \frac{\exp\left(-\rho \sum_{s=1}^{t-1} \ell(x_{k,s}, y_s)\right)}{\sum_{j=1}^{K} \exp\left(-\rho \sum_{s=1}^{t-1} \ell(x_{j,s}, y_s)\right)}, \tag{7.8}$$

where ρ is a learning parameter, called learning rate, that has to be tuned.

In our application, the K experts are therefore the PI models $\widehat{\eta}^{(k)}$ learned at Step 2., and the prediction $x_{k,t}$ is the indoor temperature obtained at time t in closed loop simulation of the building model $\widehat{\theta}_1$ with regulation $\widehat{\eta}^{(k)}$. The generation of this estimation is the same as the one described in Step 4. However, here, we are mostly interested in the *aggregated model*, rather than the prediction of the indoor temperature. Nevertheless, the standard framework of aggregation of forecasts does not assume any model on the experts, since they are seen as black boxes. This problem is theoretically involved, and our aim is not to solve it analytically in this work. Instead, we suggest two "natural" *ad hoc* solutions. Both are based on the EWA strategy:

1. aggregation of the predictions, and apply the weights $w_{t,k}$ to each component of $\widehat{\eta}^{(k)}$;
2. aggregation of the models directly, i.e. at each time instant t, the weights $w_{t,k}$ affect the model $\widehat{\eta}^{(k)}$ directly: the predicted indoor temperature is the prediction of an exponentially weighted average model $\widehat{\eta}$.

The computational difference between the two strategies is emphasized in Algorithm 7.1. The first approach is the standard EWA strategy, where the weights are subsequently applied to the parameters of the expert models. Should the model be linear in its parameters, this approach would be exact. Here, the model is nonlinear in the parameters, which is why this first method is an approximation. In the second approach, on the other hand, the weights are applied directly to the experts $\widehat{\eta}^{(k)}$ before generating the prediction $\widehat{y}_t$, which is a different context from the one studied in [CBL06, Gai15]. Hence, the nice performance bounds of the standard algorithms, not discussed here, are not available anymore *a priori*.

Both approaches result in an aggregated model $\widehat{\eta}_t$. The index t emphasizes the time-dependency of the solution, since the weights are updated sequentially in time. In order to form a time-invariant model $\widehat{\eta}_1$, we propose to use the model obtained once the mixture of weights has stabilized. In practice, we compute the temporal mean of the $\widehat{\eta}_t$ during the last n instants, where n is determined manually.

Algorithm 7.1 Aggregation of the PI models

Inputs: Experts $\{\eta^{(k)}\}_{k=1}^{K} = \{z_K^{(k)}, z_\tau^{(k)}\}_{k=1}^{K}$, R3C2 model $\widehat{\theta}$, EWA learning rate ρ, n, loss function ℓ, observations $\{\vartheta_0(t), Q_s(t), T_i(t), T_r(t)\}_{t=1}^{\mathrm{T}}$, limiting power Q_h^{max}, anti-windup.

1. For each k, simulate $\widehat{T}_i^{(k)}$ from the system $(\widehat{\theta}, \eta^{(k)})$ and the observations $\{\vartheta_0(t), Q_s(t), T_r(t)\}_{t=1}^{\mathrm{T}}$.
2. Initialize $w_{k,1} = 1/K$, $k = 1, \ldots, K$.
3. For $t = 1, \ldots, \mathrm{T}$:
   Strategy 1
   (a) prediction: compute $\widehat{T}_i(t) := \sum_{k=1}^{K} w_{k,t} \widehat{T}_i^{(k)}(t)$;
   (b) aggregated model: set $\widehat{z}_K(t) := \sum_{k=1}^{K} w_{k,t} z_K^{(k)}(t)$ and $\widehat{z}_\tau(t) := \sum_{k=1}^{K} w_{k,t} z_\tau^{(k)}(t)$;
   (c) weight update: for each k, compute $w_{k,t+1} \propto \exp\left(-\rho \sum_{s=1}^{t} \ell(\widehat{T}_i^{(k)}(s), T_i(s))\right)$.

   Strategy 2
   (a) aggregated model: set $\widehat{z}_K(t) := \sum_{k=1}^{K} w_{k,t} z_K^{(k)}(t)$ and $\widehat{z}_\tau(t) := \sum_{k=1}^{K} w_{k,t} z_\tau^{(k)}(t)$;
   (b) prediction: simulate $\widehat{T}_i(t)$ from the system $(\widehat{\theta}, \widehat{\eta}(t))$ where $\widehat{\eta}(t) := (\widehat{z}_K(t), \widehat{z}_\tau(t))$;
   (c) weight update: for each k, compute $w_{k,t+1} \propto \exp\left(-\rho \sum_{s=1}^{t} \ell(\widehat{T}_i^{(k)}(s), T_i(s))\right)$.
4. Time average of the aggregated model: compute $\widehat{z}_K := \left(\sum_{\mathrm{T}-n+1}^{\mathrm{T}} \widehat{z}_K(t)\right)/n$ and $\widehat{z}_\tau := \left(\sum_{\mathrm{T}-n+1}^{\mathrm{T}} \widehat{z}_\tau(t)\right)/n$, set $\widehat{\eta} := (\widehat{z}_K, \widehat{z}_\tau)$.

Outputs: aggregated PI parameters $\widehat{\eta}$.

Step 4/7: Simulated trajectories of the indoor temperature and heating power

Once $\widehat{\theta}_1$ and $\widehat{\eta}_1$ have been estimated, the fourth step requires the simulation of the indoor temperature and heating power under the model hypothesis. Similarly to Step 2., it is assumed that the limiting heating power Q_h^{max} and the anti-windup gain z_{ks} are known. Then, the indoor temperature $\widehat{T}_i$ and heating load curve $\widehat{Q}_u$ are simulated from a given setpoint temperature, the measured outdoor temperature and solar radiation, and free internal heat gains Q_{fg} constantly set to 0 since they are unknown. As such, the estimated load curve $\widehat{Q}_u$ is an equivalent useful load. Generally speaking, one expects that at time $t \in \mathcal{I}_0$, the order relation $\widehat{Q}_u(t) > Q_h(t)$ holds, except if there is saturation at Q_h^{max}, since the simulated useful load curve should compensate for the unknown free gains.

Step 5/7: Accuracy of the simulated trajectories

The accuracy of the simulated indoor temperature may be evaluated by computing the root mean square error (RMSE) alongside the entire trajectory. It is classically defined by the quantity

$$\mathcal{E}_{ti} := \sqrt{\frac{1}{\mathrm{T}} \sum_{t=1}^{\mathrm{T}} \Big(T_i(t) - \widehat{T}_i(t) \Big)^2}. \tag{7.9}$$

Regarding the useful load, the RMSE may be computed by restricting the comparison to the time instants included in the set

$$\mathcal{T} := \{ t | t \in I_0,\, Q(t) \leq Q_h^{max} \}, \tag{7.10}$$

which yields the following defintion

$$\mathcal{E}_{qu} := \sqrt{\frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \Big(Q(t) - \widehat{Q}_u(t) \Big)^2}, \tag{7.11}$$

where $|\mathcal{T}|$ is the cardinality of the set $\mathcal{T}$. Indeed, the first condition is that $t \in I_0$, such that the measured load curve $Q(t)$ contains only heating loads, i.e. $Q(t) = Q_u(t)$. Moreover, we have necessarily $\widehat{Q}_u(t) \leq Q_h^{max}$ at any time t, such that the comparison is not relevant whenever $Q_u(t) = Q(t) > Q_h^{max}$. In other words, if $t \in I_0$ and $Q(t) > Q_h^{max}$, we know that $Q(t)$ contains uncontrolled free heat gains that cannot be accounted for by any model $\{\widehat{\theta}_1, \widehat{\eta}_1\}$.

Step 6/7: Termination condition

For the evaluation of the convergence of the algorithm, we check whether or not the quantity

$$\delta = \frac{\Big| \Big(\mathcal{E}_{ti}(\widehat{\theta}_1, \widehat{\eta}_1) + \mathcal{E}_{qu}(\widehat{\theta}_1, \widehat{\eta}_1) \Big) - \Big(\mathcal{E}_{ti}(\widehat{\theta}_2, \widehat{\eta}_2) + \mathcal{E}_{qu}(\widehat{\theta}_2, \widehat{\eta}_2) \Big) \Big|}{\mathcal{E}_{ti}(\widehat{\theta}_1, \widehat{\eta}_1) + \mathcal{E}_{qu}(\widehat{\theta}_1, \widehat{\eta}_1)} \tag{7.12}$$

falls below a certain threshold, where we have emphasized the dependency of $\mathcal{E}_{ti}$ and $\mathcal{E}_{qu}$ on the parameters of the system, while $\widehat{\theta}_2$ and $\widehat{\eta}_2$ are the new parameters defined in the next step. This criterion takes thus into account both stabilizations of the trajectories of the indoor temperature and heating power.

Step 7/7: Potential update of $\widehat{\theta}_1$

If the termination condition is not satisfied, we propose to update the estimation $\widehat{\theta}_1$ of the R3C2 model according to the following scheme. The new parameters $\widehat{\theta}_2$ are found by minimizing the cost $J_2(\theta)$, where

$$J_2(\theta) := \lambda_1 \frac{\mathcal{E}_{ti}(\theta, \widehat{\eta}_1)}{\bar{T}_i} + \lambda_2 \frac{\mathcal{E}_{qu}(\theta, \widehat{\eta}_1)}{\bar{Q}_u} + \frac{1}{2} \left|\left| \theta - \widehat{\theta}_1 \right|\right|^2_{\Lambda_3^{-1}}, \tag{7.13}$$

where $\mathcal{E}_{ti}$ and $\mathcal{E}_{qu}$ are defined respectively in (7.9) and (7.11), $\bar{T}_i$ (resp. $\bar{Q}_u$) is the mean value of the observed series T_i (resp. Q_u, computed on $\mathcal{T}$), $||\cdot||$ denotes the norm $||x||^2_\Lambda := x^\intercal \Lambda x$ for any vector x and matrix Λ of appropriate dimensions. Hence, the first two terms are the normalized root mean squared errors computed respectively for $\widehat{T}_i$ and $\widehat{Q}_u$. Note that the computation of $\mathcal{E}_{ti}$ and $\mathcal{E}_{qu}$ requires the closed loop simulation of the system $(\theta, \widehat{\eta}_1)$ according to the procedure described in Step 4.

The third term is a regularization term, that can be seen as a prior density in a Bayesian context: the values of θ that do not lie a neighborhood of the previous $\widehat{\theta}_1$ are penalized. Consistently with this idea, we choose a diagonal matrix Λ_3, with diagonal elements $\lambda_{3,i} = (\alpha \widehat{\theta}_{1,i})^2$, with for instance $\alpha = 0.2$. Finally, the scalars λ_1 and λ_2 are also tuned manually. We set $\lambda_1 = \lambda_2 = 100$ such that $\lambda_1 \frac{\mathcal{E}_{ti}(\theta, \widehat{\eta}_1)}{\bar{T}_i}$ and $\lambda_2 \frac{\mathcal{E}_{qu}(\theta, \widehat{\eta}_1)}{\bar{Q}_u}$ are expressed in %.

This cost $J_2(\theta)$ may be minimized by an unconstrained solver, such as Matlab's `fminsearch`, which uses a simplex, derivative-free, search method.

Once $\widehat{\theta}_2$ is computed, the next step consists in updating the new PI parameters $\widehat{\eta}_2$, according to the aggregation procedure of Step 3. The algorithm is then iterated until convergence, as defined by Step 6. The detailed description of the estimation procedure suggested in this chapter is summarized in Algorithm 7.2.

7.3 Stochastic generation of domestic load curves

In the previous chapters, we have used a hybrid dataset, made of ground measurements of the meteorological conditions and of the free internal gains on one hand, and of closed loop simulations of the indoor temperature and of the heating power on the other hand. This procedure is described in Chapter 2. However, this hybrid dataset is not entirely suitable to the numerical exploration of the methods suggested in this chapter. Indeed, its total load curve is always the sum of two contributions only, namely the heating system consumption (controlled power) and the aggregated consumptions of the laptops of an office building (free heating gains). The main drawback is therefore that both consumptions do contribute to heating while there is no non-heating load: $Q(t) = Q_u(t)$ at all times.

In this section, we suggest thus an alternative way of generating a simulated electric demand in an individual *domestic* building. Our purpose is to be able to generate a domestic load curve such that

1. it contains heating and non-heating loads;
2. it is representative, "in some sense", of the behaviour of an actual household;
3. the generating process accounts for the diversity of the patterns of electricity use in an indivual dwelling.

Hence, the method should in particular be stochastic. The idea behind the notion of representativity is that the respective proportions of the heating and non-heating loads as well as the total load curve

Algorithm 7.2 Estimation of a R3C2 model with uncertain heating load curve

Inputs: observed series $\{Q(t), \vartheta_0(t), Q_s(t), T_i(t), T_r(t), I_0(t)\}_{t=1}^{\mathrm{T}}$, limiting power Q_h^{max}, anti-windup effect z_{ks}, EWA learning rate ρ, threshold δ_0.

1. Set $l = 1$, estimate the parameters $\widehat{\theta}_l$ of the R3C2 network according to an EM procedure, with $\widehat{\theta}_l = \arg_\theta \min J_1(\theta)$, $J_1(\theta)$ given in (7.6).
2. Compute K estimators $\widehat{\eta}^{(k)}$, $k = 1, \ldots, K$ of the parameters of the PI controller on the intervals I_{0k} with sufficiently many data, using an ordinary least squares algorithm based on (7.7).
3. Compute $\widehat{\eta}_l$, aggregated from the K experts $\widehat{\eta}^{(k)}$, using an exponentially weighted average procedure that depends on $\widehat{\theta}_l$, with learning rate ρ, see Algorithm 7.1.
4. Simulate $\widehat{T}_i$ and $\widehat{Q}_u$ in closed loop from the system $(\widehat{\theta}_l, \widehat{\eta}_l)$, using the observed series ϑ_0, Q_s, T_r as well as null free gains.
5. Compute the root mean square errors of $\widehat{T}_i$ and $\widehat{Q}_u$ according to (7.9) and (7.11).
6. Termination condition: compute the quantity δ defined in (7.12)
    7. if $\delta > \delta_0$, where δ_0 is a chosen threshold, then set $l \leftarrow l + 1$, update $\widehat{\theta}_l$ by minimizing the cost function $J_2(\theta)$ given in (7.13), and return to step 3;
    8. otherwise, set $\widehat{\theta} = \widehat{\theta}_l$, $\widehat{\eta} = \widehat{\eta}_l$ and terminate.

Outputs: estimated R3C2 parameters $\widehat{\theta}$, estimated PI parameters $\widehat{\eta}$.

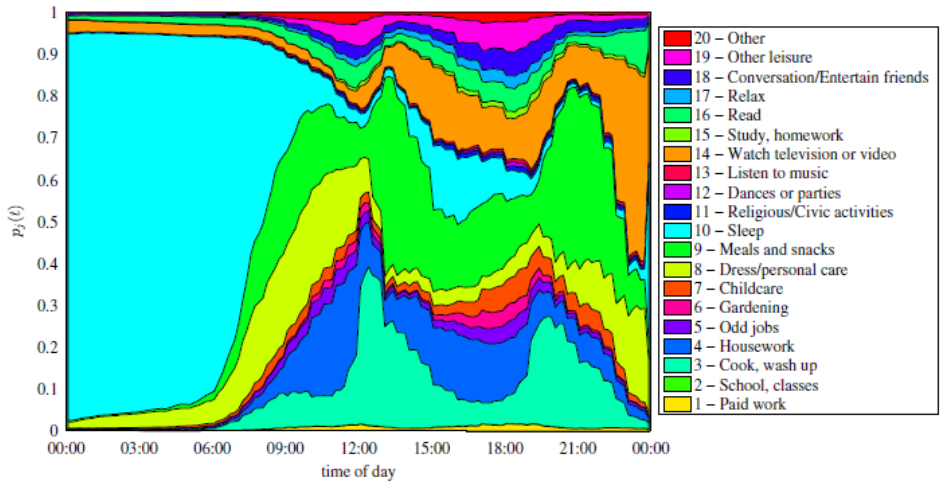

Figure 7.4: Probability profile of the merged activities whilst people are at home, French TUS 1998, screenshot of [WHSR13, Figure 1].

shape should be "reasonable" and correspond to likely situations. A literature review reveals that recent research efforts are currently suggesting such stochastic models. Indeed, it is well-known that there is a gap between the estimated energy performances of a building at the conception phase and those actually observed, the latter potentially largely exceeding the former. This gap is partly explained by the simplistic way the occupants behaviour is accounted for in the building simulation tools. New approaches, based on a stochastic representation of the occupants behaviour are therefore being proposed [TPH+]. An extensive treatment of the topic, as well as detailed models, are provided in E. Vorger's thesis [Vor14]. Recent contributions include amongst others the work of J. Widén *et al.* [WLV+09, WW10, WME12], J. Page *et al.* [PRMS08] and U. Wilke [WHSR13]. A review is also provided in [Tor14].

In this thesis, we choose to reproduce the work of I. Richardson *et al.*, Center for Renewable Energy Systems Technology (CREST), Loughborough University. This choice is motivated by the simplicity of their approach, which is consistent with our purpose, and the open-source availability in [RT10] of their model and data. The model consists of several submodels, the first one aiming at generating a daily active occupancy pattern of the dwelling, which we describe next.

7.3.1 Domestic building occupancy model

A Markov chain based method for generating daily active occupancy data is presented in [RTI08]. For a given household composition, with $n \in \{1, \dots, 6\}$ persons, the chain at time t of the day has $n+1$ possible states (empty house, 1 active occupants, ..., n active occupants). The model has two states per occupant: at home and active (state 1) and not at home or at home and not active (state 0). A time-dependent transition probability matrix is therefore computed at each time t of the day, with 10 minutes increments. The start state is generated at midnight with relevant probability distribution (i.e. most likely asleep). Then, the subsequent states till $t = 23{:}50$ are generated from the inhomogeneous Markov chain previously defined.

