Recursive Parameter Estimation of Thermostatically Controlled Loads via Unscented Kalman Filter

Eric M. Burger, Scott J. Moura

Energy, Control, and Applications Lab, University of California, Berkeley

Abstract

For thermostatically controlled loads (TCLs) to perform demand response services in real-time markets, online methods for parameter estimation are needed. As the physical characteristics of a TCL change (e.g. the contents of a refrigerator or the occupancy of a conditioned room), it is necessary to update the parameters of the TCL model. Otherwise, the TCL will be incapable of accurately predicting its potential energy demand, thereby decreasing the reliability of a TCL aggregation to perform demand response. In this paper, we investigate the potential of various unscented Kalman filter (UKF) algorithm variations to recursively identify a TCL model that is non-linear in the parameters. Experimental results demonstrate the parameter estimation of two residential refrigerators. Finally, simulation results demonstrate the incorporation of the recursive parameter estimation methods into a model predictive controller for demand response.

Keywords: Smart grid, Unscented Kalman filter, Online system identification, Recursive parameter estimation, Thermostatically Controlled Loads (TCL), Demand response

1. Introduction

1.1. Background and Motivation

Large populations of thermostatically controlled loads (TCLs) hold great potential for performing ancillary services in power systems. The advantages of responsive TCLs over large storage technologies include: (i) they are already well-established technologies; (ii) they are spatially distributed around the power system; (iii) they employ simple and fast local actuation; (iv) they are unimpaired by the outage of individuals in the population; and (v) they - on the aggregate - can produce a quasi-continuous response despite the discrete nature of the individual controls [1][2][3].

*Corresponding author: Eric M. Burger,
Email: ericburger@berkeley.edu
**Affiliation Address: Energy, Control, and Applications Lab (eCAL), 611 Davis Hall, Department of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, CA 94720, USA.

Because TCLs are controlled according to a temperature setpoint and deadband range, customers are generally indifferent to precisely when electricity is consumed. The inherent flexibility of TCLs, such as refrigerators and electric water heaters, makes them promising candidates for provisioning power system services. In fact, direct load control (DLC) and demand response (DR) programs are increasingly controlling TCLs, among other electric loads, to improve power grid stability [4][5].

1.2. Relevant Literature

Past literature on the modeling and control of TCL populations has focused on the development of aggregation methods with centralized control. Malhame and Chong's study [6] is among the first reports to use stochastic analysis to develop an aggregate model of a TCL population. The coupled Fokker-Planck equations, derived in [6], define the aggregate behavior of a homogeneous population. More recently, [7] develops a diffusion-advection partial differential equation (PDE) model and parameter identification scheme for an aggregated population of heterogeneous TCLs. In [8], the authors present a deterministic hybrid PDE-based model for heterogeneous TCL populations, analyze its stability properties, and derive a power reference tracking control law.

In [9], the author uses a linearized Fokker-Planck model to describe the aggregated behavior of a TCL population. Direct control is achieved by broadcasting a single time-varying setpoint temperature offset signal to every agent. Numerical results demonstrate how small perturbations to the setpoint can enable TCLs to perform wind generation following. The work in [10] builds upon [9] by proposing a sliding mode control algorithm for direct control of air conditioning loads. A "state bin" modeling framework is used to describe local states (On/Off) in a discrete temperature-related manner.

In [11], the authors employ a linear time-invariant (LTI) representation of a TCL population. As in [10], a "state bin" modeling framework is used and the aggregate probability mass is allowed to move through these bins. A Markov Chain-based approach is used to predict the evolution of the TCL population. In [2], the authors propose a proportional controller which, at each time step, broadcasts a switching probability, η, to all the TCLs in the population. If $\eta < 0$, all TCLs that are on must switch off with a probability of η and if $\eta > 0$, TCLs that are off switch on with a probability of η.

Recognizing that system frequency is a universally available indicator of supply-demand imbalance, a number of researchers have developed fully decentralized techniques for performing frequency services with TCLs. In [12], the authors show the suitability of TCLs to perform frequency services using system frequency as a control signal and the potential for a population of TCLs to respond to a sudden loss of generation. This demand response capability reduces the dependence of grid operators on rapidly deployable backup generation.

In [13], the authors develop a TCL model in which devices adjust their setpoints linearly according to the system frequency, allowing the population to act as a fast frequency controlled reserve. To address problems

of long-term instability, Angeli and Kountouriotis develop a decentralized stochastic controller in [14] that is capable of maintaining desynchronization among the TCLs while regulating overall power consumption. In [15], the authors present a stochastic controller whereby each TCL in the population independently targets a reference power profile. The result is a stable and fully decentralized system that requires only the locally available control signals of frequency and time.

Recent trends in the field of convex optimization, in particular the introduction of the alternating direction method of multipliers (ADMM), have enabled researchers to pursue distributed methods of load control [16][17]. With these distributed control schemes, individual TCLs coordinate amongst each other or with an aggregator to drive the population towards a shared global objective. In [3], the authors employ a sharing ADMM algorithm to coordinate the electricity demand of a TCL population to perform generation following. By allowing each TCL to locally model its dynamics and enforce constraints, the authors demonstrate the capability of their algorithm to accurately control a heterogeneous TCL population.

A consensus coordination algorithm is developed in [18] that enables the aggregated power consumption of a population of air conditioners to follow a power generation forecast. With each TCL modeling its dynamics locally, the fully distributed algorithm is able to achieve fast convergence with only limited communication between neighboring TCLs in the network. In [19], the authors present a distributed model predictive control (MPC) scheme that enables a population of TCLs to provide load balancing services. The results demonstrate that as the number of TCLs in the population increases, the distributed MPC scheme achieves a higher efficacy than an aggregation method with centralized control.

Unlike centralized control schemes with aggregation models, decentralized and distributed control schemes rely on the individual TCLs to locally optimize their behavior. Therefore, it is necessary for every agent in the population to model its own behavior and to predict its energy demand. TCLs with poorly fit models will undermine the ability of the population to accurately perform ancillary services. Given that most TCLs experience regular changes to their physical characteristics (e.g. the contents of a refrigerator, the flow through a water heater, or the occupancy of a conditioned room), a linear time-invariant model is likely to prove inadequate. Also, for TCLs like radiant heaters and air conditioners, it is not possible for the manufacturer to predetermine the physical characteristics of the spaces that will be conditioned. Therefore, to improve the performance of decentralized and distributed TCL control methods, it is advantageous to employ recursive or online parameter estimation algorithms to fit and continuously update each TCL's model.

1.3. Main Contributions

This manuscript contributes to the development of recursive parameter estimation algorithms for TCLs by investigating various unscented Kalman filters for the estimation of a TCL model that is non-linear in the parameters. We present four closely related filter methods (single, joint, dual, and triple) employing

both the standard Kalman filter (KF), and unscented Kalman filter (UKF) algorithms. Specifically, we consider: (i) a single filter approach in which one UKF estimates the TCL parameters; (ii) a joint filter approach in which one UKF simultaneously estimates both the parameters and the state; (iii) a dual filter approach in which one UKF estimates the parameters and one KF estimates the state; and (iv) a triple filter approach in which one UKF estimates the parameters, one KF estimates the state, and another KF estimates the model inputs. We present experimental parameter estimation results using real temperature data from two residential refrigerators. Finally, simulation studies demonstrate the incorporation of the single filter approach into a model predictive controller for demand response.

1.4. Paper Outline

This paper is organized as follows. Section 2 discusses the TCL model and Section 3 overviews the parameter estimation problem. Sections 4 and 5 provide background for the standard Kalman filter (KF) and the unscented Kalman filter (UKF), respectively. Section 6 formulates four filter methods for recursive parameter estimation of a TCL. Section 7 provides numerical examples of our proposed algorithms applied to real temperature data from two residential refrigerators. Section 8 presents simulation studies showing fast parameter convergence and the application of the single filter approach to the model predictive control of a TCL population for demand response. Finally, Section 9 summarizes key results.

2. TCL Model

The literature is rich with models for representing the dynamics of a diversity of thermostatically controlled loads. These models address different challenges with respect to fidelity, user or occupant behavior, unobserved inputs or states, and suitability for controls applications. For example, in [12], the authors employ a refrigerator model which represents the energy flows between 4 thermally coupled masses (fridge structure, fridge contents, freezer structure, and freezer contents), allowing the researchers to accurately estimate the thermal storage capacity of the system.

Stochastic grey-box techniques are employed in [20] and [21]. In [20], the authors estimate the thermal mass, evaporator thermal resistance, insulation thermal resistance, and coefficient of performance (COP) for a residential refrigerator using maximum likelihood estimation. An electrical power consumption to temperature model for a residential freezer is developed in [21], allowing authors to apply model predictive control to shift the demand.

In [22], the authors develop a control-oriented multi-state thermal model and parameter estimation method for heating, ventilation, and air-conditioning (HVAC) systems in commercial buildings. A piecewise linear model of a residential heating system is presented in [23] which is able to approximate higher order thermal dynamics within the system.

This paper is motivated by ongoing research into decentralized and distributed control schemes for TCL populations. Accordingly, our objective is to recursively estimate a control-oriented model that can capture the predominant dynamics and disturbance patterns within a TCL. Low order RC equivalent circuit TCL models, which have structures that are amenable to optimization and to real-time data-driven parameter estimation, are widely employed for control application.

In this paper, we use the hybrid state discrete time TCL model originally developed in [24]. The model, which is a discretization of a continuous time RC model, is widely used in the TCL population control literature, including [2][3][9][10][11][14][18][19].

$$T^{k+1} = \theta_1 T^k + (1-\theta_1)(T_\infty^k + \theta_2 m^k) + \theta_3 \tag{1a}$$

$$m^{k+1} = \begin{cases} 1 & \text{if } T^k > T_{set} + \frac{\delta}{2} \\ 0 & \text{if } T^k < T_{set} - \frac{\delta}{2} \\ m^k & \text{otherwise} \end{cases} \tag{1b}$$

where state variables $T^k \in \mathbf{R}$ and $m^k \in \{0, 1\}$ denote the temperature of the conditioned mass and the discrete state (on or off) of the mechanical system, respectively. Additionally, $k \in \mathbf{Z}$ denotes the integer-valued time step, $T_\infty^k \in \mathbf{R}$ the ambient temperature (°C), $T_{set} \in \mathbf{R}$ the temperature setpoint (°C), and $\delta \in \mathbf{R}$ the temperature deadband width (°C).

In this paper, we will define the time elapsed between each time step as $h = 1/60$ (hours). The parameter θ_1 represents the thermal characteristics of the conditioned mass as defined by $\theta_1 = \exp(-h/RC)$ where C is the thermal capacitance (kWh/°C) and R is the thermal resistance (°C/kW), θ_2 the energy transfer to or from the mass due to the systems operation as defined by $\theta_2 = RP$ where P is the rate of energy transfer (kW), and θ_3 an additive noise process accounting for energy gain or loss not directly modeled. The sign conventions in (1) assume that the TCL is providing a cooling load and that P (and thus θ_2) is negative.

As noted in [9][24], the discrete time model implicitly assumes that all changes in mechanical state occur on the time steps of the simulation. In this paper, we will assume that this behavior reflects the programming of the systems being modeled. In other words, we will assume that the TCLs have a thermostat sampling frequency of $1/h$ Hz or once per minute.

3. Parameter Estimation Background

A fundamental machine learning problem involves the identification of a nonlinear mapping

$$y^k = G(x^k, \theta) \tag{2}$$

where variable $x^k \in \mathbf{R}^X$ is the input, $y^k \in \mathbf{R}^Y$ is the output, and the nonlinear map G is parameterized by $\theta \in \mathbf{R}^\Theta$. Additionally, k denotes the integer-valued time step and X, Y, and Θ are the number of inputs, outputs, and parameters, respectively.

3.1. Batch Parameter Estimation

Learning can be performed in a batch manner by producing estimates of the parameters $\hat{\theta}$ given a training set of observed inputs and desired outputs, $\{x, y\}$. The goal of a parameter estimation algorithm is to minimize some function of the error between the desired and estimated outputs as given by $e^k = y^k - G(x^k, \hat{\theta})$.

3.2. Recursive Parameter Estimation

The parameter estimation problem can be expressed in a recursive form using a discrete-time state-space model representation

$$\theta^k = \theta^{k-1} + n^k \tag{3a}$$

$$y^k = G(x^k, \theta^k) + e^k \tag{3b}$$

where θ^k represents the parameter estimates at time step k and $n^k \in \mathbf{R}^\Theta$ corresponds to the parameter update noise (i.e. change in parameter values). The goal of a recursive parameter estimation algorithm is to produce $\hat{\theta}^k$ so as to minimize some function of the error e^k.

4. Kalman Filter Background

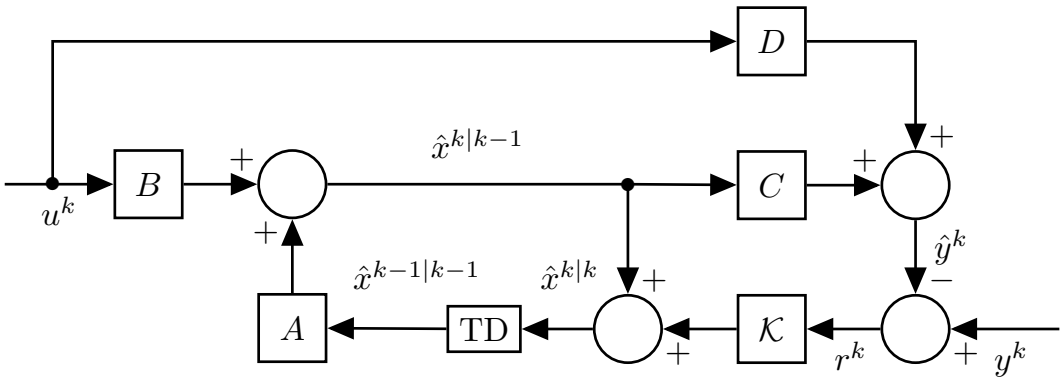

Figure 1: Kalman Filter Diagram

The Kalman filter (KF) is a recursive estimator for linear models such as the discrete-time state-space model

$$x^k = Ax^{k-1} + Bu^k + v^k \tag{4a}$$

$$y^k = Cx^k + Du^k + w^k \tag{4b}$$

where variable $x^k \in \mathbf{R}^X$ is the state of the system, $u^k \in \mathbf{R}^U$ is the known exogenous input, and $y^k \in \mathbf{R}^Y$ is the observed measurement signal. The state transition model is given by $A \in \mathbf{R}^{X \times X}$ and the control-input model by $B \in \mathbf{R}^{X \times U}$. The process noise $v^k \in \mathbf{R}^X$ has covariance $Q_v \in \mathbf{R}^{X \times X}$, $v^k \sim N(0, Q_v)$. The observation model is given by $C \in \mathbf{R}^{Y \times X}$ and the feedthrough model by $D \in \mathbf{R}^{Y \times U}$. The measurement noise $w^k \in \mathbf{R}^Y$ has covariance $Q_w \in \mathbf{R}^{Y \times Y}$, $w^k \sim N(0, Q_w)$. The variances of v^k and w^k (i.e. diagonal elements of Q_v and Q_w, respectively) must be known in order to implement a Kalman filter.

The Kalman filter (KF) algorithm consists of a prediction step and an update/correction step. The KF will model x^k as a Gaussian random variable (GRV) with estimated mean $\hat{x}^k \in \mathbf{R}^X$ and covariance $Q_x^k \in \mathbf{R}^{X \times X}$. To provide clarity, it is helpful to expand the k notation to distinguish between the state estimates produced before and after the KF correction step. Therefore, at each time step k, the predicted (a priori) state estimate, denoted as $\hat{x}^{k|k-1}$, is the mean estimate of x^k given measurements $y^0, \ldots, y^{k-1}$. The corrected (a posterior) state estimate, $\hat{x}^{k|k}$, is the mean estimate of x^k given measurements $y^0, \ldots, y^k$. To reiterate, throughout this paper, the uncorrected predictions (a priori) are denoted by $k|k-1$ or $k+1|k$ whereas the corrected predictions (a posterior) are denoted by $k|k$, $k-1|k-1$, or $k+1|k+1$.

The KF prediction step is given by

$$\hat{x}^{k|k-1} = A\hat{x}^{k-1|k-1} + Bu^k \tag{5a}$$

$$Q_x^{k|k-1} = AQ_x^{k-1|k-1}A^T + Q_v \tag{5b}$$

and the update/correction step by

$$\hat{y}^k = C\hat{x}^{k|k-1} + Du^k \tag{6a}$$

$$Q_y = CQ_x^{k|k-1}C^T + Q_w \tag{6b}$$

$$\mathcal{K} = Q_x^{k|k-1}C^T Q_y^{-1} \tag{7a}$$

$$r^k = y^k - \hat{y}^k \tag{7b}$$

$$\hat{x}^{k|k} = \hat{x}^{k|k-1} + \mathcal{K}r^k \tag{7c}$$

$$Q_x^{k|k} = Q_x^{k|k-1} - \mathcal{K}Q_y\mathcal{K}^T \tag{7d}$$

Figure 1 illustrates the KF algorithm. The block TD represents a time delay (commonly denoted in controls literature by z^{-1} or $1/z$, the Z-transform of the delay operator). To simplify notation in this paper, we will express the Kalman filter algorithm with the following 3 operator expressions

$$\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \end{bmatrix} = KF_x\left(\begin{bmatrix} A \\ B \end{bmatrix}, \begin{bmatrix} \hat{x}^{k-1|k-1} \\ Q_x^{k-1|k-1} \end{bmatrix}, u^k, Q_v\right) \tag{8a}$$

$$\begin{bmatrix} \hat{y}^k \\ Q_y^k \end{bmatrix} = KF_y\left(\begin{bmatrix} C \\ D \end{bmatrix}, \begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \end{bmatrix}, u^k, Q_w\right) \tag{8b}$$

$$\begin{bmatrix} \hat{x}^{k|k} \\ Q_x^{k|k} \\ r^k \end{bmatrix} = KF_c\left(\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \end{bmatrix}, \begin{bmatrix} \hat{y}^k \\ Q_y^k \end{bmatrix}, C, y^k\right) \tag{8c}$$

where (8a) corresponds to (5a-5b), (8b) to (6a-6b), and (8c) to (7a-7d).