The construction of the transition matrices relies on the statistics of the Time Use Survey (TUS). The TUS is a large survey conducted periodically on a country scale by the respective National Statistical Institutes of the European Union countries. It is based on more than 10 000 households completing questionnaires that describe their daily activities at time-step 10 minutes,

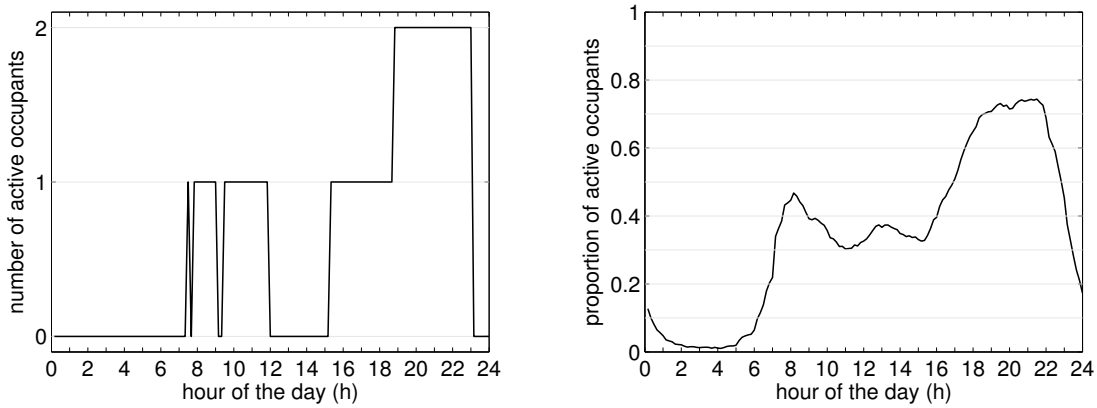

(a) 1 household (b) 2000 households

Figure 7.5: (a) Simulated active occupancy profile of a two-resident household on a weekday. (b) Simulated weekday aggregated active occupancy profile on 2000 two-resident households.

where these activities are assigned from a predefined list. When merging the activities by type, one obtains a probability profile such as the one depicted in Figure 7.4 (French TUS, 1998). By aggregating then the activities, one can retrieve the information of active occupancy in the house for each individual in the TUS database. Moreover, the TUS database indicates which people live in the same house, leading to an active occupancy pattern per day per household. This accounts for the correlation between the patterns of different individuals in the same household. Hence, the final dataset consists of the daily active occupancy profiles per household, with a distinction between weekdays and weekends. The data are aggregated according to the size of the household, which enables one to estimate the empirical transition matrices of the Markov chain.

The model is calibrated using the Time-Use Survey conducted in 2000 in the United Kingdom. The validation procedure in [RTI08] shows that the model is capable of generating series with statistics consistent with the TUS data, while being computationally efficient. An example of a simulated weekday active occupancy profile of a two-resident household is given in Figure 7.5(a), and the average over 2000 independent simulations in Figure 7.5(b). A limitation of the method is that it is based on the TUS data which provides information for individual days only, and not consecutive days. Hence the model might be used to generate occupancy patterns of multiple days, but without taking into account the likely correlations between two consecutive days. Furthermore, we mention an extension of the model to a four-state Markov chain, (at home and active - at home and not active - not at home and active - not at home and not active) provided in [MKT15]. This refinement enables one to account more precisely for the metabolic heat gains, through the state "at home and not active".

7.3.2 Generation of domestic electric load curves

Once the occupancy profile of a given dwelling is randomly generated, it remains to determine the electric demand related to the activities of the occupants. We begin with a specific submodel for the lighting demand.

Lighting load curves

The generation of domestic lighting load curves at a one-minute resolution is described in [RTID09]. Consider a given dwelling with n occupants. In order to simulate its demand related to the use of electric lighting, the main steps are:

1. *Preliminary steps:*
   - determine the active occupancy pattern, as in Section 7.3.1, and disaggregate to a one-minute time-step;
   - allocate the lighting units to the dwelling. This is based on some statistics (mean number of installed bulbs, relative proportion of the different technologies) provided by The Lighting Association. In [RTID09], the number of units is first randomly assigned, before randomly determining the associated technologies;
   - provide measurements of the received solar irradiance. Moreover, it is assumed that the occupants are likely to switch the light on below a certain level of outdoor irradiance. A threshold normally distributed around 60 W/m^2 is suggested.
2. *Switch-on event:* at each time step, i.e. every minute, compute the probability of a switch-on event. See [RTID09] for the computations. The model is such that a switch-on event cannot occur if the house is empty or the outdoor irradiance level too high. It accounts for the fact that several people in the same house will share the lighting demand, such that the demand is not proportional to the number of active occupants. A relative weight is also computed, such that some units are used more frequently than others (e.g. kitchen vs cellar);
3. *On-duration:* each time a switch-on event occurs, a random duration is drawn from an empirical distribution provided by [SRL04].

Other loads

The generation of one-minute resolution domestic electric load curves is described in [RTIC10]. It consists of generating specific load curves for each use, and aggregating these consumptions in order to obtain the total load. Similarly to the case of lighting demand, each dwelling is randomly assigned an active occupancy pattern and a set of appliances among a list of 33. The probabilities of allocation are determined according to some national statistical ownership rate data. In [RTIC10], these statistics are taken from the UK department of Energy and Climate Change (2009), the UK Market Transformation Program (2008), the Environmental Change Institute, University of Oxford, UK (2000, 2005) and the UK Ofcom (see cited references in [RTIC10]). The remaining steps are as follows:

1. *annual energy use:* determine the mean annual energy use per appliance, using the statistics of the aforementioned references;
2. *power characteristics:* each appliance is either *on* or *off*. At the time resolution of 1 minute, the daily profile is either constant during cycles or time-varying. In [RTIC10], the only time-varying load curve is that of the washing machine, where the profile is based on some measured data.
3. *appliance-activity mapping:* each appliance is mapped to one of the daily activity profiles in the Time Use Survey database. Several appliances may be mapped to the same activity. Seven activities are included in the original paper [RT10], namely "sleeping", "watching TV", "cooking", "doing laundry", "washing or dressing", "ironing" and "house cleaning".

Algorithm 7.3 Stochastic generation of a domestic load curve

Inputs: number of days D, number of occupants n, weather conditions, temperature setpoint, true R3C2 parameters.

1. Allocate a set of appliances and lighting units to the dwelling.
2. For each day d in $1, \ldots, D$,
   (a) determine the active occupancy pattern, e.g. as in [RTI08];
   (b) compute the lighting demand, e.g. as in [RTID09];
   (c) compute the appliance-specific demand, e.g. as in [RTIC10].
3. Concatenate the D days and disaggregate the demand between heating and non-heating loads.
4. Compute the metabolic gain by counting 50W per active occupant.
5. Generate the heating power from the closed loop simulation of the building system, with the sum of the heating loads and of the metabolic gains as free internal gains.

Outputs: disaggregated domestic load curve, indoor temperature.

4. *switch-on events:* the probability of activation of each appliance at time time t is the product of the activity probability at time t (determined from the associated activity profile) and an appliance-specific calibration scalar. The calibration scalar determines the average number of times an appliance is used in a year. It is adjusted such that over a very large number of stochastic simulation runs, the mean annual consumption of the appliance will be according to the national statistics.
5. *demand:* when a switch-on event occurs, use the power characteristics of the appliance to determine its electricity demand.

This method results thus in 33 one-minute appliance-specific daily load curves per dwelling, while adding these loads gives the total power consumption, for one day.

Heating system

The heating system is not fully represented in the model [RTIC10], which is why its power consumption is computed separately in our implementation. The method is the same as the one described in Chapter 2, namely:

1. generate the lighting and other appliances demand, and add the specific load curves of those appliances that contribute to heating;
2. select the true values of the five parameters of the R3C2 network;
3. choose a temperature setpoint signal and simulate the indoor temperature and the heating power from the closed-loop building system, using the free heating gains obtained at step 1, as well as the actual measurements of the outdoor temperature and solar gains.

This step concludes the overall data generating process, which is summarized in Algorithm 7.2.

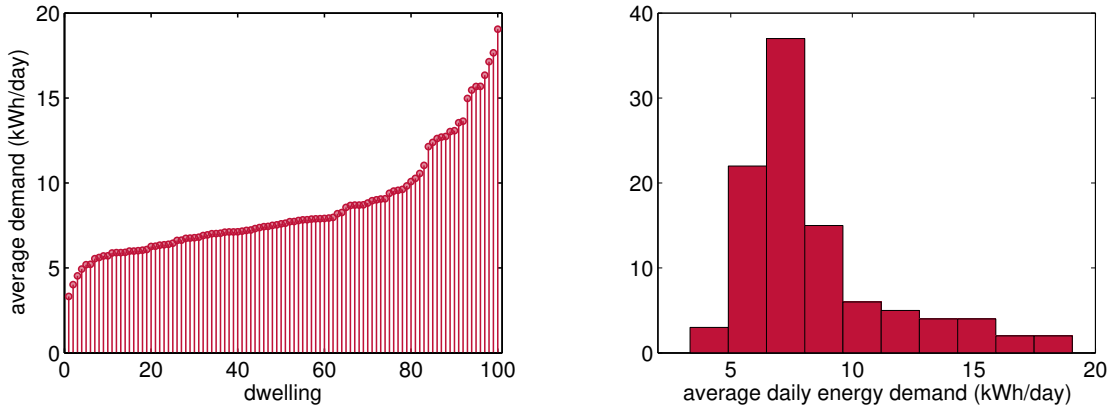

(a) (b)

Figure 7.6: Simulated average daily electricity demand of 100 two-occupant dwellings (a) ranked in increasing order (b) grouped in a histogram. The heating power and electric hot water tank are not included in the demand.

7.3.3 Validation of the model

Some numerical experiments are conducted in order to verify that the stochastic generation of domestic load curves presented in Algorithm 7.3 is consistent with the main issues raised in Section 7.3. The second step in Algorithm 7.3 is carried out using our own translation from Excel-VBA to Matlab of the freely available code in [RT10]. In particular, the power ratings of the appliances are drawn for each dwelling around the nominal values used in [RT10], which increases the variability between dwellings. Moreover, since it is assumed in this chapter that the consumption of the hot water tank may be reliably estimated, we do not include this appliance in the list of equipments of the building. The domestic specific load curves of 100 two-person dwellings are simulated for one week. The stochastic nature of this process is shown for instance by computing the mean daily total electricity demand, as in Figure 7.6. This hints thus at the fact that the generating process does account for the diversity of the patterns of electricity use in an individual dwelling. Note that the shape of Figure 7.6(a) is similar to the results described in [RTIC10].

Figure 7.7 illustrates then the difference between the simulated heating and non-heating loads, where the heating load is the aggregation of the loads of the appliances that contribute to heating, excepting the heating system which is not simulated yet. Satisfactorily for our application, the model is capable of generating non-heating and heating loads both of significant contribution to the total load curve. The more fluctuant nature of the non-heating loads is also visible on Figure 7.7(a), while it appears that the heating load on the other hand is the sum of cyclic contributions and punctual events.

Finally, the demand of the heating system is simulated three times for each of the 100 datasets, with the same meteorological conditions and setpoint temperature, but different nominal values of the parameters of the R3C2 network. In particular, the first R3C2 model has a nominal UA value of 0.18, the second 0.25 and the third 0.37, corresponding to increasingly poorly insulated buildings. Figure 7.8 shows the repartition of the average daily heating demand for each of the three experiments. It can be noted that regardless of the experiment, this demand varies highly depending on the dwelling, between 22.0 kWh/day and 26.9 kWh/day for the first experiment, 37.8 kWh/day and 46.1 kWh/day for the second, 58.9 kWh/day and 65.9 kWh/day for the third. This

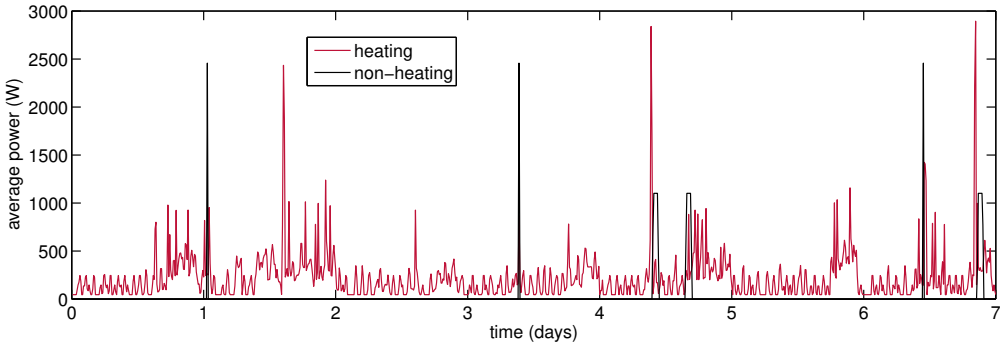

(a)

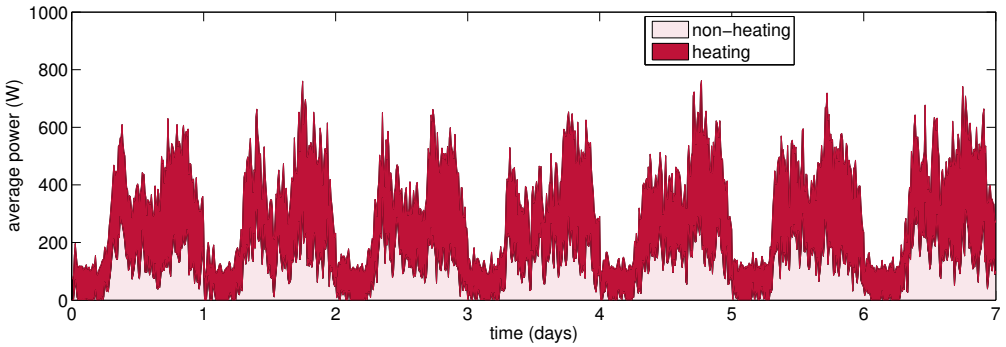

(b)

Figure 7.7: (a) Disaggregation of the demand of a dwelling between heating and non-heating loads, and (b) average over 100 dwellings. The heating power and electric hot water tank are not included in the demand.

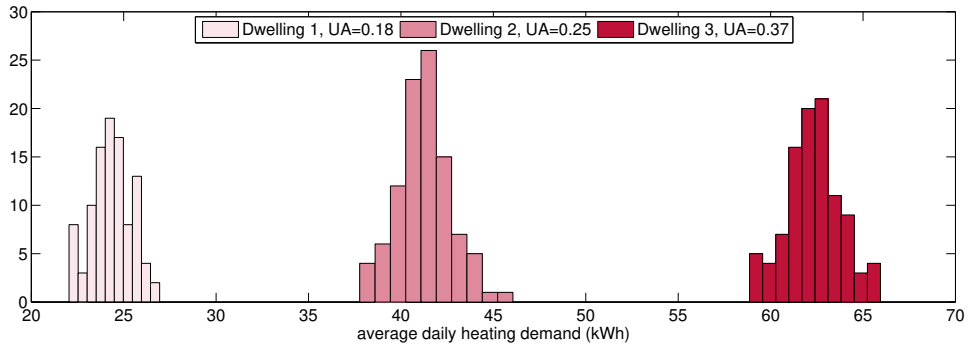

Figure 7.8: Mean daily heating demand for three dwellings, each with 100 occupancy scenarios.

discrepancy may only be explained by the specific appliance demands of the dwellings, since the other factors are all equal. Nevertheless, the three histograms do not overlap, which means that the level of insulation of the building defined by the UA value is key to the heating demand. These two remarks confirm the good behaviour of the data generating process defined by Algorithm 7.2, since they yield load curves where the dominant contribution is that of the space heating system, consistently with the national statistics, see e.g. [ADE13], and where the user behaviour explains also the energy performances of the building.

As such, the data generating process is adequate to our application. It enables one to generate domestic load curves at a dynamic time-step of one minute, and covers a significant amount of appliances. As highlighted by the numerical experiments, its intrinsic stochasticity produces diversified patterns of electricity use. Although unnecessary for our application, the level of diversity might be further increased, for instance by taking the socio-demographic characteristics of the dwellings into account as discussed and suggested in [Vor14]. A key feature of the model is to separate the TUS data according to the number of occupants in the active occupancy submodel so as to represent the correlated use of appliances. One improvement, on the other hand, would be to allow for correlations between consecutive days instead of having static parameters and independent simulations.

Moreover, the approach is representative of the behaviour of an actual dwelling, because it is based on a large-scale statistical Time-Use Survey and national ownership rate statistics. However, the raw data are extracted from the UK 2000 TUS and UK statistics collected between 2000 and 2009, which raises several issues. First of all, we mention that since these statistics are country-specific, there should be some differences between the UK and France. Nevertheless, we posit that these differences are minor compared to the fact that the activities, appliances and uses have evolved substantially since 2000. For instance, Figure 7.9 depicts the trends in terms of equipment rates of mobile and fixed-lines phones, computers, and home internet access for the French population between 1998 and 2016. Clearly, the activities related to some equipments such as a smartphone or a tablet were not included in the UK 2000 TUS, whereas there has been a meaningful surge in the ownership of home internet access. Recent trends involve thus a diversification of the uses and media. Besides, the development of the Internet of Things (IoT) comes with a new energy demand, by creating new uses and connecting devices that didn't use to be. Yet, this enhanced connectivity represents also an opportunity for energy efficiency actions, whose extent is difficult to evaluate. Overall, the energy implications of this new trend are therefore largely unpredictable.

Finally, we emphasize the fact that the metabolic heat gains are generated in a simplified manner, with a constant rate of 50 W per active occupant of the dwelling. This is another limitation of the

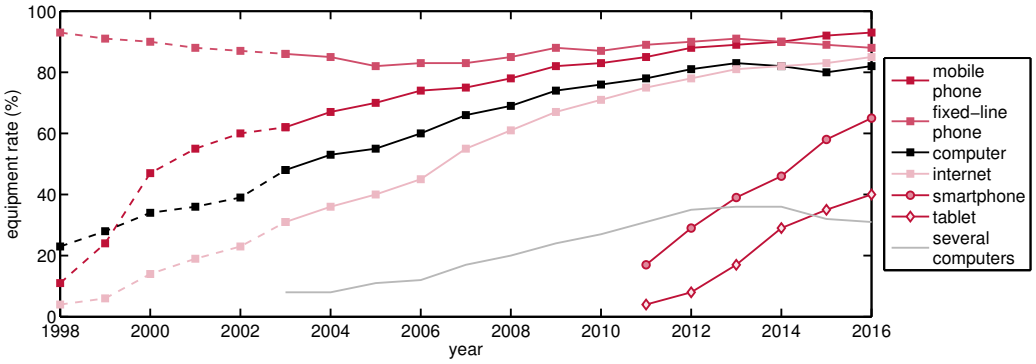

Figure 7.9: Equipment rates in France: telephony, computer, home internet access. Source: CREDOC, survey "Conditions de vies et aspirations". The survey concerns people aged 18 and more (before 2003, dashed lines) and 12 and more (since 2003, solid lines).

accuracy of the model.