5. Unscented Kalman Filter Background

The UKF is an extension to the standard Kalman filter that utilizes a deterministic sampling approach known as the *unscented transform* (UT) to characterize states which undergo a *nonlinear* transformation. The UKF builds on the intuition that it is easier to approximate a probability distribution than to approximate an arbitrary nonlinear transformation [25].

Like the Kalman filter, the UKF includes a prediction step and an update/correction step. However, with the UKF, a state distribution is approximated by a Gaussian random variable (GRV) and specified using a minimal set of sample, or sigma, points around the mean. These sigma points are selected such that they capture the true mean and covariance of the GRV. When propagated through a nonlinear transform, the sigma points accurately capture the a posterior mean and covariance of the estimated state.

In other words, rather than simply passing the previous state estimate $\hat{x}^{k-1}$ through a nonlinear transform to produce a predicted state estimate $\hat{x}^k$, the UKF transforms the set of sigma points. The predicted state estimate $\hat{x}^k$ is then recovered as a weighted mean of the transformed points. With the UT, approximations of Gaussian states are accurate to the third order for any nonlinearities [26]. For non-Gaussian states, approximations are accurate to at least the second-order for any nonlinearities.

In this paper, we employ the UKF algorithm as presented by Wan and van der Merwe [26][27][28] and summarized in the following section. Specifically, see Tables 7.3.1 and 7.3.2 in [27] for the algorithm employed in this work. We direct the reader to [25] for the original presentation of the UT and UKF. A discussion of dual estimation can be found in [26][28]. In the following subsections, we summarize the UKF algorithm as presented by Wan and van der Merwe [27].

5.1. Sigma Points and Unscented Transform

To detail the UKF algorithm, we begin by describing the generation of sigma points and the execution of the unscented transform (UT). Consider a random variable $s \in \mathbf{R}^L$ with mean $\bar{s} \in \mathbf{R}^L$ and covariance $Q_s \in \mathbf{R}^{L\times L}$ that is propagated through a nonlinear function f such that $z = f(s)$ where $z \in \mathbf{R}^Z$. To

calculate the statistics of z, we form a matrix $\mathcal{S} \in \mathbf{R}^{L\times(2L+1)}$ consisting of $2L+1$ sigma points $\mathcal{S}_i$ given by

$$\begin{aligned}\mathcal{S}_0 &= \bar{s} \\ \mathcal{S}_i &= \bar{s} + \left(\sqrt{(L+\lambda)Q_s}\right)_i, \quad i = 1,\ldots,L \\ \mathcal{S}_{i+L} &= \bar{s} - \left(\sqrt{(L+\lambda)Q_s}\right)_i, \quad i = 1,\ldots,L\end{aligned} \tag{9}$$

where $\left(\sqrt{(L+\lambda)Q_s}\right)_i$ is the ith column of the matrix square root of $(L+\lambda)Q_s$ and $\lambda = \alpha^2(L+\kappa) - L$ is a scaling parameter. Constant α determines the spread of the sigma points (usually $10^{-4} \leq \alpha \leq 1$) and constant κ is a secondary scaling parameter (usually $\kappa = 0$ or $3 - L$). We refer the reader to section 7.3 of the Wan and van der Merwe text [27] for a discussion about selecting the α and κ parameters.

The sigma points are propagated through the nonlinear function

$$\mathcal{Z}_i = f(\mathcal{S}_i) \quad i = 0,\ldots,2L \tag{10}$$

and the mean and covariance of z are approximated as a weighted mean and covariance of the a posterior sigma points

$$\bar{z} \approx \sum_{i=0}^{2L} \mathcal{W}_{m,i}\mathcal{Z}_i \tag{11}$$

$$Q_z \approx \sum_{i=0}^{2L} \mathcal{W}_{c,i}(\mathcal{Z}_i - \bar{z})(\mathcal{Z}_i - \bar{z})^T \tag{12}$$

with weights $\mathcal{W}_m$, corresponding to the a posterior mean of the sigma points, given by

$$\begin{aligned}\mathcal{W}_{m,0} &= \lambda/(L+\lambda) \\ \mathcal{W}_{m,i} &= \lambda/(2(L+\lambda)), \quad i = 1,\ldots,2L\end{aligned} \tag{13}$$

and weights $\mathcal{W}_c$, corresponding to the a posterior covariance of the sigma points, given by

$$\begin{aligned}\mathcal{W}_{c,0} &= \lambda/(L+\lambda) + (1 - \alpha + \beta) \\ \mathcal{W}_{c,i} &= \mathcal{W}_{m,i}, \quad i = 1,\ldots,2L\end{aligned} \tag{14}$$

where constant β incorporates prior knowledge of the distribution of s (for Gaussian distributions, $\beta = 2$ is optimal).

To simplify notation, we will denote the generation of sigma points and the execution of the unscented transform ((9)-(12)) with the following operator expressions

$$\mathcal{S} = UT_s(\bar{s}, Q_s) \tag{15a}$$

$$\mathcal{Z}_i = f(\mathcal{S}_i) \quad i = 0,\ldots,2L \tag{15b}$$

$$\bar{z} = UT_m(\mathcal{Z}) \tag{15c}$$

$$Q_z = UT_c(\mathcal{Z}) \tag{15d}$$

where (15a) corresponds to (9), (15c) to (11), and (15d) to (12).

5.2. Additive Unscented Kalman Filter Algorithm

The unscented Kalman filter (UKF) is a straightforward application of the UT to recursive estimation. To present the UKF algorithm, we will consider the state estimation of a discrete-time nonlinear dynamic system with additive noise given by the state-space model

$$x^k = F(x^{k-1}, u^k) + v^k \tag{16a}$$

$$y^k = H(x^k, u^k) + w^k \tag{16b}$$

where variable $x^k \in \mathbf{R}^X$ is the state of the system, $u^k \in \mathbf{R}^U$ is the known exogenous input, and $y^k \in \mathbf{R}^Y$ is the observed measurement signal. Function F is the transition model and the process noise $v^k \in \mathbf{R}^X$ has covariance $Q_v \in \mathbf{R}^{X\times X}$, $v^k \sim N(0, Q_v)$. Function H is the observation model and the measurement noise $w^k \in \mathbf{R}^Y$ has covariance $Q_w \in \mathbf{R}^{Y\times Y}$, $w^k \sim N(0, Q_w)$. For systems with non-additive noise processes, we refer the reader to the Appendix for a presentation of the augmented unscented Kalman filter.

The UKF will model x^k as a GRV with estimated mean $\hat{x}^k$ and covariance Q_x^k. At each time step k, the UKF will generate sigma points for the previous state estimate, $\hat{x}^{k-1|k-1}$, as given by

$$\mathcal{S}_x^{k-1} = UT_s(\hat{x}^{k-1|k-1}, Q_x^{k-1|k-1}) \tag{17}$$

where $\mathcal{S}_x$ is the sigma points associated with the state estimate.

In the prediction step, the UKF will propagate the sigma points through the process model and generate the a priori sigma points $\mathcal{X}_i^{k|k-1}$, state estimate $\hat{x}^{k|k-1}$, and covariance $Q_x^{k|k-1}$ as follows,

$$\mathcal{X}_i^{k|k-1} = F(\mathcal{S}_{x,i}^{k-1}, u^k) \quad i = 0, \ldots, 2X \tag{18a}$$

$$\hat{x}^{k|k-1} = UT_m(\mathcal{X}^{k|k-1}) \tag{18b}$$

$$Q_x^{k|k-1} = UT_c(\mathcal{X}^{k|k-1}) + Q_v \tag{18c}$$

$$\tag{18d}$$

In the correction step, the UKF will propagate the a priori sigma points through the measurement model to generate the measurement sigma points $\mathcal{Y}_i^k$, estimate $\hat{y}^k$, and covariance Q_y as follows,

$$\mathcal{Y}_i^k = H(\mathcal{X}_i^{k|k-1}, u^k) \quad i = 0, \ldots, 2X \tag{19a}$$

$$\hat{y}^k = UT_m(\mathcal{Y}^k) \tag{19b}$$

$$Q_y = UT_c(\mathcal{Y}^k) + Q_w \tag{19c}$$

$$\tag{19d}$$

These are used to calculate the cross-covariance Q_{xy}, the Kalman gain $\mathcal{K}$, and the observation error r^k. Finally, the state estimate and covariance are corrected, producing the a posterior estimate $\hat{x}^{k|k}$ and covariance $Q_x^{k|k}$.

$$Q_{xy} = \sum_{i=0}^{2L} \mathcal{W}_{c,i}(\mathcal{X}_i^{k|k-1} - \hat{x}^{k|k-1})(\mathcal{Y}_i^k - \hat{y}^k)^T \tag{20a}$$

$$\mathcal{K} = Q_{xy}Q_y^{-1} \tag{20b}$$

$$r^k = y^k - \hat{y}^k \tag{20c}$$

$$\hat{x}^{k|k} = \hat{x}^{k|k-1} + \mathcal{K}r^k \tag{20d}$$

$$Q_x^{k|k} = Q_x^{k|k-1} - \mathcal{K}Q_y\mathcal{K}^T \tag{20e}$$

where $r^k \in \mathbf{R}$ is the error between the measurement y^k and the estimate $\hat{y}^k$ at time step k.

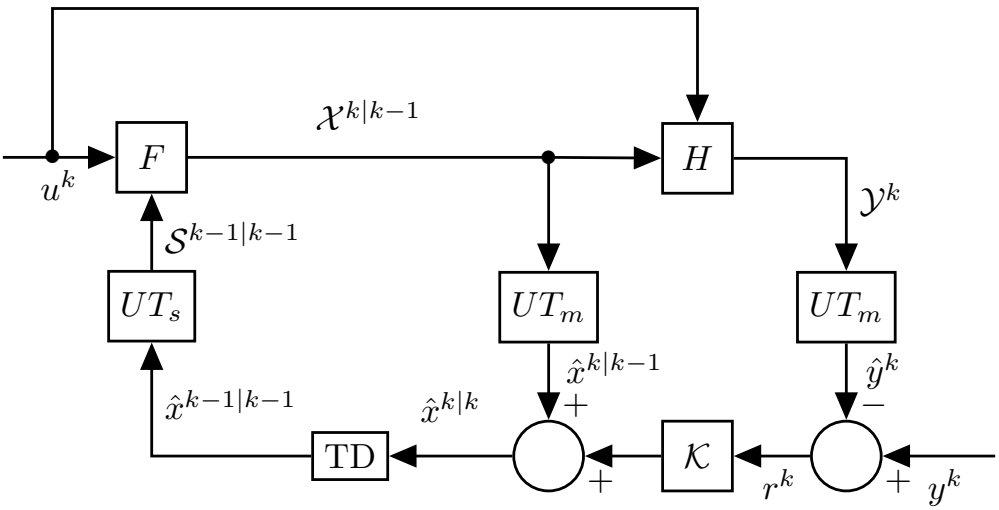

Figure 2: Additive Unscented Kalman Filter Diagram

Figure 2 illustrates the additive UKF algorithm. The block TD represents a time delay (commonly denoted in controls literature by z^{-1} or $1/z$, the Z-transform of the delay operator). To simplify notation in this paper, we will express the additive unscented Kalman filter algorithm using the following 4 operator expressions

$$\left[\mathcal{S}_x^{k-1}\right] = UKF_s^+\left(\begin{bmatrix}\hat{x}^{k-1|k-1} \\ Q_x^{k-1|k-1}\end{bmatrix}\right) \tag{21a}$$

$$\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \\ \mathcal{X}^{k|k-1} \end{bmatrix} = UKF_x^+\left(F, \mathcal{S}_x^{k-1}, u^k, Q_v\right) \tag{21b}$$

$$\begin{bmatrix} \hat{y}^k \\ Q_y^k \\ \mathcal{Y}^k \end{bmatrix} = UKF_y^+(H, \mathcal{X}^{k|k-1}, u^k, Q_w) \tag{21c}$$

$$\begin{bmatrix} \hat{x}^{k|k} \\ Q_x^{k|k} \\ r^k \end{bmatrix} = UKF_c^+\left(\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \\ \mathcal{X}^{k|k-1} \end{bmatrix}, \begin{bmatrix} \hat{y}^k \\ Q_y^k \\ \mathcal{Y}^k \end{bmatrix}, y^k\right) \tag{21d}$$

where (21a) corresponds to (17), (21b) to (18a-18c), (21c) to (19a-19c), and (21d) to (20a-20e).

6. Recursive TCL Parameter Estimation

In this section, we will present 4 closely related approaches for parameter estimation of a thermostatically controlled load (TCL) using the Kalman filter (KF) algorithm in (8) and unscented Kalman filter (UKF) algorithm in (21). In this paper, we will consider: (i) a single filter approach in which one UKF is used to estimate the parameters θ^k; (ii) a joint filter approach in which one UKF simultaneously estimates both θ^k and T^k; (iii) a dual filter approach in which one UKF estimates the parameters θ^k and one KF estimates the state T^k; and (iv) a triple filter approach in which we use one UKF to estimate θ^k, one KF to estimate T^k, and another KF to estimate the inputs, T_∞^k and m^k.

In each case, we define the function G according to the TCL model (1)

$$\begin{aligned} T^{k+1} &= \theta_1^k T^k + \begin{bmatrix} 1-\theta_1^k & 1-\theta_1^k\theta_2^k & \theta_3^k \end{bmatrix} \begin{bmatrix} T_\infty^k \\ m^k \\ 1 \end{bmatrix} \\ &= G(T^k, T_\infty^k, m^k, \theta^k) \end{aligned} \tag{22}$$

6.1. Single Filter Parameter Estimation

Using the function G given in (22), the TCL recursive parameter estimation problem can be expressed with the state-space model

$$\theta^k = \theta^{k-1} + n^k \tag{23a}$$

$$y^k = G(T^k, T_\infty^k, m^k, \theta^k) + v^k + w^k \tag{23b}$$

where (23b) combines (22) with observation model $y^k = T^{k+1} + w^k$.

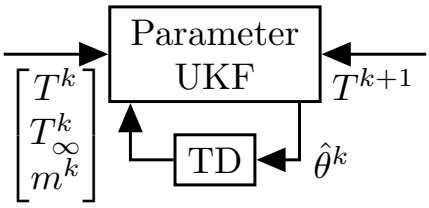

Figure 3: Single Filter Method Diagram

Figure 3 illustrates the single filter method. The block TD represents a time delay (commonly denoted $1/z$, the Z-transform of the delay operator). By employing the additive UKF algorithm in (21) with the T^{k+1} observation as y^k, we can produce $\hat{\theta}^k$, an estimate of the model parameters at time step k. Note that in the single filter case, θ corresponds to x, the variable being estimated, and G to H, the observation model. Additionally, the transition model F is given by $F(x^{k-1}, u^k) = x^{k-1}$ and T^k, T^k_∞, and m^k are effectively u^k, control and feed-through inputs at time step k.

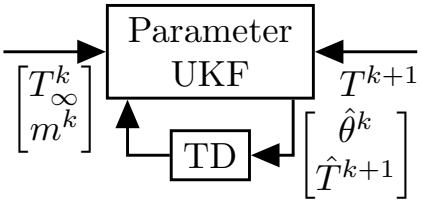

Figure 4: Joint Filter Method Diagram

6.2. Joint Filter State and Parameter Estimation

For system identification, it is often necessary to simultaneous perform state and parameter estimation from noisy observations [26]. There are two basic approaches, joint and dual estimation. In the joint estimation method, state and parameter estimation can be performed simultaneously with a single filter by estimating the state-space model

$$\begin{bmatrix} \theta^k \\ T^{k+1} \end{bmatrix} = \begin{bmatrix} \theta^{k-1} \\ G(T^k, T^k_\infty, m^k, \theta^{k-1}) \end{bmatrix} + \begin{bmatrix} n^k \\ v^k \end{bmatrix} \tag{24a}$$

$$y^k = T^{k+1} + w^k \tag{24b}$$

Figure 4 illustrates the joint filter method where the block TD represents a time delay. Just as in the single filter method, we employ the additive UKF algorithm in (21) with the T^{k+1} observation as y^k to recursively estimate the model. However, in the joint filter method, we produce estimates of both the state and the parameters ($\hat{T}^{k+1}$ and $\hat{\theta}^k$, respectively).