7.4 Numerical experiments

7.4.1 Implementation

The estimation procedure described in Algorithm 7.2 is tested on several datasets, where the data are generated according to the random process in Algorithm 7.3. More precisely, Section 7.4.2 is devoted to the analysis of a single experiment, while multiple (Monte Carlo) experiments are performed in Section 7.4.3. In both cases, the data are simulated for several days, at time-step 10 minutes, and 10 days are used in the training set. Regarding the implementation, we have used Matlab, with the gradient-based trust-region routine `fminunc` for solving Step 1. of the algorithm and the gradient-free simplex search `fminsearch` for solving Step 6. The code for the aggregation strategy in Step 3. is much inspired by P. Gaillard and Y. Goude's `R` package `opera` [GG16]. The aggregation weights are initialized at $t = 1$ with a uniform prior: $w_{k,1} = 1/K$, $k = 1, \dots, K$.

The results obtained by Algorithm 7.2 are compared to a naive estimator, obtained by applying the standard EM algorithm with total load curve Q seen as a deterministic input of the open-loop R3C2 model, instead of Q_u (i.e. by improperly including the non-heating loads in the heating power). This naive model is estimated by Algorithm 3.1. A second comparison is provided by estimating an ARX model of order 2 with respect to every input and output of the open-loop building system. The ARX model is learned with ordinary least squares.

In the single experiment analyzed in the subsequent Section 7.4.2, the dwelling has two occupants and the weather conditions are extracted from BESTLab measurements in January 2016. The purpose of this section is to study in details the different steps of the suggested algorithm, on a given random example. The chosen aggregation strategy is Strategy 1 in Algorithm 7.1, that is the weights computed classically by aggregation of the predictions of the indoor temperature are then applied to the PI parameters of the experts $\eta^{(k)}$. After a few trial-and-error tests, the learning rate ρ of the EWA strategy is set to $\rho = 1/2$. An alternative would be to design a grid-based calibration of ρ, a feature included for instance in the package `opera` [GG16].

Table 7.1: Estimated parameters by Alg. 7.2 on a single experiment, compared to a naive estimation and an ARX model of order 2. The parameters z_K and $z_\tau = 1/\tau$ of the PI controller are not learned by the naive method and ARX model.

Parameter	**True value**	**Naive**	**ARX(2)**	**Alg. 7.2**
UA	0.249	0.262	0.240	0.243
g	0.429	0.200	0.135	0.390
τ_1	16.7	60.9	6.4	28.3
τ_2	2.9	4.8	0.17	3.4
z_K	2.76	-	-	2.27
z_τ	0.020	-	-	0.025

7.4.2 Analysis of a single experiment

The estimated physical parameters of the R3C2 network are reported in Table 7.1, whereas the corresponding trajectories, on the learning set, of the useful load $\widehat{Q}_u$ and indoor temperature $\widehat{T}_i$ are depicted in Figure 7.10. In order to avoid overfitting, the trajectories simulated with the estimated parameters at each step of Algorithm 7.2 are obtained without any model nor observation noise; this explains their smoother shape. The results show that the algorithm presented in this chapter yields a substantial improvement in the estimation of the parameters, compared to the naive standard approach. In similarity with the results obtained in Chapter 6, the ARX identification also fails into catching the dynamic time constants of the model, even if its estimation of UA is accurate. Likewise, the solar transmittance g given by the identified ARX model is largely underestimated. In the end, the results emphasize the fact that the ARX model is not adapted to a small amount of closed-loop generated data. Moreover, the simulated trajectories $\widehat{Q}_u$ and $\widehat{T}_i$ are also consistent with the actual observations. Although it cannot - and should not - account for the sudden peaks caused by some uncontrollable free gains, the simulated heating power matches the observed useful load rather well, with a RMSE $\mathcal{E}_{qu} = 0.32$ kW (normalized RMSE: 16.3%). The indoor temperature agrees even more closely with the actual measurements, with a RMSE of $\mathcal{E}_{ti} = 0.31°$C (normalized RMSE: 1.7%).

Furthermore, the progression of the root mean square errors $\mathcal{E}_{ti}$ and $\mathcal{E}_{qr}$ is represented in Figure 7.10(c), starting from the estimated parameters at Step 1 till the convergence of the algorithm. Similarly, Figure 7.10(d) shows the improvement of the estimation of the parameters of the R3C2 network over the iterations of Algorithm 7.2. Both plots illustrate the benefit of introducing the correction of $\widehat{\theta}_1$ obtained at Step 1, although this improvement is not monotonic. Besides, it can be noted that the RMSE computed on the trajectory of the indoor temperature is hardly improved, being already quite small. Hence, the second term $\mathcal{E}_{qu}$ in the cost $J_2(\theta)$ in (7.13) is key to enable the improvement of the estimation.

The estimated useful load amounts to an energy consumption of 527 kWh over the time intervals covered by $\mathcal{T}$ (1140 instants out of 1440). On the other hand, the actual consumption of the heating system during this same period is 478 kWh so that we have $\widehat{Q}_u > Q_h$ in average, as expected. More interestingly, the difference between these two consumptions is 49 kWh, whereas the actual free heating gains are 51 kWh. This confirms the good performances of our algorithm also at a static time-scale, since it is able to accurately estimate the total amount of energy used for space heating.

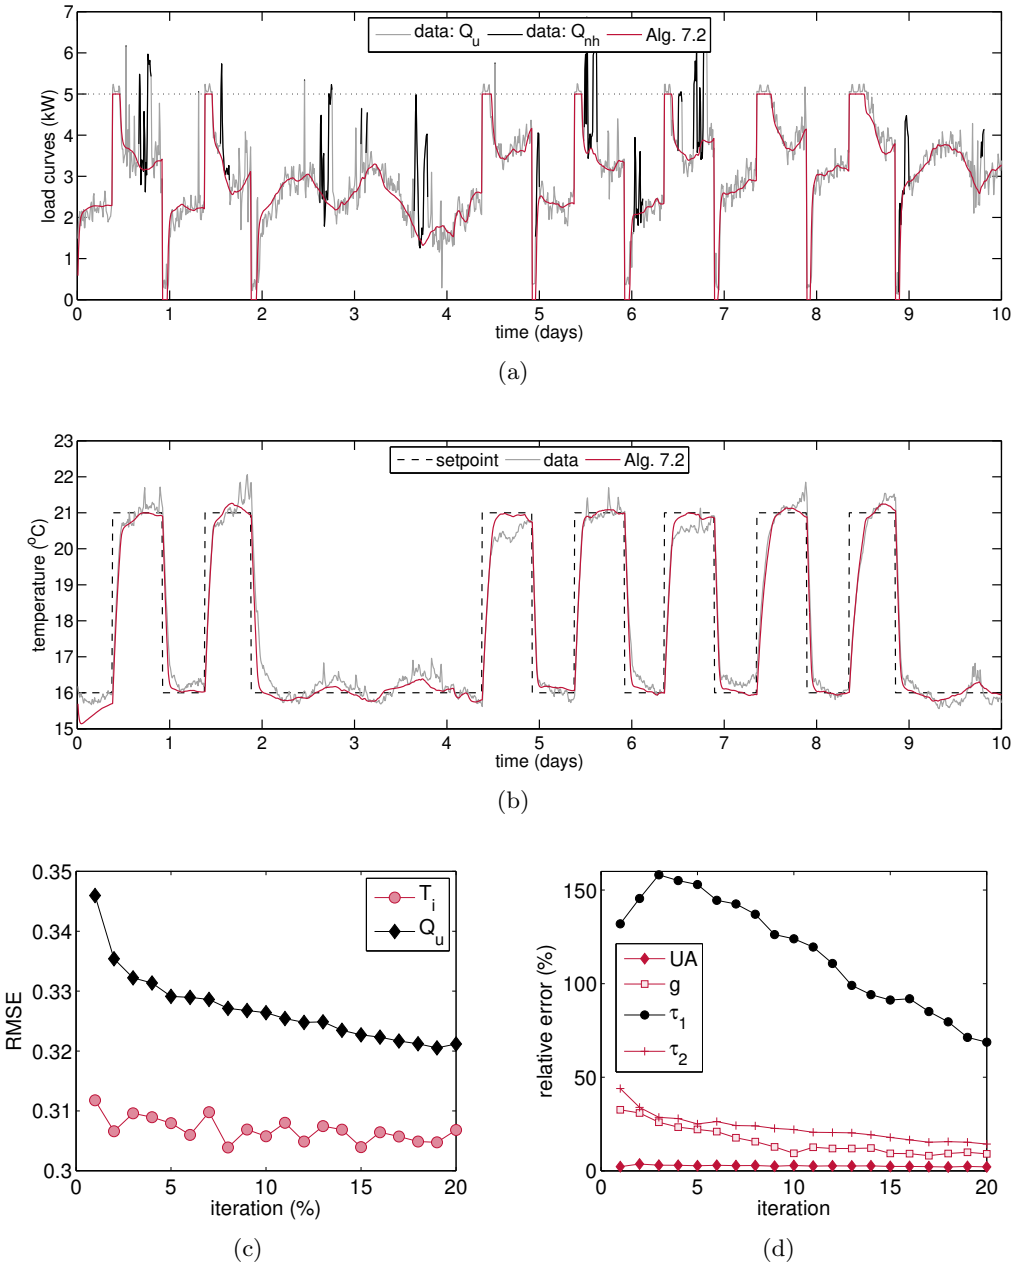

Figure 7.10: Estimation from an uncertain load curve of (a) the heating load and (b) the indoor temperature. The actual measurements are in light grey, the simulated trajectories in red. The non-heating loads are in black, the dotted line represents the limiting power Q_h^{max}. Evolution over the iterations of Algorithm 7.2. (c) of the RMSE of the estimated indoor temperature (circles) and of the useful load (diamonds), and (d) of the estimation errors of the physical parameters of the R3C2 network.

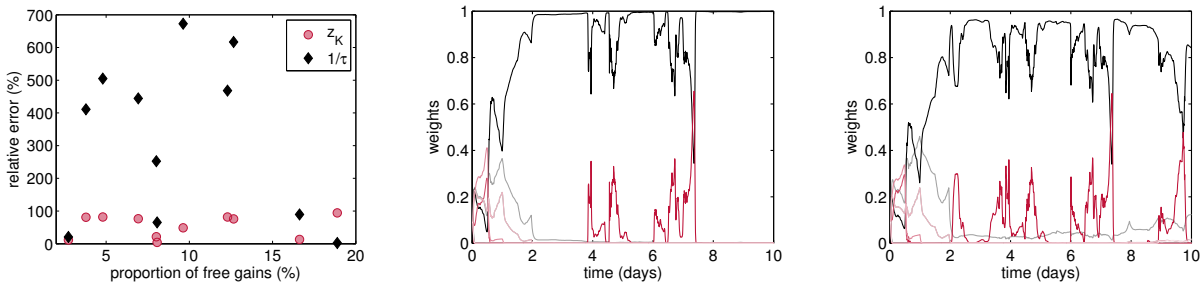

(a) (b) (c)

Figure 7.11: (a) Accuracy of the 11 expert models and associated aggregation weights at the (b) first (c) last iteration of Algorithm 7.2. The least important weights are in light grey and not visible on these last two plots.

Regarding the estimation of the PI parameters, note that the naive and ARX models do not model the regulation, so that no comparison can be made with Algorithm 7.2 for z_K and z_τ. Hence, in the following, we analyze more specifically the aggregation of the expert models suggested at Step 3 of Algorithm 7.2. In this experiment, we obtain a set of $K = 11$ experts $\widehat{\eta}^{(k)}$, computed at Step 2, whereas the number of intervals for which the load curve contributes entirely to heating is $n_0 = 18$. The accuracy of these experts varies greatly, although there is no apparent relationship between the proportion of free heating gains in I_{0k} and the accuracy of $\widehat{\eta}^{(k)}$, as shown in Figure 7.11(a). The best expert however has a good accuracy for both PI parameters z_K (11.5%) and $z_\tau := 1/\tau$ (20.6%) and is estimated from the time interval with lowest amount of free heating gains (bottom left points in Figure 7.11(a)). Other experts may be more accurate on either one of the two parameters, but much less on the other. For instance, one expert has an accuracy of 4.1% on z_K but only 65.3% on z_τ. Another has an accuracy of 2.2% on z_τ but 94.5% on z_K. This emphasizes the need for an efficient aggregation strategy able to target the best experts.

The weight of this best expert in the aggregation strategy is depicted in black in Figures 7.11(b) and (c), where Figure 7.11(b) (respectively 7.11(c)) depicts the time-evolution of the sequential weights estimated from the model $\widehat{\theta}$ obtained at the first (respectively the last) iteration of Algorithm 7.2. In both plots, the effect of the uniform initialization of the weights requires several iterations before being forgotten, between one and two days. The best expert is distinctly given the most important weight after this "burn-in" phase. The final time-independent weights are the average of the weights in Figures 7.11(c) over the last two days. They yield an aggregated model $\widehat{\eta}$, given in Table 7.1, with accuracy 17.8% on z_K end 25.5% on z_τ, which is close to the best expert. These results validate thus the first aggregation strategy described in Algorithm 7.1.

Although each step of Algorithm 7.2 may be improved, the results presented for this single experiment are already promising. The identification of the PI parameters on distinct intervals and their aggregation into a single estimator seems relevant, whereas in the end, the information brought by the connected objects contribute to better estimations of the R3C2 network.

7.4.3 Multiple experiments

In order to further validate our approach, we proceed now to the analysis of 50 experiments. Each experiment has a randomized set of nominal parameters of the R3C2 network, as well as a randomized starting time between November 2015 and January 2016 for the weather conditions extracted from the BESTLab database. Similarly to the previous results, the estimation $\widehat{\theta}$ provided

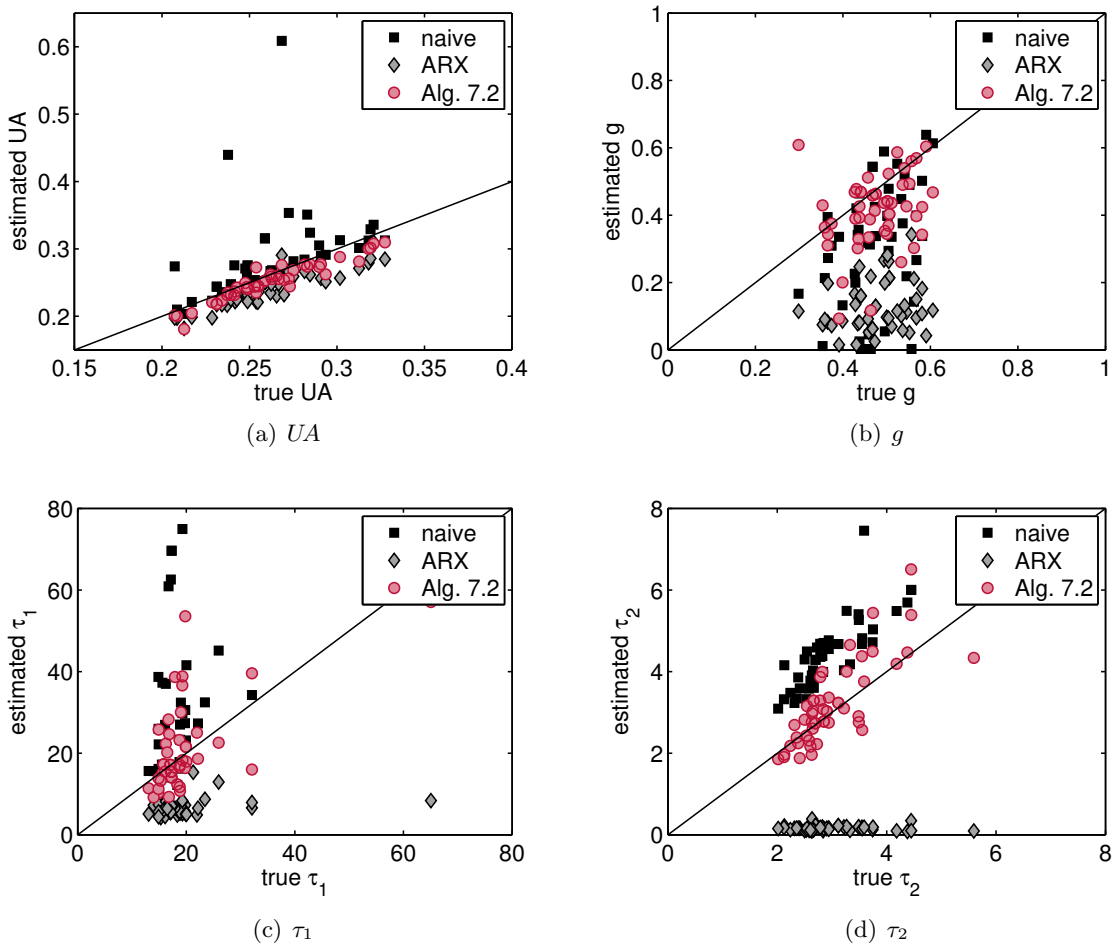

(a) UA (b) g

(c) τ_1 (d) τ_2

Figure 7.12: Estimation of the physical parameters of the R3C2 network from an uncertain load curve, 50 experiments. Black squares: naive EM algorithm, grey diamonds: ARX model of order 2, red circles: Algorithm 7.2.

by Algorithm 7.2 is compared to a naive EM estimation and an ARX model of order 2, where in particular both the naive and ARX methods use the aggregated load curve Q as input.

The 50 estimations of the parameters UA, g, τ_1 and τ_2 are plotted in Figure 7.12 agains their respective nominal values, and the corresponding boxplots of the estimation errors are in Figure 7.13. These plots confirm that the algorithm suggested in this chapter outperforms a naive EM estimator or an ARX model of order 2. In particular, the gain is significant for the heat transfer coefficient UA, with an accurate estimation with reduced bias and variance of the estimation error. Algorithm 7.2 brings also a consequential improvement of the fast time constant τ_2: the ARX model constantly largely underestimates it, whereas the naive EM overestimates it. The fact that the naive EM overestimates it was already noticed in Chapter 3: the improved is here due to the correcting step 7 in Algorithm 7.2. The estimations of the slow time constant τ_1 and of the solar transmittance g seem more difficult for all three algorithms, although Algorithm 7.2 behaves more robustly and yields a reasonable order of magnitude of these two quantities. Hence, these first simulations illustrate the benefits of the method suggested in this chapter.