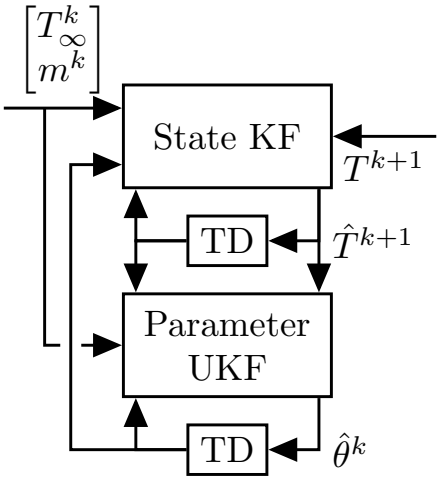

Figure 5: Dual Filter Method Diagram

6.3. Dual Filter State and Parameter Estimation

In the dual estimation method, a separate state-space representation is used for the states and parameters. For a TCL, the state model is given by

$$T^{k+1} = G(T^k, T_\infty^k, m^k, \theta^k) + v^k \tag{25a}$$

$$y^k = T^{k+1} + w^k \tag{25b}$$

and the parameter model by (23).

Figure 5 illustrates the dual filter method where the block TD represents a time delay. Because the function G is linear in the states, we can estimate the state model (25) using the KF algorithm in (8) with the T^{k+1} observation as y^k to produce $\hat{T}^{k+1}$. Again, the parameter model (23) is estimated using the additive UKF algorithm in (21). We tie the two filters together by using the estimated state of one filter as the control input and/or observation in another filter. Specifically, for the state filter, we use the previous parameter estimate $\hat{\theta}^{k-1}$ in the transition model. For the parameter filter, we use the previous estimate $\hat{T}^k$ as input in the observation model and the current estimate $\hat{T}^{k+1}$ as the observation y^k (rather than the T^k and T^{k+1} observations, respectively).

6.4. Triple Filter Input, State, and Parameter Estimation

Lastly, we consider an estimation approach in which separate filters are used to estimate the inputs, states, and parameters. For simplicity, we will refer to this as the triple filter approach. The input model is given by

$$\begin{bmatrix} T_\infty^k \\ m^k \end{bmatrix} = \begin{bmatrix} T_\infty^{k-1} \\ m^{k-1} \end{bmatrix} + \begin{bmatrix} p_1^k \\ p_2^k \end{bmatrix} \tag{26a}$$

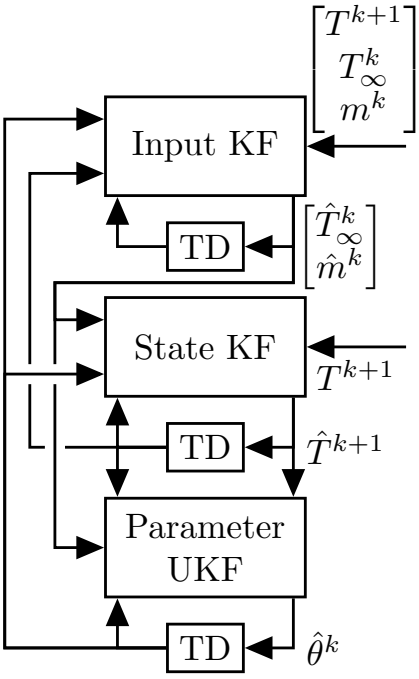

Figure 6: Triple Filter Method Diagram

$$y^k = \begin{bmatrix} G(T^k, T^k_\infty, m^k, \theta^{k-1}) \\ T^k_\infty \\ m^k \end{bmatrix} + \begin{bmatrix} v^k + w^k \\ q^k_1 \\ q^k_2 \end{bmatrix} \tag{26b}$$

where $p \in \mathbf{R}^2$ and $q \in \mathbf{R}^2$ are process and measurement noises, respectively, associated with the inputs T_∞ and m. Again, the state model is given by (25) and the parameter model by (23).

Figure 6 illustrates the triple filter method where the block TD represents a time delay. The input model (26) is estimated using the KF algorithm in (8) with the T^{k+1}, T^k_∞, and m^k observations as y^k to produce $\hat{T}^k_\infty$ and $\hat{m}^k$. To tie the three models together, the input estimates $\hat{T}^k_\infty$ and $\hat{m}^k$ are used in the transition model of the state filter and the observation model of the parameter filter. The previous parameter estimate $\hat{\theta}^{k-1}$ is used in the observation model of the input filter and the transition model of the state filter. Lastly, the state estimate $\hat{T}^k$ is used in the observation model of the input and parameter filters and $\hat{T}^{k+1}$ serves as the observation y^k in the parameter filter.

7. TCL Estimation Experimental Studies

In this section, we present parameter estimation results for TCL_1, a 500W residential refrigerator, and TCL_2, a 100W mini-fridge. Each TCL is instrumented with two DS18B20 digital temperature sensors to measure the ambient temperature T_∞ and internal refrigerator temperature T. The sensors have a $-55°$C to $+125°$C temperature range and a $\pm 0.5°$C accuracy from $-10°$C to $+85°$C. A current sensor is used to

	TCL_1			TCL_2		
	$\hat{\theta}_1^f$	$\hat{\theta}_2^f$	$\hat{\theta}_3^f$	$\hat{\theta}_1^f$	$\hat{\theta}_2^f$	$\hat{\theta}_3^f$
Single	0.998	-50.329	0.004	0.987	-36.193	-0.141
Joint	0.997	-49.762	0.005	0.985	-37.277	-0.159
Dual	0.998	-50.872	0.004	0.988	-36.822	-0.133
Triple	0.997	-52.268	0.005	0.986	-37.141	-0.153

Table 1: Final Parameter Estimates for TCL_1 and TCL_2

measure the state (on or off) of the compressor, m. Measurements were taken at 1 minute intervals for a period of 7 days. In this study, the temperature of the freezer is neither measured nor modeled.

TCL_1 is observed under typical operating conditions for a residential refrigerator. Therefore, the door is opened randomly and the contents of the refrigerator change regularly. By contrast, TCL_2 is empty except for 18 liters of water, which compose the thermal mass being conditioned by the unit. The door of TCL_2 remains closed for the duration of the study.

For both TCLs, we have implemented four parameter estimation methods using the standard and unscented Kalman filters: single filter, joint filter, dual filter, and triple filter. The final parameter estimates $\hat{\theta}^f$ are presented in Table 1. Normally, for system identification, we would seek to measure the performance of each algorithm by first learning the parameters using a training dataset and then testing the parameters using a separate validation dataset. However, an advantage of recursive parameter estimation is that we can continuously improve the parameter estimates and potentially adapt to changes in the mechanical system. Therefore, for the single, dual, joint, and triple filter methods, we represent the performance as the root mean squared error (RMSE) over the last 300 time steps (i.e 5 hours). We will denote this moving window RMSE as RMSE300. To measure the parameter estimation error, we employ the residual error r^k of the parameter filter for each of the four methods.

Figure 7 presents the parameter estimation results using the single filter method. The top subplots show the parameter estimates $\hat{\theta}^k$ produced by the parameter filter at each time step k. For each parameter, the center line is the mean or expected value of the estimate and the top and bottom lines illustrate the variance relative to the mean. Eventually, all parameter estimates converge and the variances decrease. The bottom subplots depict the residual error r^k. As shown, the parameter estimation for TCL_1 converges in about 1500 time steps while TCL_2 converges in about 3000 time steps.

The parameter estimation convergence rate of the UKF can be adjusted by changing the initial parameter estimates ($\hat{\theta}_1^0$, $\hat{\theta}_2^0$, and $\hat{\theta}_3^0$) and by tuning the process and measurement noise covariance matrices (Q_v and Q_w). We have chosen initial parameter estimates and covariance matrices that result in relatively slow

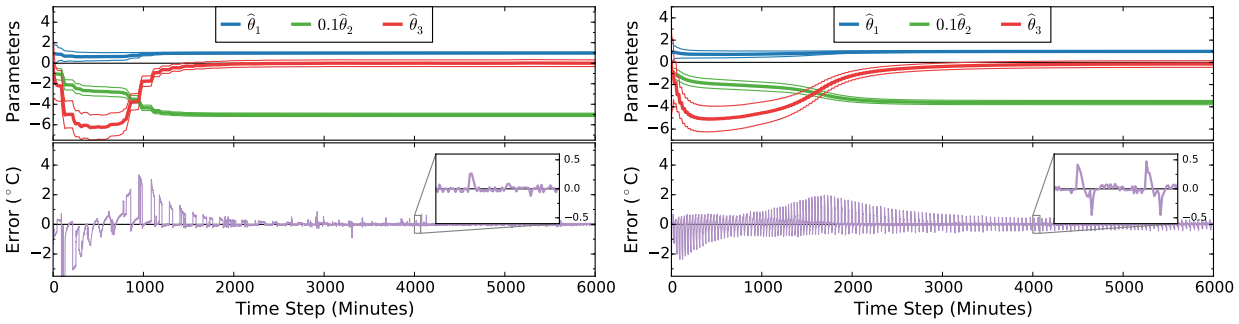

(a) TCL_1 (b) TCL_2

Figure 7: Single Filter Parameter Estimation

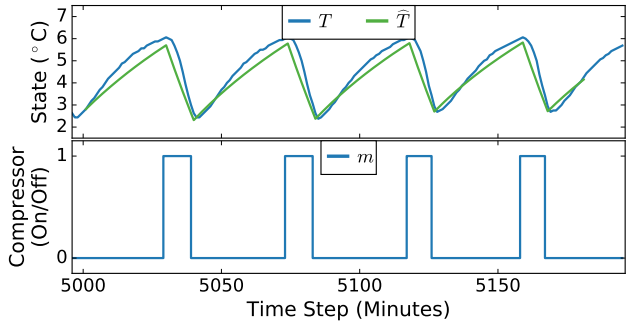

Figure 8: TCL_2 Temperature Forecast

convergence in order to produce results that illustrate the suitability of the UKF to perform parameter estimation of a TCL.