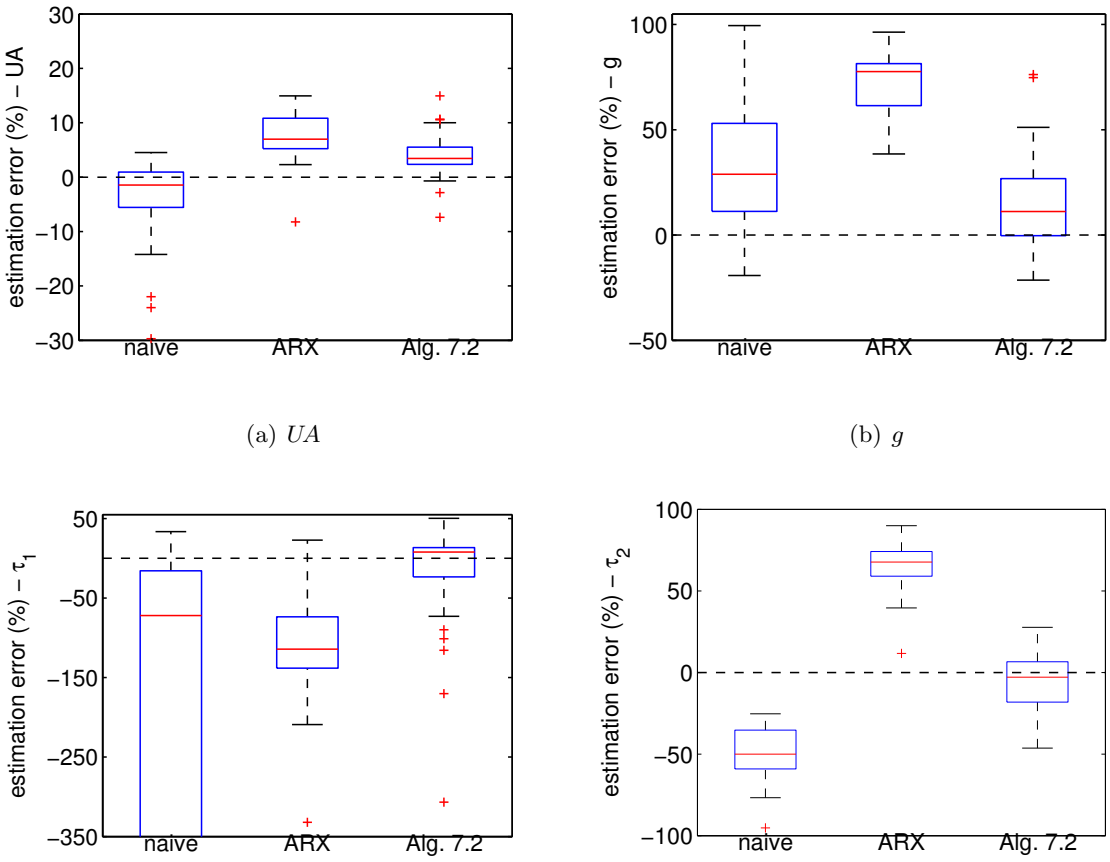

(a) UA (b) g

(c) τ_1 (d) τ_2

Figure 7.13: Estimation errors of the physical parameters of the R3C2 network from an uncertain load curve, 50 experiments.

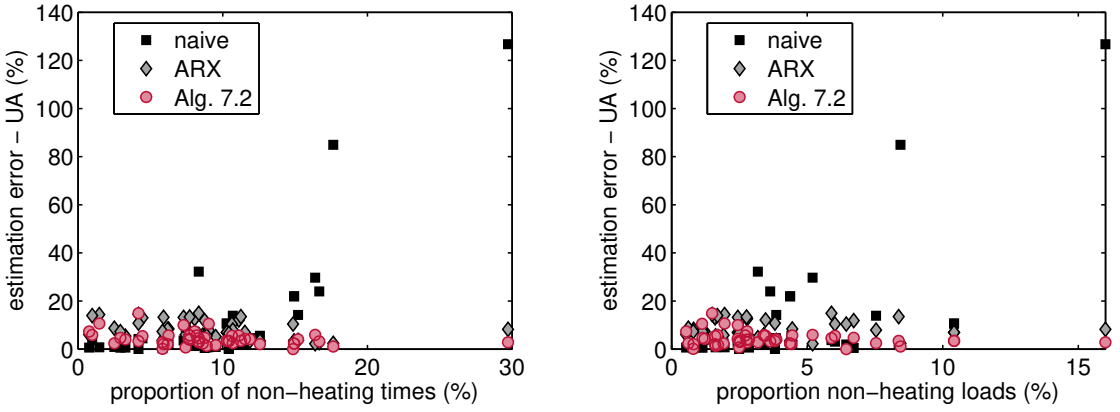

Figure 7.14: Estimation error of UA against (left) the frequency of non-heating events, (right) the proportion of non-heating loads in the aggregated load curve. Black squares: naive EM algorithm, grey diamonds: ARX model of order 2, red circles: Algorithm 7.2. 50 experiments.

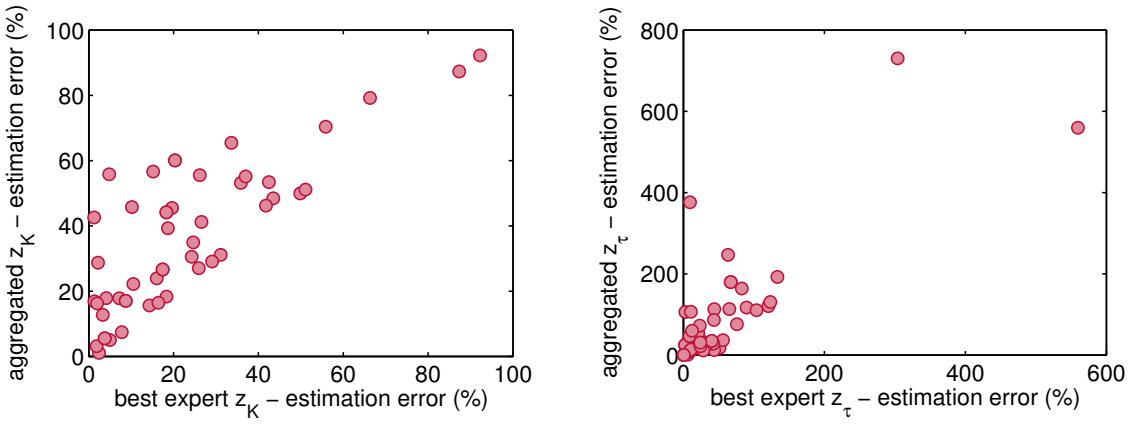

Figure 7.15: Estimation error of the aggregated $\widehat{z}_K$ (left) and $\widehat{z}_\tau$ (right), against the best expert $\widehat{z}_K^{(k)}$ (left) and $\widehat{z}_\tau^{(k)}$ (right), respectively; 50 experiments.

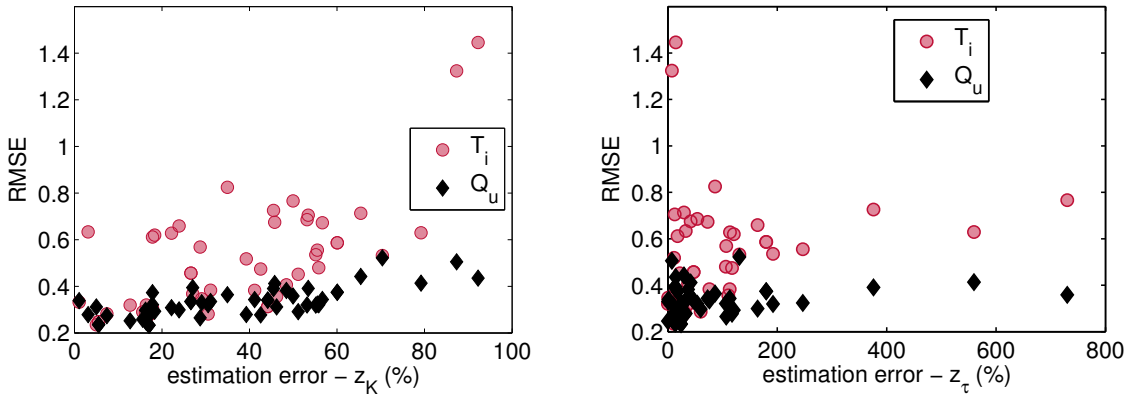

Figure 7.16: RMSE of the estimated trajectories of T_i (red circles) and Q_u (black diamonds), against the estimation error of the gain z_K (left) and the time constant z_τ (right) of the PI controller; 50 experiments.

Besides, Figure 7.14 displays the absolute estimation error of UA in terms of frequency of non-heating events (left plot) and total proportion of non-heating loads in the aggregated load curve (right plot). From these plots, it seems that the estimation of UA by Algorithm 7.2 and by the ARX model is robust to the experimental conditions, since it does not depend on the actual frequency or volume of non-heating loads. On the contrary, the performances of the naive EM algorithm deteriorate strongly whenever non-heating events occur more than 10% of the time, or whenever the proportion of non-heating loads exceeds 4% of the total energy consumption. When there is almost no non-heating load, then the naive EM algorithm performs better than Algorithm 7.2, as can be expected since in such case the approximation $Q \simeq Q_u$ is valid. Similar trends are observed for the other parameters g, τ_1 and τ_2, although the corresponding plots are not displayed here.

Finally, we focus more specifically on the behaviour of Algorithm 7.2 presented in this chapter. Regarding the estimation of the regulation $\widehat{\eta}$, Figure 7.15 depicts the estimation error of the gain $\widehat{z}_K$ (resp. the inverse of the time constant $\widehat{z}_\tau$) against the estimation error of the best expert $\widehat{z}_K^{(k)}$

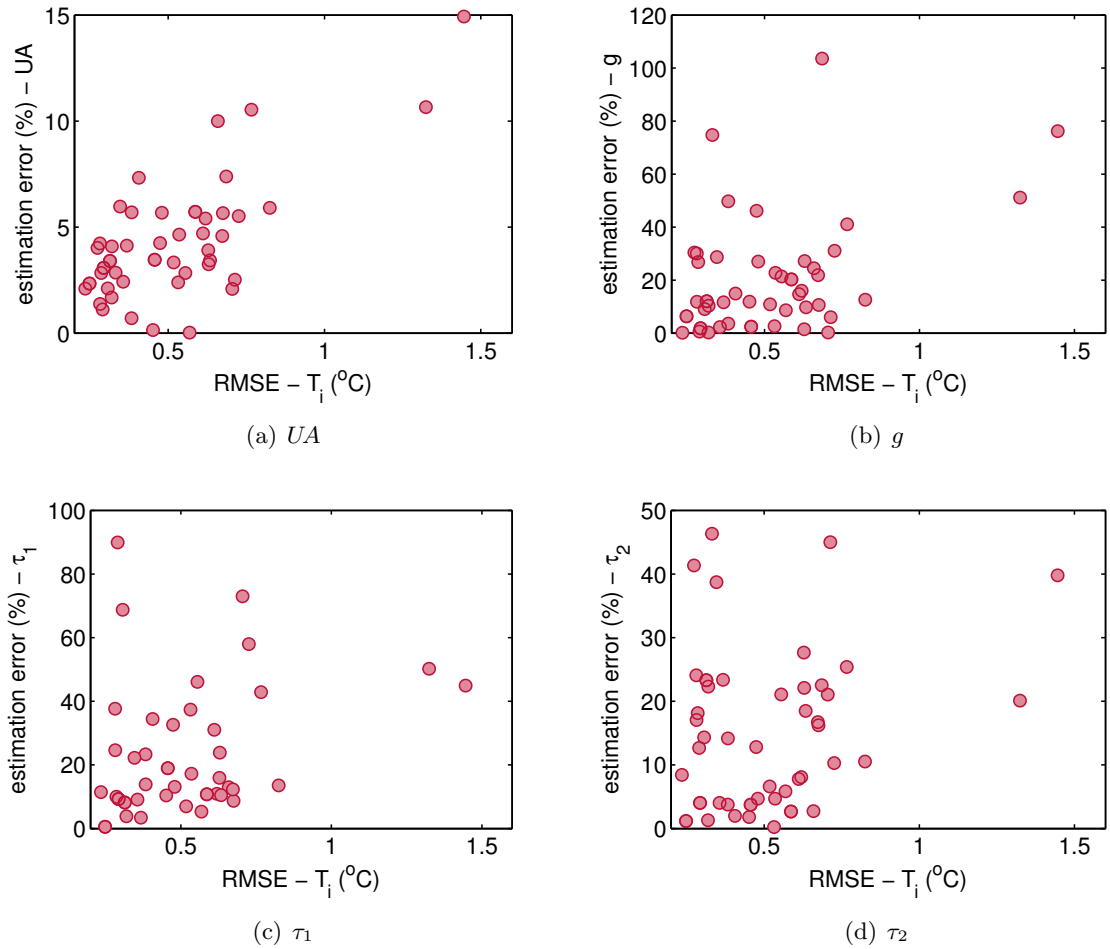

(a) UA (b) g

(c) τ_1 (d) τ_2

Figure 7.17: Estimation errors of (a) UA, (b) g, (c) τ_1 and (d) τ_2 as a fonction of the final RMSE $\mathcal{E}_{qu}$ of the estimated trajectory of the indoor temperature T_i; 50 experiments.

(resp. $\widehat{z}_\tau^{(k)}$). These plots emphasize the importance of the estimation of the experts $\widehat{\eta}^{(k)}$, since the aggregation strategy does not outperform the best expert. Note that as explained in the analysis of the single experiment, the best $\widehat{\eta}^{(k)}$ may not combine both the best $\widehat{z}_K^{(k)}$ and the best $\widehat{z}_\tau^{(k)}$, which is why some aggregated $\widehat{z}_K$ or $\widehat{z}_\tau$ lie above the main diagonal in Figure 7.15.

Furthermore, poorly estimated PI parameters lead naturally to less accurate estimated trajectories of the indoor temperature and the useful load, as shown in Figure 7.16. Since Figure 7.17 confirms on the other hand that smaller values of the RMSE $\mathcal{E}_{ti}$ do correspond to an increased accuracy of the estimation of UA, g, τ_1 and τ_2, this highlights thus the impact of the estimation of the PI experts at Step 2. of Algorithm 7.2. Similar results are obtained with $\mathcal{E}_{qu}$ (plots not shown). Another consequence of Figure 7.17 is that the quality of the estimation of the R3C2 model may be assessed through the final values of $\mathcal{E}_{ti}$ and $\mathcal{E}_{qu}$.

Table 7.2: Identified physical parameters for the western BESTLab cell, with unknown free heating and non-heating gains in the aggregated load curve.

Parameter	**U**	τ_2	τ_2	**g**
Estimation	0.73	37.8	8.7	0.45

7.4.4 BESTLab data

Similarly to the numerical study in Chapter 6, the suggested Algorithm 7.2 is tested against BESTLab data. Indeed, both cells used in the experiment have two radiators, a controlled one for heating and a second one that represents free heating gains. The first radiator is therefore regulated such that the measured indoor temperature of the cell follows a predetermined setpoint. On the other hand, the second radiator is not regulated: it is activated according to a predetermined schedule. When activated, this second radiator heats at full power, provided that the indoor temperature does not exceed 27°C. Note that in this situation, the weight of the free heating gains represented by the latter radiator is too important compared to the controlled heating system, since both radiators have the same nominal power (500 W). Finally, there is no non-heating load in the cells. We choose to simulate such loads for a given cell by randomly picking the free heating gains from the other cell. In the sequel, we analyze the results of the estimation for a 10-day experiment in the western cell.

Figure 7.18 shows the estimated trajectories of the heating power and indoor temperature. This case study emphasizes the limits of Algorithm 7.2, with total load curve containing a small proportion of controlled heating power. In this example, the energy consumed by the heating system represents only 30% of the total consumption whereas the free heating gains account for 49%, and the other loads for 21%. Hence, most of the load curve is explained by uncontrollable loads, as can be seen from Figure 7.18. Note that this experimental scenario is much unlikely for an actual building. The performances of the algorithm may therefore be seen as being relative to a limit case. In these conditions, the notion of setpoint temperature becomes paradoxical and Algorithm 7.2 faces a contradiction: on one hand, it is asked to minimize the RMSE on the trajectory of the indoor temperature, which clearly does not follow the setpoint, and on the other hand, Algorithm 7.2 should generate a saturated command equivalent to the useful load curve Q_u and adapted to the setpoint signal. Figure 7.18(b) shows that the estimated indoor temperature follows closely the setpoint, rather than the actual measurement. This may be exploited in the following manner: the difference between the observed and estimated trajectories may be attributable to the free heating gains only, such that an estimation procedure of the Q_{fg} could be considered. This is a promising track for future work, that has not been explored in the thesis.

Table 7.2 gives the point estimates of the four physical parameters of the identified R3C2 model. The U value is slightly underestimated compared to the findings of Chapter 3. Yet, this value still corresponds to a low-energy building, which shows that the method is able to correctly characterize the building. The slow time constant seems also to be underestimated. However, the estimated fast time constant τ_2 and solar transmittance coefficient g fall within the bounds of the error bars in Figure 3.8 (10 days of data), which reassures of the good behaviour of Algorithm 7.2 in a highly degraded context of observation.

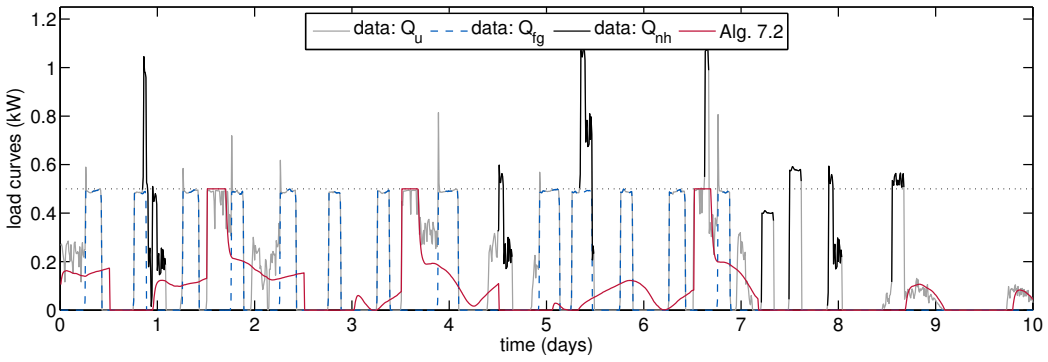

(a)

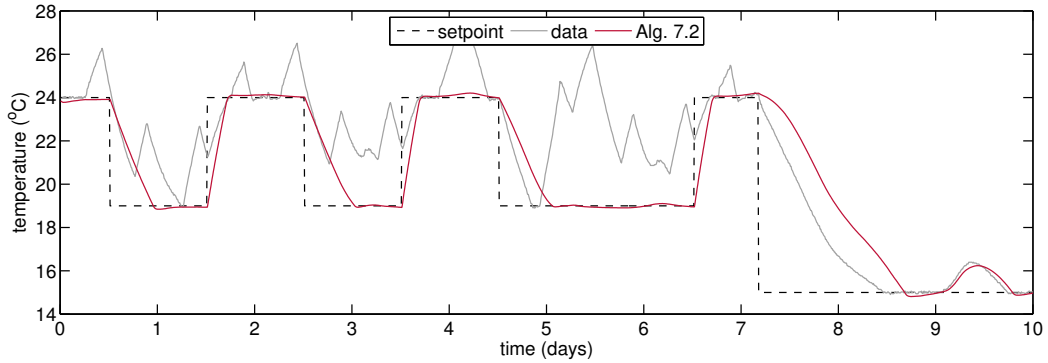

(b)

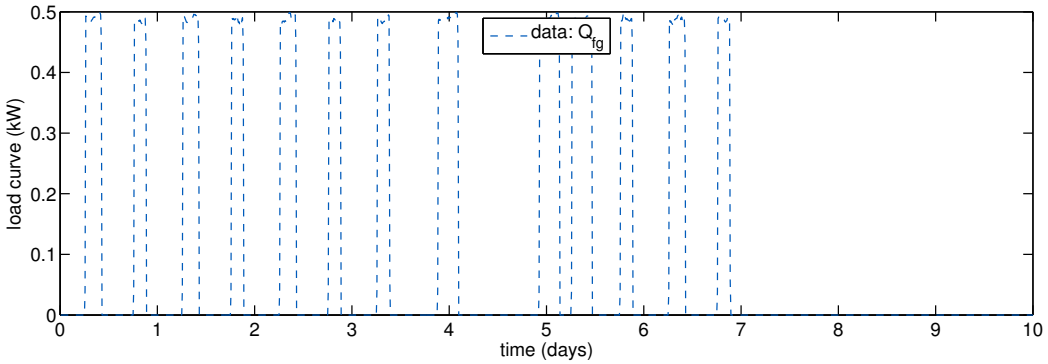

(c)

Figure 7.18: Estimated trajectories of (a) the heating power and (b) the indoor temperature in BESTLab western cell, with unknown free heating and non-heating gains in the aggregated load curve. (c) The unknown free heating gains.

7.5 Summary and discussion

This chapter, mostly prospective, is a first attempt to account for the uncertainty of the heating load curve in the identification of building models. The uncertainty is caused by the fact that an energy utility may only measure the total load curve, instead of the heating loads only. Given the main assumption that the connected objects of the building enable one to observe a timeline of all the non-heating events, we suggest an iterative procedure, summarized in Figure 7.3 and Algorithm 7.2, to estimate both the open-loop building model, here the R3C2 network, and the indoor temperature controller. Even if the results described in the previous sections are promising, each step of the algorithm may be improved. We have thus highlighted the impact of the accuracy of the PI experts $\widehat{\eta}^{(k)}$ on the accuracy of the aggregated $\widehat{\eta}$, and thereby also on the parameters of the R3C2 network. Here, we have used an ordinary least squares identification of $\widehat{\eta}^{(k)}$, improperly since we include the uncontrolled free heat gains in the output of the controller. One might therefore want to explore alternative strategies, accounting for this stochasticity. For instance, censored models could be worth investigating. Similarly, the updating step 7 in Figure 7.3 may be performed differently, for instance with a more suitable cost or a finer optimizer. Our experience suggest that the key ingredient is the inclusion of a term penalizing the difference between the observed useful load Q_u and the simulated one $\widehat{Q}_u$ in the cost $J_2(\theta)$ (7.13), rather than only the difference between the obseved and simulated indoor temperatures. However, the estimation with BESTLab data illustrated the fact that this makes sense mostly when the free heating gains are not overrepresented in the load curve. The termination test could also involve the stabilization of the estimated parameters rather than the stabilization of the costs.

Moreover, the suggested algorithm depends on the modelling framework defined by the hypotheses given in Section 7.1.2. The advantage of this framework is that it does not require any knowledge about the specific consumptions of each and every appliance in the building. We have only assumed that the consumption of the hot water tanks is known, yet this might be relaxed: if it is hidden, then it is sufficient to know when the water drains occur and stop, similarly to other non-heating loads. Different assumptions would lead naturally to other algorithms. For instance, still in the context of the connected objects, one could assume that the consumption Q_h of the regulated space heating space could be available at a dynamic time step, whether it be by a smart heater or a dedicated connected plug. Then, the information is more accurate compared to our framework, since at all times t, the inequalities $Q_h(t) \leq Q_u(t) \leq Q(t)$ hold, with known and accurate lower and upper bounds.

Future work should also investigate the issue of metabolic heat gains. Indeed, those are neglected in this work, with numerical experiments corresponding to dwellings with 2 occupants only (simulated data) or no occupant at all (BESTLab data). However, the occupation of office buildings is much more consequent, representing thus a substantial heat gain. It seems unlikely to count precisely the number of occupants at a dynamic time step. One idea would be to include a probability distribution on the number of occupants, with moments given by an algorithm infering the occupancy given a set of connected environmental sensors, see [AAP+16] for an example.

Part III

Conclusion and perspectives

Chapter 8

Conclusions and perspectives

In this thesis, we have developed a modelling framework, as well as a set of estimation algorithms, in order to address the issue of data availability for building identification. We believe that this is a question of interest, because building identification is often prevented in practice due to the lack of measurements of the physical quantities involved in the model. The main question that underlies the thesis is the following: how can an energy utility create building identification based energy services, such as thermal diagnostic or optimal control, considering that the only data known with certainty is the total load curve of the building? The emerging solution sees the utilities providing their clients with connected smart thermostats and weather stations. It becomes thus plausible to have access to dynamic measurements of the indoor and outdoor air temperatures. Yet, these sensors come with some uncertainty, while other quantities such as the useful heating load curve or the solar gains remain hidden.

The framework elaborated in the thesis, and based first of all on a state-space representation of a building by an equivalent R3C2 electrical network and the Expectation-Maximization algorithm, has proven to be flexible enough to account for this uncertainty. We have shown in Chapter 3 that although being actually exploited in closed loop, the open-loop building model may be accurately identified without explicitly modelling the regulation. However, if the indoor temperature is not measured at all, a closed loop model becomes necessary. Consequently, Chapter 4 presents an estimation algorithm that uses the Hamiltonian Monte Carlo sampler to solve the underlying nonlinear identification problem. Although the indoor temperature is getting more accessible with the spread of connected thermostats, this result may be of interest for an utility, since it can be seen as non-intrusive.

On the other hand, if the indoor temperature is measured, open-loop approaches are thus appropriate. They are employed in Chapters 5 to 7, where different versions of the EM algorithm are designed successively for the different inputs of the R3C2 model. Firstly, the uncertainty related to the connected outdoor temperature sensor is caused by the fact that the energy utility does not know the exact placement of the sensor. This leads to a possible bias of the measurements, which is addressed in Chapter 5 by appending a dynamic model of the outdoor temperature to the R3C2 model. The resulting switching linear state space is identified by an EM algorithm almost identical to the reference one presented in Chapter 3. Next, Chapter 6 focuses on the solar radiation. Since this quantity is often unknown in practice, due to the prohibitive cost of the pyranometer, we suggest to replace this sensor by a cheap connected temperature sensor intentionally left exposed to the sun. The Rao-Blackwellised Particle Smoother EM algorithm exploits then the difference of measurements between this temperature sensor and the reference sheltered one to accurately estimate both the building model and the received solar radiation. Ultimately, we assume in Chapter 7 that a timeline of events specifying when the non-heating loads are either switched On or Off

is available. This enables the joint identification of the open-loop building model and the indoor temperature regulation from the aggregated total load curve that improperly contains non-heating loads.

This work shows thus promising results on simulated and real datasets. This opens many perspectives. Amongst others, the following ones may be considered.

1. In order to get closer to the practical constraints of an energy utility, the next modelling step consists in assuming that several inputs or outputs of the building model are simultaneously uncertain. For instance, we may consider a dwelling equipped with a smart thermostat, two connected outdoor temperature sensors and a smart electricity/gas meter. Then, the deterministic variables would be the indoor and outdoor air temperatures, whereas the solar gains and the useful heating load curves are hidden. Since the energy utilities already equip their clients with smart thermostats and connected weather stations, the only additional sensor would be the outdoor temperature sensor. Besides, we believe that solving this identification problem should come at little effort, by suitably mixing the strategies introduced in Chapters 6 and 7. In particular, the EM algorithm used in the first step of Algorithm 7.2 may be replaced by the Rao-Blackwellised Particle Smoother EM of Algorithm 6.3.

2. Another line of work, consistent with the context of smart buildings, consists in taking into account time-varying nature of the building structure. Indeed, opening doors and windows alter the thermal properties of the building. By using connected contacts, or detection algorithms based for example on the gradient of the indoor temperature, we may extend the R3C2 network such that it includes electrical switches. Such a model has been suggested recently in [SBPW16] and is worth exploring further.

3. Similarly, we may question the hypothesis that the model comprises a unique thermal zone. Cheap connected temperature sensors make it possible to consider multizone models. The R3C2 network adapts easily to the multizone context, by adding other nodes, resistances and capacitors. This results in an augmentated linear state-space model, whose order can be reduced using aggregation techniques, as for instance in [DBMM10]. The methods designed in this thesis should apply likewise to this higher-order model.

4. Should the future development of the connected objects meet with the elevated expectations they provoke, more detailed models of some subcomponents of the building are also worth considering. For instance, [CSM+13] identifies an equivalent electrical network representing a refrigerator. Such models could either be appended to the R3C2 network, or studied separately.

Appendix A

Estimation of the solar radiation: kriging model

A.1 Introduction

Estimation and prediction of the solar radiation has gained growing attention from researchers in the past few years, due to the great number of applications, for instance in the design of solar plants. Most articles in the litterature address the problem of prediction, that is given past measurements at one site, predict the forthcoming irradiations at the same site. Here, our problem is slightly different: we want to reconstruct the radiations formerly received at any given location in a given area, typically in France. The difficulty lies first in the availability of the data. Indeed, solar radiations are measured by pyranometers, which are expensive sensors. Hence, it is not conceivable that past measurements would be available everywhere, and we shall typically assume that no measurement at all is available. We face thus the following problem:

how to construct a statistical model for the solar radiation at a given location, without the corresponding learning dataset?

To address this issue, we shift our focus from estimation of the solar radiation to estimation of the cloud coverage. This is due to the fact that there exist simple physical models, depending on time and geographical coordinates, that estimate the solar radiation. These models are accurate enough provided that they are corrected by a measure of the cloud coverage. Similarly to the approach described in Chapter 6, we consider Perrin de Brichambaut's model in [PdBV82]:

$$G_{h,clr} = \kappa G_0 (\cos\theta_z)^{1.15} \tag{A.1}$$

where $G_{h,clr}$ is the global horizontal solar flux under clear sky conditions, κ is a constant set in the litterature to $\kappa = 0.81$ ([RW99]), θ_z is the azimuth angle (that depends only on the latitude, longitude and solar time) and G_0 the extraterrestrial solar constant (theoretical radiation at the top of the atmosphere, depends on the day number). This model is corrected by a measure $N \in [0,1]$ of the cloud coverage according to

$$G_{h,cld} = G_{h,clr}(1 - aN^b). \tag{A.2}$$

If the sky is completely clear ($N = 0$), 100% of the initial model $G_{h,clr}$ is kept, whereas a fraction $(1 - a)$ remains even if the sky is totally overcast ($N = 1$). For example, the values $a = 0.75$ and $b = 3.4$ are obtained from a fit to a 10-year hourly record at Hambourg in [KC80]. The benefit in considering cloud coverage rather that solar radiation is that there exist open access records of the former in many locations in the world (see Section A hereafter and Figure A.1). It suggests therefore

the idea of using historical cloud cover data at a given number of locations in order to estimate their dependency in space and time through a correlation model. This allows one to subsequently estimate the cloud coverage at any new location, at any time. This technique is called kriging, and has been applied first in the context of mining, see [Cre91] for a reference textbook.

The underlying idea that some data are "closer" in a certain sense to a subset of the dataset has been mainly considered in the solar prediction litterature in a temporal context. Indeed, many contributions (e.g. [BC13, AGG16, JPML16]) operate a preliminary classification step with a clustering algorithm, and separate models may be fitted within each cluster. Yet, using spatial patterns and/or applying kriging techniques to the field of solar estimation and prediction is not so common. We review briefly some of the contributions of the literature in this line of work. [GFDD16] is a first example, in which the authors aim at improving the performances of Numerical Weather Prediction (NWP) models in Spain, by downscaling them in space and time, using machine learning techniques. The prediction at a given location is based on the SVR algorithm with the forecasts of the total cloud cover and of the solar radiation at 4 nearby corner points in the NWP grid as features. More closely related to this work, a sophisticated kriging method is implemented in [YGD+13]. the authors analyze the spatio-temporal covariance structure of the solar radiation measured at 10 locations in the island of Singapore. They show how to transform the data in order to make reasonable the assumptions of time stationarity and space isotropy. Their results are encouraging, all the more given that Singapore has an unstable climate, making it tricky to address the prediction of the solar flux. Compared to our framework, the limit is that their study does require the mesure of the solar radiation at known locations. Moreover, they benefit from a dense network in space, which is not our case here, as shown by Figure A.1 afterwards. A second article in this framework is [ISW12], which uses also a dense network of measurement of the solar radiation around Osaka, Japan: their spatial resolution is < 5 km, the time resolution is 5 minutes. They model the space-time covariance matrix originally, with an exponential function that decays with a Mahalanobis distance, and spatial anisotropy is included in the model to take the cloud movement into account.

In this work, we suggest a way to reconstruct the solar radiation received at a given location using open access data only, rather than expensive on-site measurements. We use the framework of isotropic ordinary kriging. Our aim is not to provide the "absolute best" model, but rather to show that the methodology is suitable for this constrained context. We will therefore present some simple, standard isotropic spatial models, and future works can include a refinement of those models.

The rest of this appendix is organized as follows. After describing the dataset in Section A.2, we study two classes of models. Both are purely spatial models, relying on the assumption of time independence, and are analyzed in Section A.3. Finally, we provide a comparison of those models in Section A.5, and we discuss our findings in Section A.6.

A.2 Data

We use some open source data from METAR (METeorological Aerodrome Report) weather reports of past observations in 100 stations in France; see positions in Figure A.1. One can find such archived reports emitted from various locations in the world for instance from the website mesonet.agron.iastate.edu.

The chosen dataset consists of half-hourly measurements of several meteorological variables, including cloud cover, air temperature, dew point, relative humidity, atmospheric pressure, wind speed and wind direction. They were collected during two months, in November and December 2015. There are therefore up to 48 observations per day per station and per variable. Yet, the

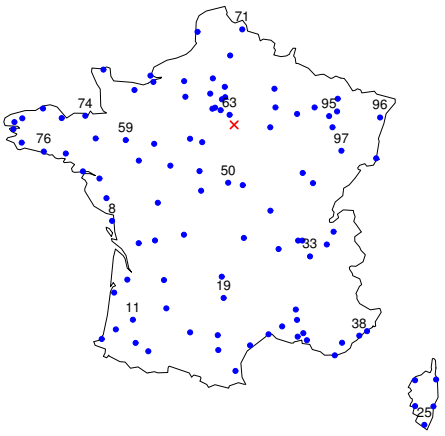

Figure A.1: Positions of 100 weather stations delivering Metar reports in Metropolitan France (blue dots) and of EDF Lab Les Renardières (red cross). The numbered stations form the test set.

number of daily observations is very often smaller than 48, due to numerous missing reports. In total, there are 62% of missing cloud coverage data, and 36% for the other variables. The cloud coverage is actually reported under the form of an abbreviation that corresponds to oktas (eighths), it corresponds to the fraction of the sky occupied by the clouds (from 0/8 to 8/8). The following transformation is applied: the value is set to 0 when there is no significant cloud, 1 for a full coverage, whereas "FEW" corresponds to 1.5/8, "SCT" to 3.5/8, "BKN" to 6/8. Hence, these data take a finite number of values between 0 and 1. However, for the prediction, we will consider that the index varies continuously between 0 and 1 since it makes more sense physically.

A second dataset consists of records at time step 1 minute of the global, direct and diffuse solar irradiances at a specific location, namely EDF Lab Les Renardières, with latitude 48.38 N and longitude 2.83 E (see Figure A.1). The air temperature, relative humidity, atmospheric pressure and wind speed were also collected at the same rate.

A.3 Spatial models

In this section and the following, we consider a regionalized variable $Y(\cdot)$, that is a random variable that depends on the space and time coordinates. If Y is univariate, it represents the cloud cover process. In the multivariable case, Y includes the cloud cover and other covariables, such as the wind speed for instance. $s_1 \in D, \ldots, s_n \in D$ denote a set of $n \geq 1$ locations in $D \subset \mathbb{R}^2$ where the variable $Y(s_i, t)$ has been measured at time $t \in \mathbb{R}$. $s_0 \in D$ is the site at which we wish to predict the cloud cover. $|.|$ denotes the Euclidian norm of a vector, whether it be $s \in D$ or $t \in \mathbb{R}$.