For use in controls applications, it is important for a model to be able to accurately predict the states of a system. Figure 8 presents a 3 hour temperature forecast $\hat{T}$ for TCL_2 given the parameter estimates produced by the single filter method at time step $k = 5000$. To generate the forecast, we assume perfect knowledge of the inputs (m and T_∞). The root mean squared error (RMSE) of the 3 hour temperature forecast $\hat{T}$ with respect to the measured temperature T is 0.47°C.

Figure 9 presents the parameter estimation results using the joint filter method and Fig. 10 for the dual filter method. Again, the top subplots show the parameter estimates $\hat{\theta}^k$. The bottom subplots depict the residual error r^k. The state estimates $\hat{T}^k$ are presented in the center subplots. Both TCLs exhibit similar convergence characteristics compared to the single filter method. Upon convergence, the difference between the measured temperature T and estimated temperature $\hat{T}$ becomes negligible. A sample of the temperature estimates produced by the dual filter are shown in Fig. 11.

As shown in Table 1, the differences in the final TCL_1 and TCL_2 parameter estimates for the single, joint, and dual filters are small enough to be considered negligible. Thus, with respect to parameter estimation, the joint and dual filters show little to no advantage over the single filter method. In other words,

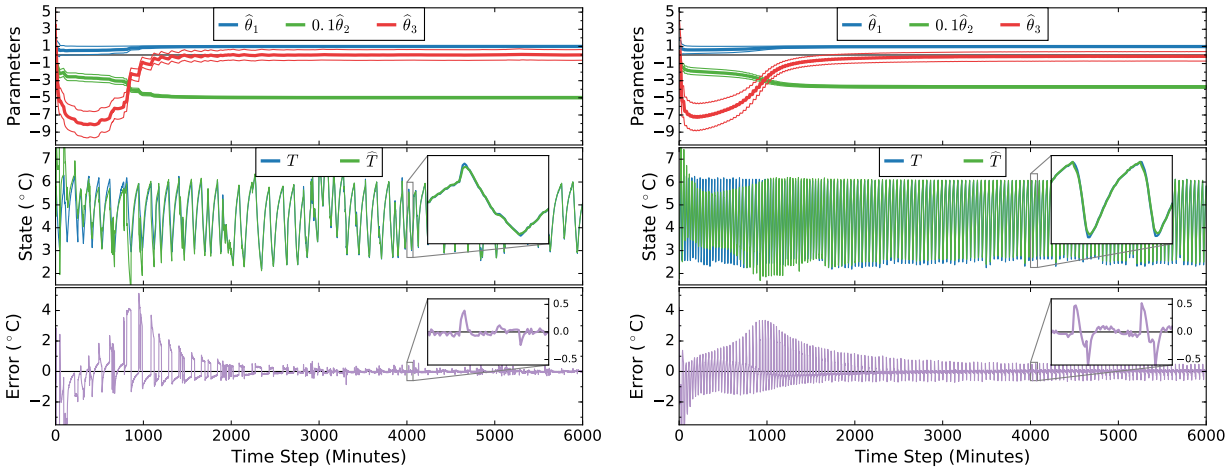

(a) TCL_1 (b) TCL_2

Figure 9: Joint Filter Parameter Estimation

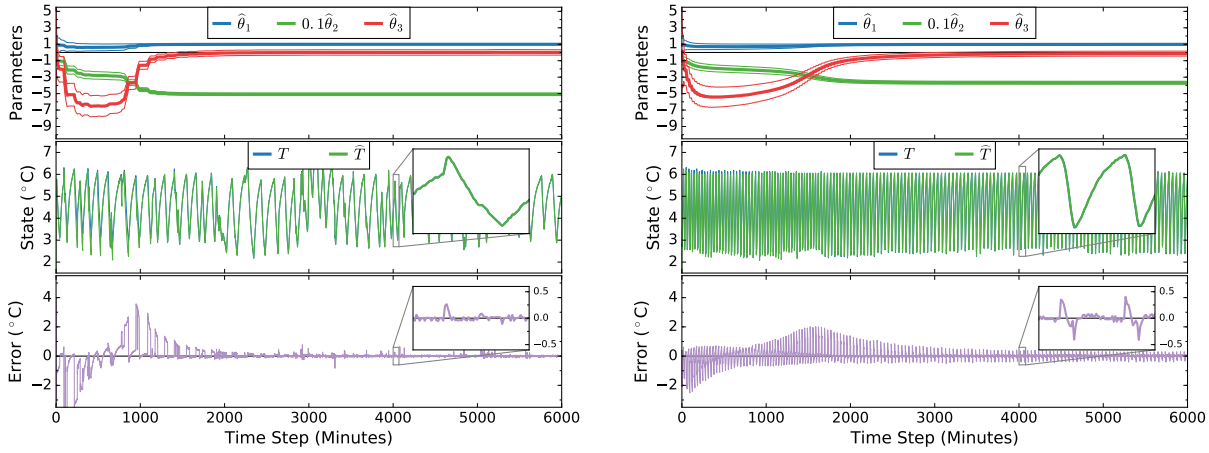

(a) TCL_1 (b) TCL_2

Figure 10: Dual Filter Parameter Estimation

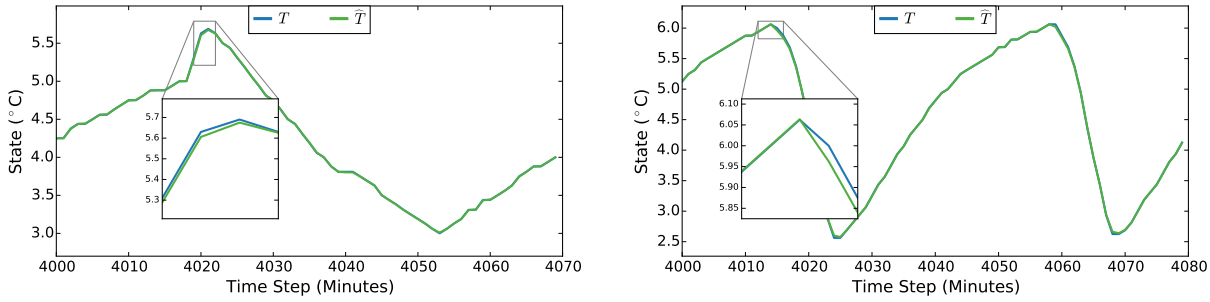

(a) TCL_1 (b) TCL_2

Figure 11: Dual Filter Temperature State Estimate

filtering the temperature measurement T does not appear to significantly improve the performance of our parameter estimation algorithm. Considering that we are estimating the parameters of two residential-sized refrigerators, this is not a surprising outcome. For a larger or noisier TCL, the joint and dual filters may yield a greater advantage, but for the TCLs used in this study, the estimation of T is simply unnecessary.

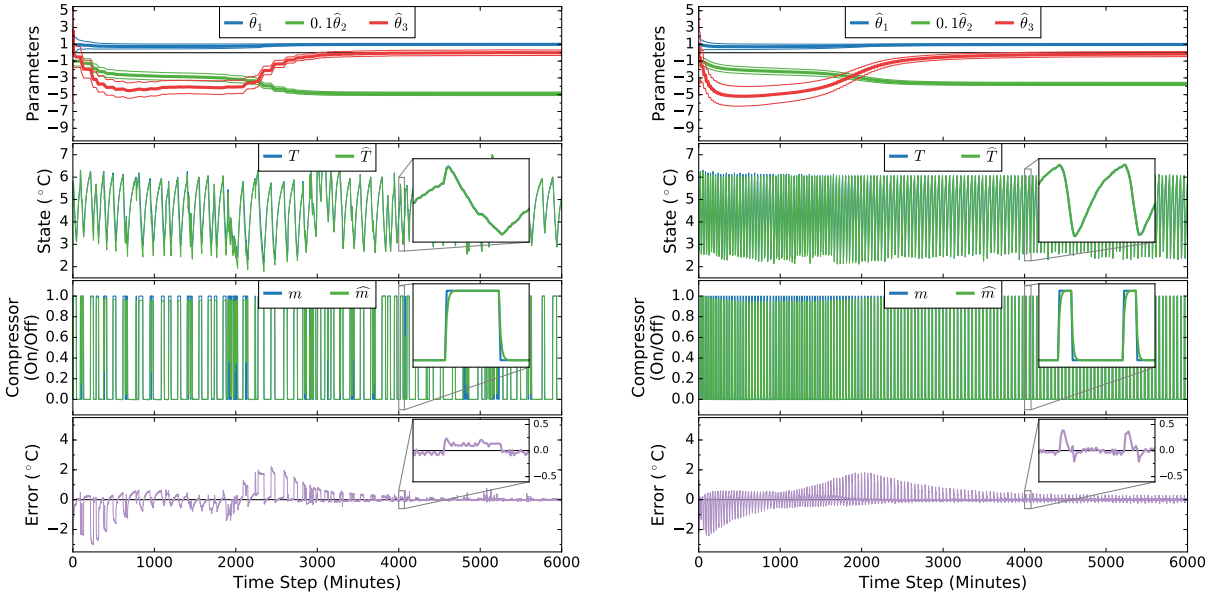

(a) TCL_1 (b) TCL_2

Figure 12: Triple Filter Parameter Estimation

Figure 12 presents the parameter estimation results using the triple filter method. In addition to the parameter estimate $\hat{\theta}^k$, state estimate $\hat{T}^k$, and parameter filter residual error r^k, the figure displays the compressor state estimates $\hat{m}^k$. The input filters also produce ambient temperature estimates $\hat{T}_\infty^k$. However, due to the negligible difference between the ambient temperature observations and estimates, we have excluded the results from this manuscript.

A sample of the compressor state estimates are plotted in Fig. 13. As shown in Fig. 13, the estimated compressor states $\hat{m}^k$ resemble a first-order system response rather than the discrete on/off state m^k given by the TCL model (1). One could argue that $\hat{m}^k$ better represents the thermodynamics of a TCL like the refrigerators used in this study. Specifically, while the compressor may instantaneously turn on or off (providing a step input), it takes some amount of time for the refrigeration cycle to start or stop removing heat (resulting in a first-order response). We could elect to model the refrigeration cycle as a first-order linear time-invariant system and to estimate the system parameters, however using the triple filter method to estimate m^k is a simple way to achieve a comparable result.

The moving window root mean square error (RMSE) results for each filter method are presented in

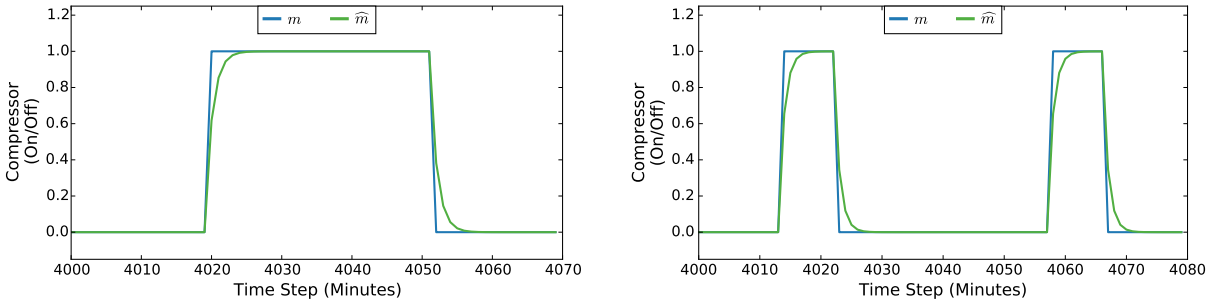

(a) TCL_1 (b) TCL_2

Figure 13: Triple Filter Compressor State Estimate

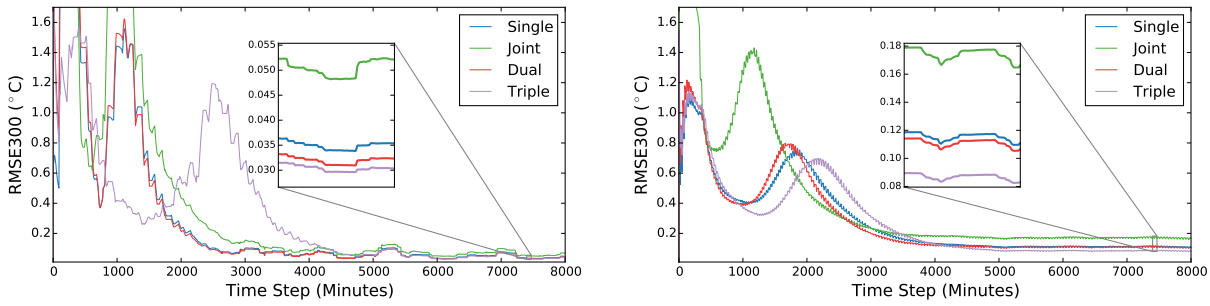

(a) TCL_1 (b) TCL_2

Figure 14: Parameter Estimation Moving Window RMSE

Figure 14. As stated previously, the RMSE300 at each time step is a measure of the RMSE over the last 300 time steps (i.e 5 hours). In this way, the RMSE300 is a function of the parameter filter's residual error r^k but provides a clearer means of comparing the performance of each filter method. As shown in Figure 14, the triple filter is the slowest to converge but performs slightly better than the other methods. The single and joint filter performances are comparable in TCL_1 whereas the joint method has the worst performance for TCL_2.

Overall, each filter method succeeds in performing parameter estimation and converges to comparable values. After convergence, the differences in the RMSE300 values are small enough to be considered negligible. In short, based on the experimental results, the joint, dual, and triple filter methods provide an insignificant advantage with respect to parameter estimation relative to the single filter method.

It is inconclusive whether this lack of advantage is attributable to the methods (i.e. filtering the states and inputs does not improve the parameter estimates) or to the systems (i.e. there was little noise in the refrigerators studied and therefore state and input estimation was unnecessary). Nonetheless, this manuscript presents a compiled collection of parameter estimation and filtering methods for online learning of TCLs.

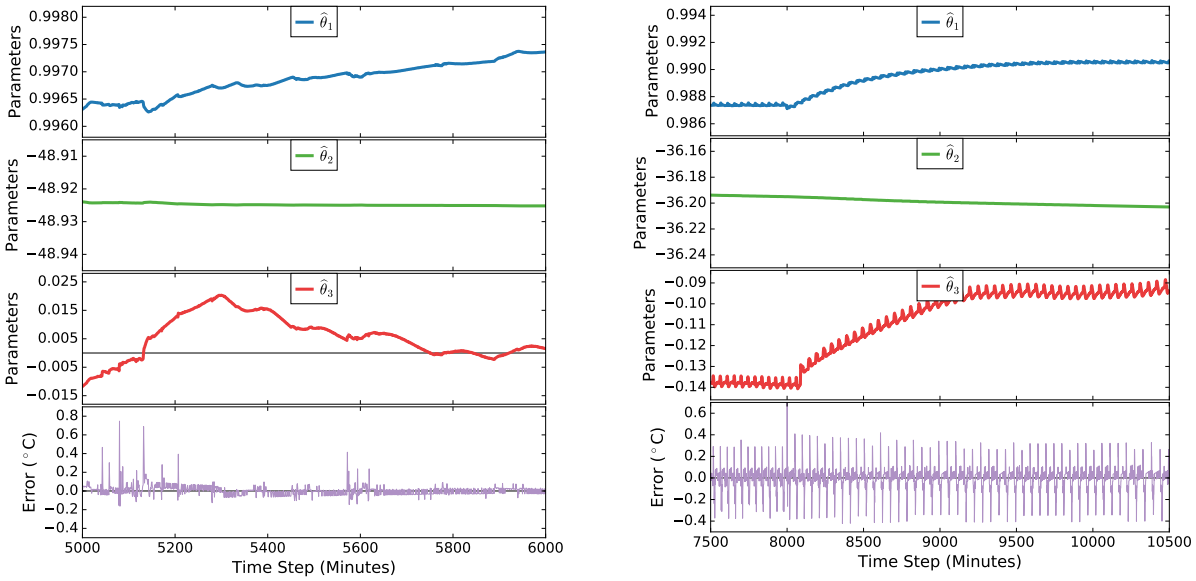

(a) TCL_1 (b) TCL_2

Figure 15: Adaptive Parameter Estimation of Experimental TCLs

Utilizing a recursive system identification technique provides the added benefit of allowing the parameter estimates to continuously adapt to changes in the system. This point is illustrated in Figure 15. Because TCL_1 is subject to normal residential use, the capacitance of the refrigerator regularly changes by relatively

small, random, and unobserved magnitudes. Figure 15(a) shows how the parameter estimates produced by the single filter method, particularly $\hat{\theta}_3$, respond to observations from the temperature sensors. In other words, the UKF updates the parameter estimates in order to minimize the residual error, thus allowing the model to dynamically adapt.