A.3.1 Univariate case

As a first approach, we consider a purely spatial model for the cloud cover data only. That is, we assume that the data are independant both in time and of other covariates. The model assumption

is the following:

$$Y(s) := \mu + \delta(s), \qquad s \in D, \tag{A.3}$$

where $Y(\cdot)$ is scalar μ is an unknown constant mean, and $\delta(\cdot)$ a zero-mean random process with second-order moments, called correlated error process. We assume that for all $(s_1, s_2) \in D^2$, the variogram

$$2\gamma(|s_1 - s_2|) := \text{var}\left[Y(s_1) - Y(s_2)\right] \tag{A.4}$$

is well defined (and depends therefore only on the space increment). The aim is to predict $Y(s_0)$ according to the weighted average

$$\hat{Y}(s_0) = \sum_{i=1}^{n} \lambda_i Y(s_i), \qquad \sum_{i=1}^{n} \lambda_i = 1. \tag{A.5}$$

The latter condition guarantees the uniform unbiasedness of the predictor. This procedure is called ordinary kriging, and is described in details in [Cre91]. The main task is to obtain a valid variogram model. In particular, we need to make sure that the variogram is conditionally negative definite. Consequently, the variogram fitting process is usually carried out in two steps:

1. estimate an experimental variogram, according to the formula

$$\gamma(h(k)) = \frac{1}{2|N(h(k))|} \sum_{N(h(k))} \left(y(s_i) - y(s_j)\right)^2, \tag{A.6}$$

   where $N(h(k)) = \left\{(s_i, s_j) \in D^2 | (s_i, s_j) \in h(k)\right\}$ and $|N(h(k))|$ is the cardinal of the set $N(h(k))$;

2. fit a valid parametric model to the experimental variogram. Standard models, including the Gaussian, exponential, linear, etc. families are given in [Cre91].

Here, we estimated an averaged experimental variogram. This means that an experimental variogram was computed at each time step for the same spatial blocks $h(k)$. The averaged variogram is made of averages in time within each block. The center of the block is the average distance between the pairs of data in the block. Due to erratic missing data, this center varies slightly in time, hence the averaged variogram also has averaged block centers. The fitted model is exponential, of the form

$$\gamma(h) := b\left(1 - \exp\left(-\frac{|h|}{a}\right)\right), \qquad h \in \mathbb{R}^2, \tag{A.7}$$

for some $a > 0$, $b > 0$, with an additional nugget effect (discontinuity at the origin, with a non-zero right value). a is called the range of the model, b the (partial) sill. The model is nonlinear in the parameters a, b and the nugget effect. Those are estimated by minimizing a weighted least square criterion, with weights $w_k \propto \frac{|N(h(k))|}{|h(k)|^2}$, by an iterative gradient descent algorithm (Gauss-Newton algorithm). Figure A.2 displays the experimental variogram and the fitted Gaussian model.

Once the variogram model $\gamma(\cdot)$ is fitted, the kriging equations [Cre91] give the prediction at any location s_0. The weights λ_i, $i = 1, \ldots, n$ are found by minimizing the error variance estimation

$$E\left[\left(Y(s_0) - \hat{Y}(s_0)\right)^2 | Y_1, \ldots, Y_n\right] = E\left[\left(Y(s_0) - \sum_{i=1}^{n} \lambda_i Y(s_i)\right)^2\right]. \tag{A.8}$$

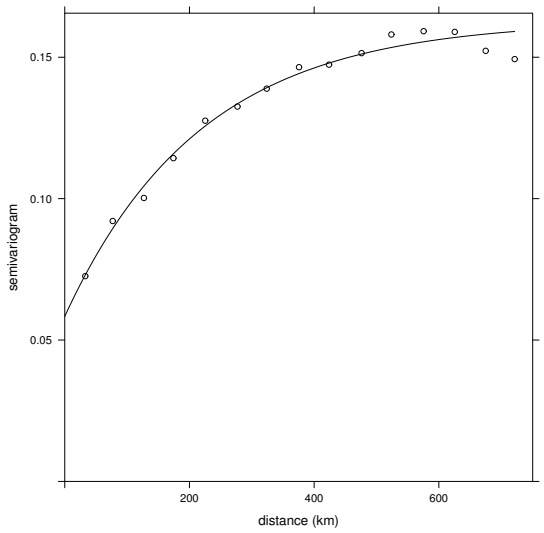

Figure A.2: Averaged experimental variogram of the cloud coverage data in December 2015, and fitted exponential model.

With m the Lagrange multiplier associated to the constraint $\sum_{i=1}^{n} \lambda_i = 1$ and

$$\boldsymbol{\lambda} := (\lambda_1, \dots, \lambda_n)^T, \tag{A.9}$$

$$\boldsymbol{\gamma} := (\gamma(s_0 - s_1), \dots, \gamma(s_0 - s_n))^T, \tag{A.10}$$

$$\Gamma := [\gamma(s_i - s_j)],\ 1 \leq i, j \leq n, \tag{A.11}$$

one obtains

$$\boldsymbol{\lambda}^T = \left(\boldsymbol{\gamma} + \mathbf{1}\frac{(1 - \mathbf{1}\Gamma^{-1}\boldsymbol{\gamma})}{\mathbf{1}^T\Gamma^{-1}\mathbf{1}}\right)^T \Gamma^{-1}, \qquad m = -\frac{(1 - \mathbf{1}\Gamma^{-1}\boldsymbol{\gamma})}{\mathbf{1}^T\Gamma^{-1}\mathbf{1}}. \tag{A.12}$$

$$\tag{A.13}$$

The minimized mean-squared prediction error, called the kriging variance, is

$$\sigma_k^2(s_0) = \boldsymbol{\lambda}^T\boldsymbol{\gamma} + m = 2\sum_{i=1}^{n} \lambda_i\gamma(s_0 - s_i) - \sum_{i,j=1}^{n} \lambda_i\lambda_j\gamma(s_i - s_j). \tag{A.14}$$

For numerical evaluation of the model, we randomly split the data from December 2015 between training set (85 stations) and test set (15 stations, with location shown in Figure A.1). The code is written in `R`, using more particularly the package `gstat` with the functions `fit.variogram` for variogram fitting and `krige` for the ordinary kriging. We perform leave-one-out cross validation (LOOCV) on the training set, with `gstat::krige.cv`. The relative errors are plotted in Figure A.3. The total RMSE (root mean squared error) over the 15 test stations is 0.3339, whereas the total MAE (mean absolute error) is 0.2805. These performance indicators vary depending on the station, see Figure A.4(c) and A.4(d): it seems that the estimation is more accurate when there are closer neighbours in the training set. However, this dependency is less obvious from a boxplot perspective, see Figure A.4(b).

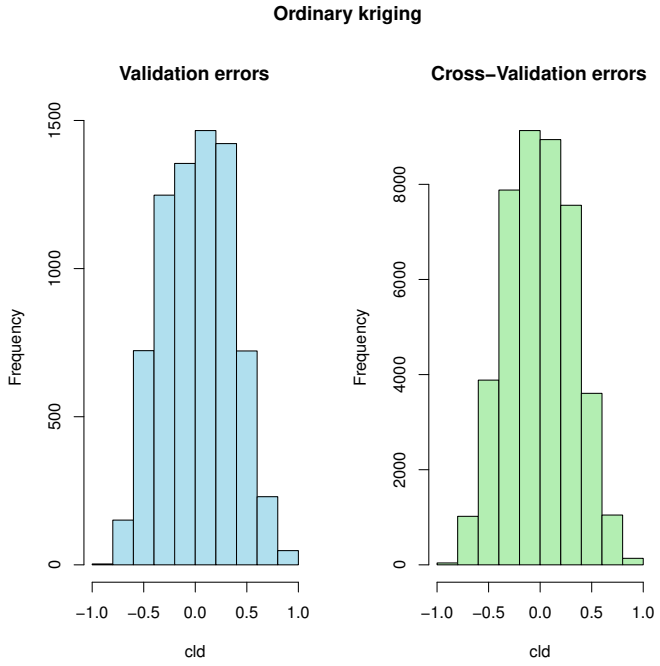

Figure A.3: (Left) Histogram of the relative errors for the 15 test locations. (Right) Histogram of the relative errors for the 85 LOOCV predictions for the training set.

A.3.2 Multivariate model

A first improvement of the univariate spatial model consists in keeping the time independency assumption, but adding one or several covariables to the model. Let q denote the number of covariables, resulting in a random vector Y of dimension $q+1$ when adding the cloud cover data. The theory of ordinary kriging adapts easily to the multivariate case, see e.g. [VHC93]. First of all, the variogram is extended to a cross variogram according to:

$$2\gamma_{ij}(|s_1 - s_2|) := E\left[(Y_i(s_1) - Y_i(s_2))\,(Y_j(s_1) - Y_j(s_2))\right], \tag{A.15}$$

for $1 \leq i, j \leq q+1$ and for all $s_1, s_2 \in D$, and where it is implicitly assumed that each component Y_i has a constant mean. In the multivariate case, it is more difficult to define a valid cross-variogram model. The cross-structures between variables are not independent of the simple structures of these variables. Hence, the cross-variograms $\gamma_{ij}(\,\cdot\,)$ cannot be modeled independently. A standard class of valid models is the linear model of coregionalization (LMC, see [GV92]): all simple and cross structures are linear combinations of the same (univariate) variogram models. For Y with constant mean, the LMC is

$$\Gamma(h) := E\left[(Y(s+h) - Y(s))(Y(s+h) - Y(s))^\intercal\right] = \sum_{k=1}^{m} S_k g_k(h), \tag{A.16}$$

where m is the number of known variogram structures g_k (for instance spherical, Gaussian, exponential, etc.) and the S_k's are non-negative unknown matrices of sills.

Similarly to the univariate case, the task of fitting the cross-variogram is processed in two steps:

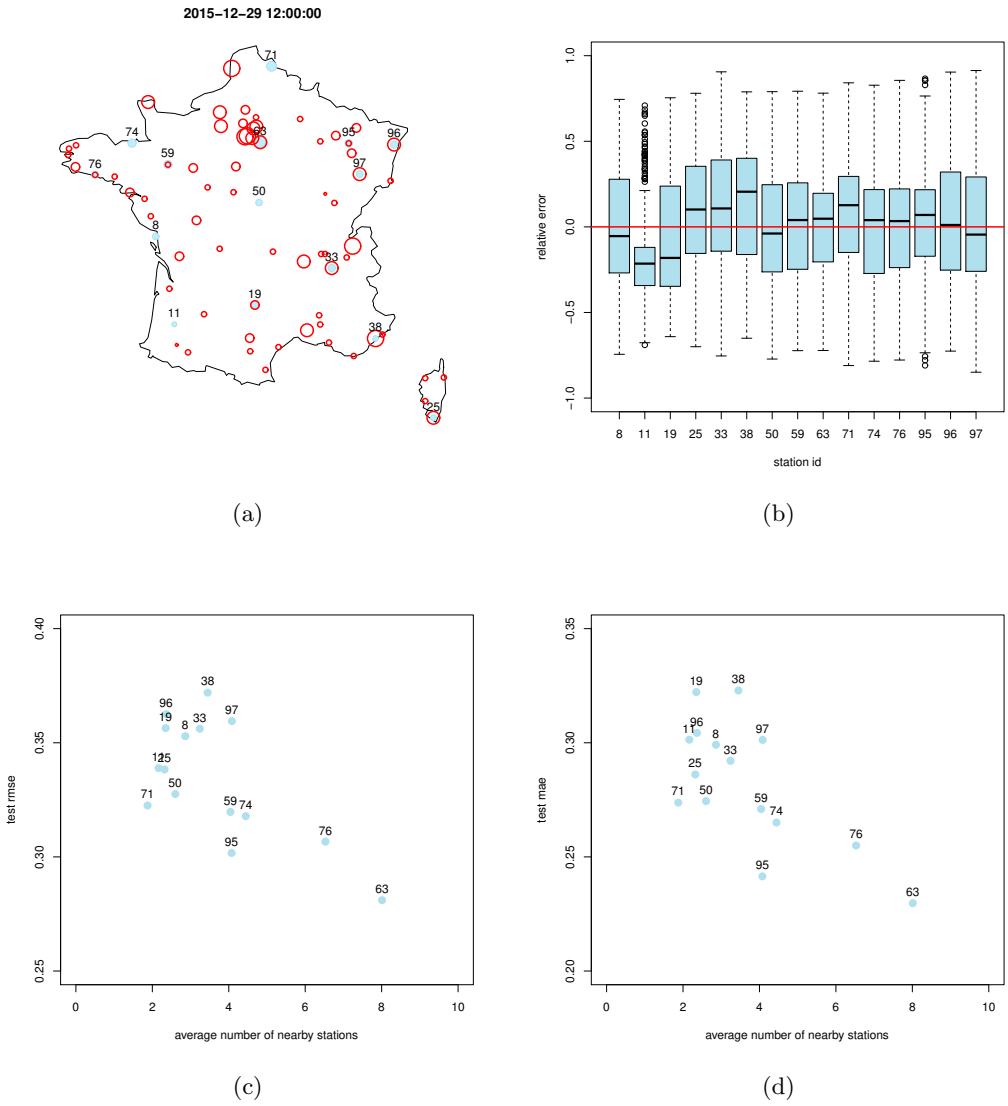

(a) (b)

(c) (d)

Figure A.4: Prediction of the cloud cover by univariate ordinary kriging. (a) An example of estimation at fixed time. Each circle is centered at sample point with radius proportional to the actual cloud cover. Test predictions are the filled discs, the red circles represent measured values. (b) Boxplot of the relative estimation errors for the test set, computed (c) RMSE and (d) MAE at test locations, depending on the number of close neighbours (< 150 km) in the training set.

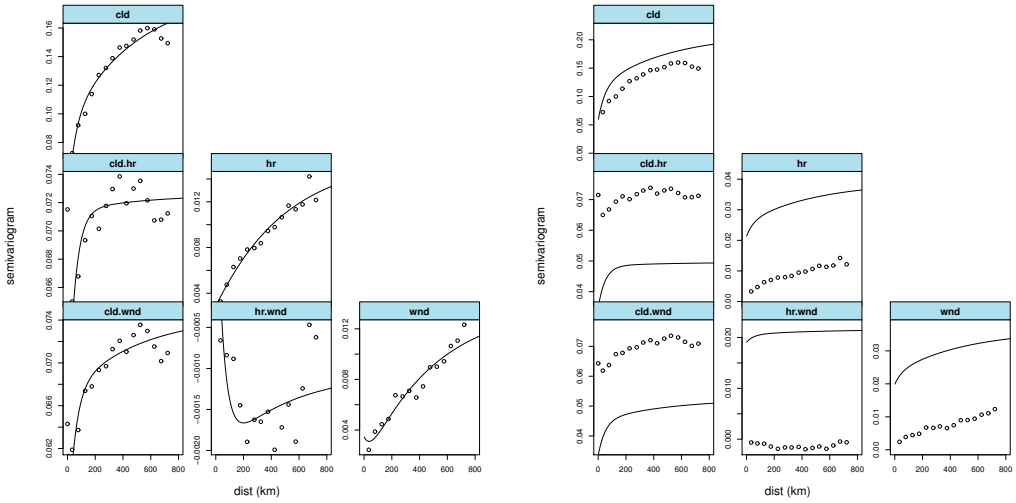

(a) (b)

Figure A.5: Averaged experimental cross-variogram of the cloud cover, relative humidity and wind speed data (points) in December 2015, and fitted linear model of coregionalization (line) (a) with unconstrained and (b) positive semidefinite matrices of partial sills. The basic components of the LMC are Nug, Exp(50), Exp(500).

estimation of the experimental cross-variogram, according to the formula

$$\gamma_{ij}(h(k)) = \frac{1}{2|N(h(k))|} \sum_{N(h(k))} \left(y_i(s_a) - y_i(s_b)\right)\left(y_j(s_a) - y_j(s_b)\right), \tag{A.17}$$

and then fit a model cross-variogram. This fitting step is not trivial, because of the constraint that each S_k must be non-negative definite. [GV92] suggests a method to estimate the matrices S_k in the LMC (A.16), by minimizing a weighted sum of squares in an iterative procedure. They emphasize that it is important to choose an appropriate combination of the elementary models g_k. The range of the elementary models is also a key parameter (it is chosen *a priori* and not fitted). We refer to [PDLF04, Eme10] for a study of other iterative algorithms.

An example of fitted cross-variogram is given in Figure A.5(b), with vector Y including the cloud cover, the hygrometry (normalized between 0 and 1) and the wind speed (normalized between 0 and 1). We used our own version of the function `gstat::fit.lmc` in order to handle missing values. This function uses a simple fitting procedure: first, each variogram model is fitted to a direct or cross variogram, by minimizing a weighted sum of squares; next each of the partial sill coefficient matrices is approached by its closest positive semidefinite matrix in least square sense. The matrices of partial sills are therefore the projections of the fitted models on the cone of positive semidefinite matrices. In Figure A.5, the model is the sum of three components: a nugget effect, an exponential model with short range (50 km) and a second exponential model with longer range (500 km). The difference between Figures A.5(a) and A.5(b) shows how the constraints of positive semidefiniteness deteriorates the fit.

The cokriging equations are obtained in the multivariate case similarly to the univariate case, that is by looking for the best linear unbiased predictor that minimizes the variance of the error prediction. Finally, the estimator is not constrained to the segment $[0, 1]$. Hence, in order to make the estimations physically plausible, we simply projected them onto $[0, 1]$.

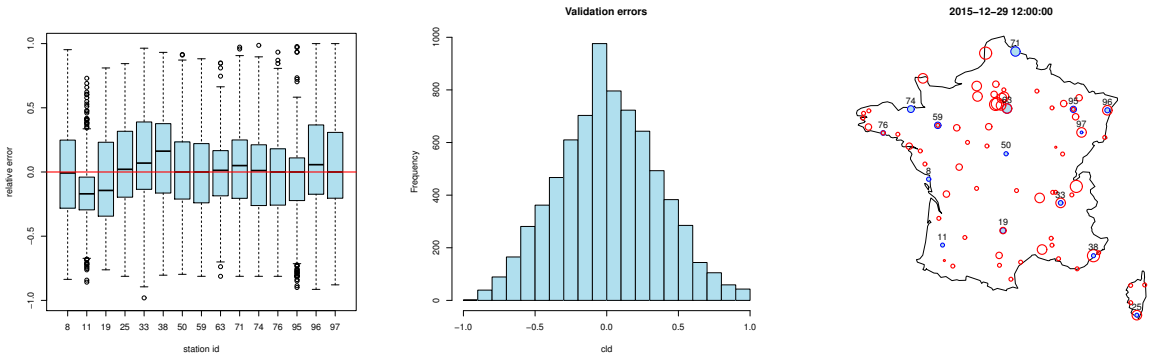

(a) (b) (c)

Figure A.6: Prediction of the cloud cover by multivariate ordinary kriging. (a) Boxplot of the relative estimation errors for the test set with the multivariable model with components cloud cover, hygrometry, wind speed. (b) Overall corresponding histogram. (c) An example of estiamtion at fixed time. Each circle is centered at sample point with radius proportional to the actual cloud cover. Test predictions are the filled discs, the red circles represent measured values.

We evaluated numerically the multivariate spatial model with the same learning and test sets as in the univariate case. The covariables are the hygrometry and the wind speed, with cross-variogram model given in Figure A.5. The distributions of the relative prediction errors are given in Figure A.6(a) (by station) and Figure A.6(b) (global histogram). Compared to the univariate case, the errors are still high. The global RMSE is actually 0.3486, the global MAE is 0.2778, that is the multivariate model has (slightly) worse performances than the univariate model, in terms of RMSE. The effect of resorting to covariates in this case is more visible on the variance of the prediction error. Indeed, this variance is much smaller in the multivariate case than in the univariate: in average, we obtained a variance of 0.1147 for the univariate model against 0.0422 for the multivariate model.

Given these first results, we see that the multivariable model needs improvement. An in-depth study of the multivariate spatial model would investigate some of the following issues:

- the impact of the choice of covariables: there are probably other meteorological variables that are more strongly correlated to the cloud cover data;
- the comparison of the LMC obtained via our simple procedure to those obtained using the methodology described in either [GV92] or [PDLF04, Eme10];
- the comparison of the LMC to other theroretical multivariate variogram models.

A.4 Comparison of the models

Several spatial models for the cloud cover were presented in Section A.3. They were tested in several Metar stations against the actual value of the cloud cover. Yet, our primary goal is the estimation of the global solar radiation. Hence, this section details how to obtain a prediction of the global solar radiation from the cloud cover data, and compares the different spatial models against measurements of the solar radiation at Les Renardières.

Table A.1: Perrin de Brichambault's model with parameter $kappa = 0.81$, cloud correction with $a = 0.67$ and $b = 1.08$, evaluated at Les Renardières in Dec. 2015. The RMSE and MAE are given in W/m^2. S-U is the spatial univariate kriging model, S-M the spatial multivariate. GT is the approximate ground truth, with cloud cover directly estimated from the clearness index.

Model	RMSE	NRMSE	MAE	NMAE	R^2
S-U	29.4	0.51	23.4	0.40	0.54
S-M	30.5	0.53	23.7	0.41	0.55
GT	3.92	0.20	1.94	0.10	0.98

A.4.1 Prediction of the solar irradiance at Les Renardières

To infer solar radiation from cloud cover, we adopted the strategy of correcting a clear-sky model by an adjustment factor that takes into account the effective cloudiness of the sky. The chosen clear-sky model, denoted PdBV after Perrin de Brichambaut and Vauge [PdBV82, RW99] uses the solar zenithal angle as a unique input and reads

$$G_{h,clr} = \kappa G_0 (\cos\theta_z)^{1.15} \tag{A.18}$$

with the same notations as in the introduction to this appendix. The model is corrected by a measure $N \in [0, 1]$ of the cloud cover according to

$$G_{h,cld} = G_{h,clr}(1 - aN^b). \tag{A.19}$$

Note that the same clear-sky and cloud correction model are employed in Chapter 6. Our approach combines thus two types of errors, namely the cloud cover estimation error and the model error, due to the assumption that the horizontal global solar flux is accurately described by eq. (A.18)-(A.19). Unfortunately, no measure of the cloud cover is available at Les Renardières, such that it is difficult to evaluate the second error. However, there are some models that interpolate linearly the cloud cover N from the clearness index k. k is defined as the ratio of the received global solar radiation during a given period to the extraterrestrial global horizontal radiation for this period, such that it can be computed at Les Renardières. Such linear models have a simple structure, N being set to 1 for $k \leq k_{cld}$, 0 for $k \geq k_{clr}$ and linearly interpolated inbetween. Good values of k_{cld} and k_{clr} are explored in [FXM+09], we can use for instance $k_{cld} = 0.25$ and $k_{clr} = 0.8$.

Here, we have chosen $k_{cld} = 0.2$ and $k_{clr} = 0.7$. Then, the parameters a and b from (A.19) were computed using a least square fit to the data recorded in November 2015. This gives $a = 0.67$ and $b = 1.08$. These values are lower than the values $a = 0.75$ and $b = 3.4$ obtained from a fit to a 10-year record at Hambourg in [KC80], or to the values $a = 0.79$ and $b = 3.21$ fitted to hourly measurements at the lowland site of Baumgarten, Germany, in [IM02].

We then compared the estimations of the global solar radiation at Les Renardières in December 2015 (during daytime hours only), with cloud cover given by the different kriging models and fitted correction parameters $a = 0.67$ and $b = 1.08$. The results are summarized in Table A.1 and in Figure A.7. Additionally to the RMSE and MAE that were already used in the previous sections, we propose the normalized RMSE (NRMSE), the normalized MAE (NMAE) and the goodness of fit (R^2) as error metrics for evaluating the different models. The NRMSE (respectively the NMAE) is the RMSE (respectively the MAE) divided by the mean of the observed global solar radiation during daytime.

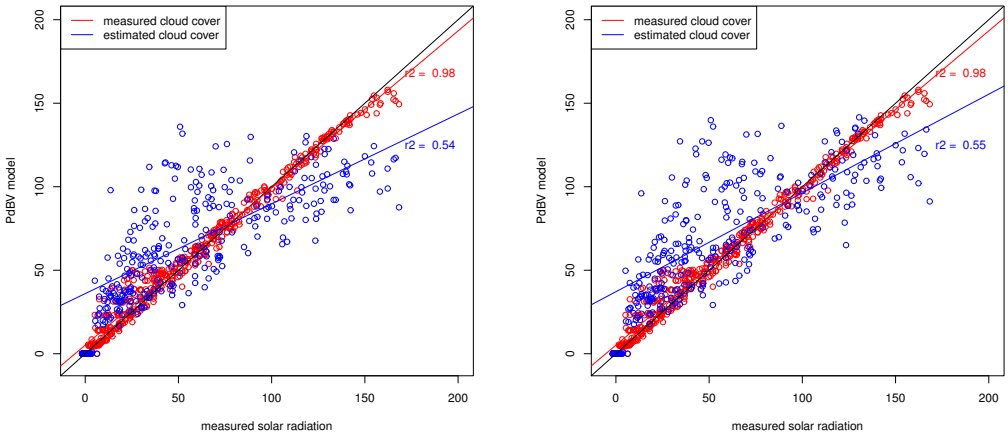

(a) (b)

Figure A.7: Solar radiation predicted in December at Les Renardières by the PdBV model combined with cloud correction, against actual values. The cloud cover is estimated from the actual solar radiation (red points) or from (a) the univariate, (b) the multivariate, spatial ordinary kriging technique (blue points).

Figure A.7 justifies the choice of PdBV's model. Indeed, if it is corrected by a good estimation of the cloud cover (red points), then the estimation of the global solar radiation follows closely the actual measure (red points close to the main diagonal, correlation coefficient $R^2 = 0.98$). This is confirmed by Table A.1: the PdBV model with experimental cloud cover data has good error metrics, NRMSE of 20% and NMAE of 10%. Hence, resorting to predicting the cloud cover rather than the solar radiation directly is a reasonable strategy.

We have seen previously that all three kriging models suffered from severe bias. This remains true here, with NRMSE over 50% and NMAE over 40%. This means that the kriging models for estimating the cloud cover need to be further explored in order to achieve better results.

A.4.2 Comparison to the literature

Standard non-kriging learning methods of the literature, based on available records of the global horizontal solar radiation, largely outperform our approach. For instance, [JPML16] typically obtains NRMSE of 23% and 25%, and gives similar values from other algorithms (between 22% and 37% depending on the studies). In [GFDD16], the best models for hourly predictions have MAE between 14.84 W/m^2 and 21.79 W/m^2. Usually, these studies use data from very sunny locations (Spain and Mediterranean countries), which is a rather favorable context.

[YGD+13] uses anisotropic space-time simple kriging for hourly time-forward predictions of the solar radiation in 10 stations deployed in Singapore. Their results are given in terms of RMSE of the hourly clearness index data at 10 stations for two weeks of predictions in late November 2012 (and from two weeks of learning data in early Nov. 2012). Their best covariance model achieves an overall RMSE of 0.144, whereas the simple persistance method yields 0.193. In our setting, different because we do not predict the future but reconstruct the present, we obtained RMSEs of 0.175 (univariate spatial ordinary kriging), 0.184 (multivariate spatial ordinary kriging) for the half-hourly clearness indices.

Interestingly, some articles address the issue of estimating the solar radiation in the context of building performance simulation. For instance, [OS14] compares different method for restoring missing gaps in solar radiation measurements. Given a sequence of hourly measurements of the solar radiation, they want to fill gaps of several days in these records. Hence, for a given gap, measurements of the solar radiation prior to and posterior to the gap are still available. In Oklahoma City North, which has a warm and humid climate, and for a gap of 20 days the different methods had NRMSE between 52.7% and 60.3%, NMAE between 28.0% and 36.0%, R^2 between 0.76 and 0.88. Those results are more in line with our findings, and illustrate the difficulty of estimating the solar radiation in the absence of measurements.

A.5 Discussion and conclusion

The aim of this work was to propose a statistical model for the global solar radiation at a given specific location that takes the spatial dependence structure into account. Moreover, we added the constraint that a proper learning set, with historical records of the solar radiation, was not available. The underlying idea being to build an estimator that requires a minimal specific instrumentation. This is why we chose to estimate the cloud cover instead, and correct a simple clear-sky model to obtain the estimated solar radiation. Our approach combines thus two sources of errors, the physical model and the cloud cover adjustment model. Moreover, the cloud cover data itself is difficult to measure, resulting in a dataset with poor accuracy, combined with a loose spatial resolution (100 stations in France). Hence, a reasonable target was to obtain a first "rather good" estimation, with evaluation of the uncertainty, rather than the "best" estimator.

Beyond the intrinsic constraints of the problem, we emphasize the fact that the kriging method depends on several modelling choices. First of all, at time step 30 minutes, we assumed that a purely spatial model was sufficient for the cloud cover estimating. Nevertheless, it would be worthwhile investigating the influence of the time-step on the model, although higher frequency measurements of the cloud cover do not exist yet to the best of our knowledge. We also chose a model with constant unknown mean and isotropic regionalized error process, which corresponds to the framework of ordinary kriging. Other models exist (e.g. universal kriging), and in particular we might question the assumption of spatial isotropy. We have already noted that anisotropy was indeed taken into account in similar studies such as [ISW12] or [YGD+13], for denser spatial networks.

From a practical point of view, kriging requires also to select carefully several parameters: the space (and, possibly, time) resolution and the cutoff of the experimental variograms, the initialization of the parameters for variogram fitting. We have chosen them "manually", but one could include a cross-validation procedure to select them more accurately.

As a conclusion, this work shows that it is difficult to estimate the solar radiation without resorting to supervised learning. Using cloud cover data, freely accessible, within the kriging frameworks looked promising. However, the first basic models are still largely perfectible, and this short discussion highlights the main steps that could be improved in the estimation process.

Bibliography

[AAC+15] Fatima Amara, Kodjo Agbossou, Alben Cardenas, Yves Dubé, and Sousso Kelouwani. Comparison and simulation of building thermal models for effective energy management. *Smart Grid and Renewable Energy*, 6(4):95–112, 2015.

[AAP+16] Manar Amayri, Abhay Arora, Stephane Ploix, Sanghamitra Bandhyopadyay, Quoc-Dung Ngo, and Venkata Ramana Badarla. Estimating occupancy in heterogeneous sensor environment. *Energy and Buildings*, 129:46–58, 2016.

[ADE13] ADEME. Les chiffres clés du bâtiment 2013, 2013.

[ADFDJ03] Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. An introduction to MCMC for machine learning. *Machine learning*, 50(1-2):5–43, 2003.

[ADST04] Christophe Andrieu, Arnaud Doucet, Sumeetpal S. Singh, and Vladislav B. Tadic. Particle methods for change detection, system identification, and control. *Proceedings of the IEEE*, 92(3):423–438, 2004.

[AGG16] R Azimi, M Ghayekhloo, and M Ghofrani. A hybrid method based on a new clustering technique and multilayer perceptron neural networks for hourly solar radiation forecasting. *Energy Conversion and Management*, 118:331–344, 2016.

[BB06] Stephen J. Blundell and Katherine M. Blundell. *Concepts in thermal physics*. Oxford University Press, 2006.

[BC13] Khalil Benmouiza and Ali Cheknane. Forecasting hourly global solar radiation using hybrid k-means and nonlinear autoregressive neural network models. *Energy Conversion and Management*, 75:561–569, 2013.

[BdSACM16] Peder Bacher, Philip Anton de Saint-Aubain, Lasse Engbo Christiansen, and Henrik Madsen. Non-parametric method for separating domestic hot water heating spikes and space heating. *Energy and Buildings*, 130:107–112, 2016.

[Bee16] Nathanael Beeker. *Modeling and control of electric hot water tanks: from the single unit to the group*. PhD thesis, Mines ParisTech, 2016.

[Ber13] Thomas Berthou. *Développement de modèles de bâtiment pour la prévision de charge de climatisation et l'élaboration de stratégies d'optimisation énergétique et d'effacement*. PhD thesis, École nationale supérieure des Mines de Paris, 2013.

[BHZ14] Michael C Burkhart, Yeonsook Heo, and Victor M Zavala. Measurement and verification of building systems under uncertain data: A gaussian process modeling approach. *Energy and Buildings*, 75:189–198, 2014.

[BM11] Peder Bacher and Henrik Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.

[BP66] Leonard E Baum and Ted Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.

[BPR+13] Alexandros Beskos, Natesh Pillai, Gareth Roberts, Jesus-Maria Sanz-Serna, Andrew Stuart, et al. Optimal tuning of the hybrid monte carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.

[BPSSS11] Alexandros Beskos, Frank J. Pinski, Jesús Marıa Sanz-Serna, and Andrew M. Stuart. Hybrid Monte Carlo on Hilbert spaces. *Stochastic Processes and their Applications*, 121(10):2201–2230, 2011.

[BSSM14] Thomas Berthou, Pascal Stabat, Raphael Salvazet, and Dominique Marchio. Development and validation of a gray box model to predict thermal behavior of occupied office buildings. *Energy and Buildings*, 74:91–100, 2014.

[Cai87] Peter E Caines. *Linear stochastic systems*. John Wiley & Sons, Inc., 1987.

[CBL06] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[CDRR12] Lou Chesné, Thierry Duforestel, Jean-Jacques Roux, and Gilles Rusaouën. Energy saving and environmental resources potentials: Toward new methods of building design. *Building and Environment*, 58:199–207, 2012.

[CGM07] Olivier Cappé, Simon J Godsill, and Eric Moulines. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924, 2007.

[CMR05] Olivier Cappé, Éric Moulines, and Tobias Rydén. *Inference in hidden Markov models*, volume 6. Springer, 2005.

[Cre91] Noel Cressie. *Statistics for spatial data*. Wiley series in probability and mathematical statistics. John Wiley & Sons, 1991.

[CSM+13] Giuseppe Tommaso Costanzo, Fabrizio Sossan, Mattia Marinelli, Peder Bacher, and Henrik Madsen. Grey-box modeling for system identification of household refrigerators: A step toward smart appliances. In *4th International Youth Conference on Energy (IYCE)*, pages 1–5. IEEE, 2013.

[DB13] John A Duffie and William A Beckman. *Solar engineering of thermal processes*. Wiley New York, fourth edition, 2013.

[DBMM10] Kun Deng, Prabir Barooah, Prashant G Mehta, and Sean P Meyn. Building thermal model reduction via aggregation of states. In *American Control Conference (ACC), 2010*, pages 5118–5123. IEEE, 2010.

[DGA00] Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential monte carlo sampling methods for bayesian filtering. *Statistics and computing*, 10(3):197–208, 2000.

[DGM+11] Randal Douc, Aurélien Garivier, Eric Moulines, Jimmy Olsson, et al. Sequential Monte Carlo smoothing for general state space hidden markov models. *The Annals of Applied Probability*, 21(6):2109–2145, 2011.

[DJ09] Arnaud Doucet and Adam M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, 12:656–704, 2009.

[DKPR87] Simon Duane, Anthony D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987.

[DLR77] Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (methodological)*, pages 1–38, 1977.

[DM01] Randal Douc and Catherine Matias. Asymptotics of the maximum likelihood estimator for general hidden Markov models. *Bernoulli*, 7(3):381–420, 2001.

[DMS14] Randal Douc, Éric Moulines, and David Stoffer. *Nonlinear time series: Theory, Methods and Applications with R examples*. CRC Press, 2014.

[EG09] Mabrouka El Guedri. *Caractérisation aveugle de la courbe de charge électrique: Détection, classification et estimation des usages dans les secteurs résidentiel et tertiaire*. PhD thesis, Université Paris Sud-Paris XI, 2009.

[Eme10] Xavier Emery. Iterative algorithms for fitting a linear model of coregionalization. *Computers & Geosciences*, 36(9):1150–1160, 2010.

[FABG14] Samuel F Fux, Araz Ashouri, Michael J Benz, and Lino Guzzella. EKF based self-adaptive thermal model for a passive house. *Energy and Buildings*, 68:811–817, 2014.

[FL99] Urban Forssell and Lennart Ljung. Closed-loop identification revisited. *Automatica*, 35(7):1215–1241, 1999.

[FMXY12] Xi Fang, Satyajayant Misra, Guoliang Xue, and Dejun Yang. Smart grid - the new and improved power grid: a survey. *IEEE communications surveys & tutorials*, 14(4):944–980, 2012.

[FVLA02] Gilles Fraisse, Christelle Viardot, Olivier Lafabrie, and Gilbert Achard. Development of a simplified and accurate building model based on electrical analogy. *Energy and Buildings*, 34(10):1017–1031, 2002.

[FXM+09] G N Flerchinger, Wei Xiao, Danny Marks, T J Sauer, and Qiang Yu. Comparison of algorithms for incoming atmospheric long-wave radiation. *Water Resources Research*, 45:W03423, 2009.

[Gai15] Pierre Gaillard. *Contributions to online robust aggregation: work on the approximation error and on probabilistic forecasting. Applications to forecasting for energy markets*. PhD thesis, Université Paris-Sud 11, 2015.

[GAO+16] Santosh K Gupta, Sam Atkinson, Ian O'Boyle, John Drogo, Koushik Kar, Sandipan Mishra, and John T Wen. BEES: Real-time occupant feedback and environmental learning framework for collaborative thermal management in multi-zone, multi-occupant buildings. *Energy and Buildings*, 125:142–152, 2016.

[GCL06] Valentine Genon-Catalot and Catherine Laredo. Leroux's method for general hidden Markov models. *Stochastic processes and their applications*, 116(2):222–243, 2006.

[GDW04] Simon J. Godsill, Arnaud Doucet, and Mike West. Monte Carlo smoothing for non-linear time series. *Journal of the american statistical association*, 99(465), 2004.