To more directly test the adaptive characteristics of our recursive system identification technique, we removed the 18 liters of water from TCL_2 at roughly time step $k = 8000$, thereby producing a step reduction in the thermal capacitance of the system. Figure 15(b) illustrates how the single UKF method responded to this step change by increasing both $\hat{\theta}_1$ and $\hat{\theta}_3$. According to the TCL model (1), we expect that decreasing the capacitance in the TCL increases $\hat{\theta}_1$, which represents the thermal characteristics of the conditioned mass ($\theta_1 = \exp(-h/RC)$). As stated previously, $\hat{\theta}_3$ represents a noise process accounting for energy gain or loss that is not directly modeled.

8. TCL Simulation Studies

8.1. Adaptive Parameter Estimation

In this section, we demonstrate the capability of the single UKF method to quickly converge to new parameter estimates in response to changes in the system. In this simulation study, we first generate data of a residential refrigerator with time-varying parameters. We then employ the single UKF method to recursively estimate the parameters using the temperature and compressor state signals. We have tuned the values of the covariance matrices Q_v and Q_w to allow the UKF to converge to new parameter estimates more quickly than in the previous TCL studies.

The refrigerator simulated in this study has parameters which vary periodically due to the loss or gain of thermal capacitance and changes in the unmodeled noise process. The change in thermal capacitance is represented by randomly drawing a value for θ_1 from the uniform distribution [0.985, 0.995] once an hour and drawing θ_3 from the uniform distribution [-0.16,-0.08] once every 3 hours. We assume that θ_2 is constant and equal to -40. Internal temperature and compressor state data is generated by simulating the deadband control of the system for 7 days using the random time-varying parameters and employing the ambient temperature data collected by TCL_1.

Next, we apply the single UKF method to recursively estimate the model parameters based on the simulated internal temperature and compressor state signals. Figure 16 presents examples of the time-varying system parameters and the parameter estimates produced by the single UKF method. The figure illustrates how the UKF is able to converge to new parameter estimates within a few time steps. We observe that the UKF often updates both θ_1 and θ_3 in response to changes in the thermal capacitance of the system. This observation is comparable to how the parameter estimates changed in response to removing thermal mass from TCL_2, as shown in Figure 15 above.

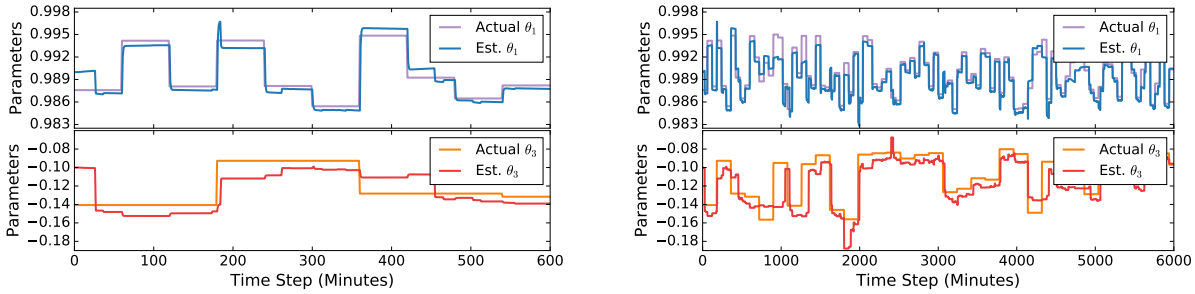

(a) Parameter estimates over 10 hours

(b) Parameter estimates over 100 hours

Figure 16: Adaptive Parameter Estimation of Simulated TCL

To quantify the advantage of using an adaptive modeling approach, we compare the performance of 3 observers which model the fridge and produce 1-hour ahead forecasts of the internal temperature and compressor state. The first observer, referred to as the "static" observer, employs a model with fixed parameters ($\theta_1 = 0.99$, $\theta_2 = -40$, and $\theta_3 = -0.12$). The second observer, referred to as the "adaptive" observer, employs the parameter estimates produced by the single UKF method. Note that the adaptive observer uses the model parameters as estimated at the time step that the forecast was produced and that the parameters do not vary within a 1-hour ahead forecast. Finally, the third observer, referred to as the "informed" observer, has perfect knowledge of the system parameters at the time that each forecast is produced. Note that the informed observer does not have knowledge of how the parameters will change in the future and thus the parameters do not vary within a 1-hour ahead forecast.

For each observer, we produce a 1-hour ahead forecast of the internal temperature and compressor state starting at each minute over a 7 day period. The performance of each observer is quantified as the root mean squared error (RMSE) of the forecasts compared to the simulated data with random time-varying parameters. The RMSE of the internal temperature forecasts for the static, adaptive, and informed observers are 1.37°C, 1.12°C, and 0.73°C, respectively, and for the compressor states are 0.402, 0.305, and 0.234, respectively. Based on these results, the adaptive observer reduces the temperature forecast RMSE by 24.1% and the compressor state forecast RMSE by 18.2% relative to the static observer. Assuming that the electrical power demand of the TCL is linearly proportional to the compressor state, the adaptive observer also reduces the power demand forecast RMSE by 18.2% relative to the static observer.

8.2. TCL Demand Response

In this section, we present a simulation study in which a population of residential refrigerators optimizes its power demand according to a price signal from a demand response event using model predictive control (MPC). Because the individual TCLs all respond to the same price signal and do not coordinate with each other or with a grid entity, the demand response approach can be described as decentralized model predictive control. The objective of this study is to illustrate the impact of model fidelity on the decision making of

Parameter	Refrigerator
Thermal resistance, R (°C/kW)	[80, 100]
Thermal capacitance, C (kWh/°C)	[0.4, 0.8]
Energy transfer rate, P (kW)	[-1, -0.2]
Coefficient of performance, COP	2
Temperature setpoint, T_{set} (°C)	[1.7, 3.3]
Deadband width, δ (°C)	[1, 2]
Ambient temperature, T_a (°C)	21

Table 2: TCL parameter ranges adopted from [29]

the TCLs. Specifically, we illustrate the advantage of using the single UKF method to recursively estimate a system's parameters rather than employing a fixed set of parameters for model predictive control. By more accurately representing the dynamics of a TCL, a controller with recursive parameter estimation is capable of identifying and implementing a better solution.

To avoid using electricity during the hour long event, the fridges are able to shift their demand by decreasing their temperature setpoints by 1°C for a certain amount of time. Note that the setpoint will never deviate by more than -1°C from the original temperature setpoint. The power demand of a TCL at each time step is defined by

$$p^k = \frac{|P|}{COP} m^k \tag{27}$$

where $p^k \in \mathbf{R}$ is the electric power demand (kW) and COP the coefficient of performance.

To determine when to change the setpoints, each refrigerator employs a receding horizon model predictive controller. Every 10 minutes, the controller generates a temperature setpoint schedule for the next 2 hours that minimizes the electricity costs. The refrigerator then implements the first 10 minutes of the 2-hour setpoint schedule. This procedure is repeated every 10 minutes over one day.

Setpoint changes are applied over 10 consecutive time steps. In other words, the 2-hour horizon is divided into 12 segments of 10 minutes. For each segment, the controller must decide whether to employ the original setpoint temperature or to decrease the setpoint by 1°C. Therefore the controller must consider 2^{12} possible 2-hour setpoint schedules. To find the optimal setpoint schedule, the controller simulates the refrigerator with each of the 2^{12} setpoint schedules and selects the schedule that has the lowest electricity cost.

In this study, we simulate the dynamics of 1000 refrigerators. To generate a diverse population of refrigerators, we employ published model parameter ranges, given in Table 2 and adopted from [29]. For each

refrigerator in the population, we randomly draw the R, P, T_{set}, and δ values from a uniform distribution between the maximum and minimum values shown in the table. These values remained fixed for a specific TCL. To represent periodic loss or gain of thermal capacitance and changes in the unmodeled noise process, we randomly vary the C and θ_3 parameters of each TCL. Once an hour, a new value of C is randomly drawn from the uniform distribution in Table 2. Once every 3 hours, a new θ_3 is drawn from the uniform distribution [-0.16,-0.08]. Finally, we assume a constant COP of 2 and ambient temperature T_∞ of 21°C.

To quantify the impact of model accuracy on the decision making, we implement 3 model predictive controllers. The first controller, referred to as the "static" controller, employs a model with fixed parameters ($C = 0.6$ and $\theta_3 = -0.12$). The second controller, referred to as the "adaptive" controller, employs parameter estimates produced recursively using the single UKF method. The third controller, referred to as the "informed" controller, has perfect knowledge of the system parameters at the time that each forecast is produced. We employ each of these controllers to simulate the TCLs in the population and to minimize the electricity costs by adjusting the temperature setpoints using the procedure described above.