[GFDD16] Yvonne Gala, Ángela Fernández, Julia Díaz, and José R Dorronsoro. Hybrid machine learning forecasting of solar radiation values. *Neurocomputing*, 176:48–59, 2016.

[GG16] Pierre Gaillard and Yannig Goude. `opera`: Online Prediction by ExpeRts Aggregation., 2016.

[Gof13] Jeanne Goffart. *Impact de la variabilité des données météorologiques sur une maison basse consommation. Application des analyses de sensibilité pour des entrées temporelles.* PhD thesis, Université de Grenoble, 2013.

[GSS93] Neil J. Gordon, David J. Salmond, and Adrian F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113, 1993.

[Gus10] Fredrik Gustafsson. Particle filter theory and practice with positioning applications. *Aerospace and Electronic Systems Magazine, IEEE*, 25(7):53–82, 2010.

[GV92] Michel Goulard and Marc Voltz. Linear coregionalization model: tools for estimation and choice of cross-variogram matrix. *Mathematical Geology*, 24(3):269–286, 1992.

[HD88] Edward James Hannan and Manfred Deistler. *The statistical theory of linear systems*, volume 70. SIAM, 1988.

[HGP12] Ion Hazyuk, Christian Ghiaus, and David Penhouet. Optimal temperature control of intermittently heated buildings using model predictive control: Part I - building modeling. *Building and Environment*, 51:379–387, 2012.

[HGP14] Ion Hazyuk, Christian Ghiaus, and David Penhouet. Model predictive control of thermal comfort as a benchmark for controller performance. *Automation in Construction*, 43:98–109, 2014.

[IM02] M G Iziomon and H Mayer. Assessment of some global solar radiation parameterizations. *Journal of Atmospheric and Solar-Terrestrial Physics*, 64(15):1631–1643, 2002.

[ISW12] Tsuyoshi Inoue, Tetsuo Sasaki, and Takashi Washio. Spatio-temporal kriging of solar radiation incorporating direction and speed of cloud movement. In *The 26th Annual Conference of the Japanese Society for Artificial Intelligence, Yamaguchi City*, 2012.

[JMA08] María José Jiménez, Henrik Madsen, and Klaus Kaae Andersen. Identification of the main thermal characteristics of building components using Matlab. *Building and Environment*, 43(2):170–180, 2008.

[JPML16] Pedro F Jiménez-Pérez and Llanos Mora-López. Modeling and forecasting hourly global solar radiation using clustering and classification techniques. *Solar Energy*, 135:682–691, 2016.

[KC80] Fritz Kasten and Gerhard Czeplak. Solar and terrestrial radiation dependent on the amount and type of cloud. *Solar energy*, 24(2):177–189, 1980.

[KDSM09] Nicholas Kantas, Arnaud Doucet, Sumeetpal Sindhu Singh, and Jan Marian Maciejowski. An overview of sequential Monte Carlo methods for parameter estimation in general state-space models. In *15th IFAC Symposium on System Identification (SYSID), Saint-Malo, France (invited paper)*, volume 102, 2009.

[KS05] Jari Kaipio and Erkki Somersalo. *Statistical and computational inverse problems*, volume 160. Springer, 2005.

[Laj11] Christian Lajaunie. Simulation stochastique par trajectoires hamiltoniennes. Technical report, Mines ParisTech, 2011.

[LBS+16] Fredrik Lindsten, Pete Bunch, Simo Särkkä, Thomas B Schön, and Simon J Godsill. Rao-Blackwellized particle smoothers for conditionally linear gaussian models. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):353–365, 2016.

[Lin11] Fredrik Lindsten. *Rao-Blackwellised particle methods for inference and identification*. PhD thesis, Linköping University, 2011.

[Liu08] Jun S. Liu. *Monte Carlo strategies in scientific computing*. Springer Science & Business Media, 2008.

[Lju87] Lennart Ljung. *System identification: theory for the user*. Englewood Cliffs, NJ: Prentice-Hall, 1987.

[LLH08] Céline Levy-Leduc and Zaïd Harchaoui. Catching change-points with lasso. In *Advances in Neural Information Processing Systems*, pages 617–624, 2008.

[LMDP14] Audrey Le Mounier, Benoit Delinchant, and Stephane Ploix. Determination of relevant model structures for self-learning energy management system. In *Building Simulation Optimisation*, 2014.

[LS10a] Fredrik Lindsten and Thomas B. Schön. Identification of mixed linear/nonlinear state-space models. In *49th IEEE Conference on Decision and Control (CDC)*, pages 6377–6382. IEEE, December 2010.

[LS10b] Fredrik Lindsten and Thomas B. Schön. Inference in mixed linear / nonlinear state-space models using sequential Monte Carlo. Technical report, Linköping University Electronic Press, 2010.

[LS13] Fredrik Lindsten and Thomas B Schön. Backward simulation methods for Monte Carlo statistical inference. *Foundations and Trends® in Machine Learning*, 6(1):1–143, 2013.

[LW94] Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.

[LW14] Xiwang Li and Jin Wen. Review of building energy modeling for control and operation. *Renewable and Sustainable Energy Reviews*, 37:517–537, 2014.

[Mac89] Paul B. Mackenzie. An improved hybrid Monte Carlo method. *Physics Letters B*, 226:369–371, 1989.

[Mal12] Paul Malisani. *Dynamic control of energy in buildings using constrained optimal control by interior penalty.* PhD thesis, Ecole Nationale Supérieure des Mines de Paris, 2012.

[MCPF10] Paul Malisani, François Chaplais, Nicolas Petit, and Dominique Feldmann. Thermal building model identification using time-scaled identification methods. In *49th IEEE Conference on Decision and Control*, pages 308–315. IEEE, 2010.

[Mej11] Olfa Mejri. *Développement de méthodes de diagnostic énergetique des bâtiments.* PhD thesis, Bordeaux 1, 2011.

[MK08] Geoffrey McLachlan and Thriyambakam Krishnan. *The EM algorithm and extensions*, volume 382 of *Wiley Series in Probability and Statistics*. John Wiley & Sons, second edition, 2008.

[MKT15] Eoghan McKenna, Michal Krawczynski, and Murray Thomson. Four-state domestic building occupancy model for energy demand simulations. *Energy and Buildings*, 96:30–39, 2015.

[MR91] Alain Marti and Jean-Pierre Rignac. Méthode de résolution du système différentiel utilisé par le simulateur thermique de stratège 2.1. Technical report, EDF R&D, 1991.

[MZ04] A Mechaqrane and M Zouak. A comparison of linear and neural network ARX models applied to a prediction of the indoor temperature of a building. *Neural Computing & Applications*, 13(1):32–37, 2004.

[NBR+14] Iban Naveros, Peder Bacher, Diego Pablo Ruiz, María José Jiménez, and Henrik Madsen. Setting up and validating a complex model for a simple homogeneous wall. *Energy and Buildings*, 70:303–317, 2014.

[Nea11] Radford M. Neal. MCMC using hamiltonian dynamics. In Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng, editors, *Handbook of Markov Chain Monte Carlo*. CRC Press, 2011.

[NGRC15] Iban Naveros, Christian Ghiaus, Diego Pablo Ruiz, and Sergio Castaño. Physical parameters identification of walls using arx models obtained by deduction. *Energy and Buildings*, 108:317–329, 2015.

[Øks03] Bernt Øksendal. *Stochastic Differential Equations*. Springer, 6th edition, 2003.

[OS14] Oluwqseyi T Ogunsola and Li Song. Restoration of long-term missing gaps in solar radiation. *Energy and Buildings*, 82:580–591, 2014.

[PCV+13] Samuel Prívara, Jiří Cigler, Zdeněk Váňa, Frauke Oldewurtel, Carina Sagerschnig, and Eva Žáčeková. Building modeling as a crucial part for building predictive control. *Energy and Buildings*, 56:8–22, 2013.

[PdBV82] Christian Perrin de Brichambaut and Christian Vauge. *Le gisement solaire: évaluation de la ressource énergétique*. Technique et Documentation (Lavoisier), 1982.

[PDLF04] Bernard Pelletier, Pierre Dutilleul, Guillaume Larocque, and James W Fyles. Fitting the linear model of coregionalization by generalized least squares. *Mathematical Geology*, 36(3):323–343, 2004.

[Pet69] Ted Petrie. Probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics*, 40(1):97–115, 1969.

[PLOP08] Luis Pérez-Lombard, José Ortiz, and Christine Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.

[PR15] Nicolas Petit and Pierre Rouchon. Dynamique et contrôle des systèmes - notes de cours, Mars 2015.

[PRMS08] Jessen Page, Darren Robinson, Nicolas Morel, and Jean-Louis Scartezzini. A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings*, 40:83–98, 2008.

[PSC+16] Marcelo Pereyra, Philip Schniter, Emilie Chouzenoux, Jean-Christophe Pesquet, Jean-Yves Tourneret, Alfred O Hero, and Steve McLaughlin. A survey of stochastic simulation and optimization methods in signal processing. *IEEE Journal of Selected Topics in Signal Processing*, 10(2):224–241, 2016.

[PŠFC11] Samuel Prívara, Jan Široký, Lukáš Ferkl, and Jiří Cigler. Model predictive control of a building heating system: The first experience. *Energy and Buildings*, 43(2):564–572, 2011.

[PWM16] Marco Pritoni, Jonathan M Woolley, and Mark P Modera. Do occupancy-responsive learning thermostats save energy? a field study in university residence halls. *Energy and Buildings*, 127:469–478, 2016.

[Ric91] Véronique Richalet. *Caractérisation énergétique des bâtiments sur site. Identification de modèles dynamiques. Méthodes de signature energétique.* PhD thesis, INPG, 1991.

[RT10] Ian Richardson and Murray Thomson. Domestic electricity demand model-simulation example. http://hdl.handle.net/2134/5786, 2010. Loughborough University Institutional Repository.

[RTI08] Ian Richardson, Murray Thomson, and David Infield. A high-resolution domestic building occupancy model for energy demand simulations. *Energy and buildings*, 40(8):1560–1566, 2008.

[RTIC10] Ian Richardson, Murray Thomson, David Infield, and Conor Clifford. Domestic electricity use: A high-resolution energy demand model. *Energy and buildings*, 42(10):1878–1887, 2010.

[RTID09] Ian Richardson, Murray Thomson, David Infield, and Alice Delahunty. Domestic lighting: A high-resolution energy demand model. *Energy and Buildings*, 41(7):781–789, 2009.

[RW99] Christelle Rigollier and Lucien Wald. Selecting a clear-sky model to accurately map solar radiation from satellite images. In *EARSeL Symposium 1999*, pages 131–137. Balkema, 1999.

[Sär13] Simo Särkkä. *Bayesian filtering and smoothing*, volume 3. Cambridge University Press, 2013.

[SBPW16] Lisa Scanu, Pierre Bernaud, Stéphane Ploix, and Etienne Wurtz. Méthodologie pour la comparaison de structures de modèles simplifiés. In *Conférence IBPSA France*. IBPSA, 2016.

[Sch10] Mark Schmidt. *Graphical model structure learning with L1-regularization*. PhD thesis, University of British Columbia, 2010.

[SFR07] Mark Schmidt, Glenn Fung, and Rmer Rosales. Fast optimization methods for L1 regularization: A comparative study and two new approaches. In *European Conference on Machine Learning (ECML)*, 2007.

[SGN05] Thomas Schön, Fredrik Gustafsson, and Per-Johan Nordlund. Marginalized particle filters for mixed linear/nonlinear state-space models. *IEEE Transactions on Signal Processing*, 53(7):2279–2289, 2005.

[ŠOCP11] Jan Široký, Frauke Oldewurtel, Jiří Cigler, and Samuel Prívara. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9):3079–3087, 2011.

[Söd07] Torsten Söderström. Errors-in-variables methods in system identification. *Automatica*, 43(6):939–958, 2007.

[SRL04] Melody Stokes, Mark Rylatt, and Kevin Lomas. A simple model of domestic lighting demand. *Energy and Buildings*, 36(2):103–116, 2004.

[SS82] Robert H. Shumway and David S. Stoffer. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of time series analysis*, 3(4):253–264, 1982.

[SS91] Robert H Shumway and David S Stoffer. Dynamic linear models with switching. *Journal of the American Statistical Association*, 86(415):763–769, 1991.

[SWN11] Thomas B. Schön, Adrian Wills, and Brett Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.

[TLH+17] Wai Lip Theo, Jeng Shiun Lim, Wai Shin Ho, Haslenda Hashim, and Chew Tin Lee. Review of distributed generation (DG) system planning and optimisation techniques: Comparison of numerical and mathematical modelling methods. *Renawable and Sustainable Energy Reviews*, 67:531–573, 2017.

[Tor14] Jacopo Torriti. A reivew of time use models of residential electricity demand. *Renewable and Sustainable Energy Reviews*, 37:265–272, 2014.

[TPH+] Khadija Tijani, Stephane Ploix, Benjamin Haas, Julie Dugdale, and Quoc Dung Ngo. Dynamic bayesian networks to simulate occupant behaviours in office buildings related to indoor air quality. In *IBPSA India 2015*.

[VHC93] Jay M Ver Hoef and Noel Cressie. Multivariable spatial prediction. *Mathematical Geology*, 25(2):219–240, 1993.

[VL78] Charles F. Van Loan. Computing integrals involving the matrix exponential. *IEEE Transactions on Automatic Control*, 23(3):395–404, 1978.

[VNB+16] Marika Vellei, Sukumar Natarajan, Benjamin Biri, Julian Padget, and Ian Walker. The effect of real-time context-aware feedback on occupants' heating behaviour and thermal adaptation. *Energy and Buildings*, 123:179–191, 2016.

[Vor14] Éric Vorger. *Étude de l'influence du comportement des habitants sur la performance énergétique du bâtiment*. PhD thesis, Mines ParisTech, 2014.

[WHHF16] Ouyang Wu, K Hariprasad, Biao Huang, and J Fraser Forbes. Identification of linear dynamic errors-in-variables systems with a dynamic uncertain input using the EM algorithm. In *IEEE 55th Conference on Decision and Control*, pages 1229–1234. IEEE, 2016.

[WHSR13] Urs Wilke, Frédéric Haldi, Jean-Louis Scartezzini, and Darren Robinson. A bottom-up stochastic model to predict building occupants' time-dependent activities. *Building and Environment*, 60:254–264, 2013.

[WLV+09] Joakim Widén, Magdalena Lundh, Iana Vassileva, Erik Dahlquist, Kajsa Ellegård, and Ewa Wäckelgård. Constructing load profiles for household electricity and hot water from time-use data—modelling approach and validation. *Energy and Buildings*, 41(7):753–768, 2009.

[WME12] Joakim Widén, Andreas Molin, and Kajsa Ellegård. Models of domestic occupancy, activities and energy use based on time-use data: deterministic and stochastic approaches with application to various building-related simulations. *Journal of Building Performance Simulation*, 5(1):27–44, 2012.

[WSN08] Adrian Wills, Thomas B. Schön, and Brett Ninness. Parameter estimation for discrete-time nonlinear systems using EM. In *Proceedings of the 17th IFAC World Congress, Seoul, Korea*, pages 1–6, 2008.

[WW10] Joakim Widén and Ewa Wäckelgård. A high-resolution stochastic model of domestic activity patterns and electricity demand. *Applied Energy*, 87(6):1880–1892, 2010.

[YGD+13] Dazhi Yang, Chaojun Gu, Zibo Dong, Panida Jirutitijaroen, Nan Chen, and Wilfred M Walsh. Solar irradiance forecasting using spatial-temporal covariance structures and time-forward kriging. *Renewable Energy*, 60:235–245, 2013.

[Zay11] Chadia Zayane. *Identification d'un modèle de comportement thermique de bâtiment à partir de sa courbe de charge*. PhD thesis, École Nationale Supérieure des Mines de Paris, 2011.

[ZGIR12] Ahmed Zoha, Alexander Gluhak, Muhammad Ali Imran, and Sutharshan Rajasegarar. Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey. *Sensors*, 12(12):16838–16866, 2012.

[ZM12] Hai-Xiang Zhao and Frédéric Magoulès. A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews*, 16(6):3586–3592, 2012.

Building identification within a connected object ecosystem

Tahar NABIL

RESUME : Cette thèse s'intéresse au problème de l'identification de modèle thermique d'un bâtiment intelligent, dont les objets connectés pallient la non-mesure des grandeurs physiques d'intérêt. Un premier algorithme traite de l'estimation boucle ouverte du système de bâtiment exploité en boucle fermée. Cet algorithme est ensuite modifié pour intégrer l'incertitude de mesure des données. Nous suggérons ainsi une méthode en boucle fermée, non-intrusive car s'affranchissant de la nécessité de mesurer la température intérieure. Puis, nous revenons à des approches en boucle ouverte. Les différents algorithmes permettent respectivement de réduire le biais contenu dans la mesure de température extérieure par une sonde connectée, de remplacer le coûteux capteur de flux solaire par un capteur de température extérieure, et enfin d'utiliser la courbe de charge totale, et non désagrégée, en tirant profit de signaux On/Off des objets connectés.

MOTS-CLEFS: Identification de bâtiment, objets connectés, systèmes boucle ouverte et boucle fermée, modèles d'état linéaire et non-linéaire, Espérance-Maximisation.

ABSTRACT: This thesis is devoted to the problem of the identification of a thermal model of a smart building, whose connected objects alleviate the lack of measurements of the physical quantities of interest. The first algorithm deals with the estimation of the open-loop building system, despite its actual exploitation in closed loop. This algorithm is then modified to account for the uncertainty of the data. We suggest a closed-loop estimation of the building system as soon as the indoor temperature is not measured. Then, we return to open-loop approaches. The different algorithms enable respectively to reduce the possible bias contained in a connected outdoor air temperature sensor, to replace the costly solar flux sensor by another connected temperature sensor, and finally to directly use the total load curve, without disaggregation, by making the most of the On/Off signals of the connected objects.

KEY-WORDS: Building identification, connected objects, open and closed loop systems, linear and non-linear state space models, Expectation-Maximization.

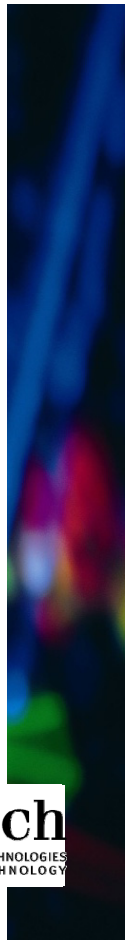

TELECOM ParisTech

EDiTE

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY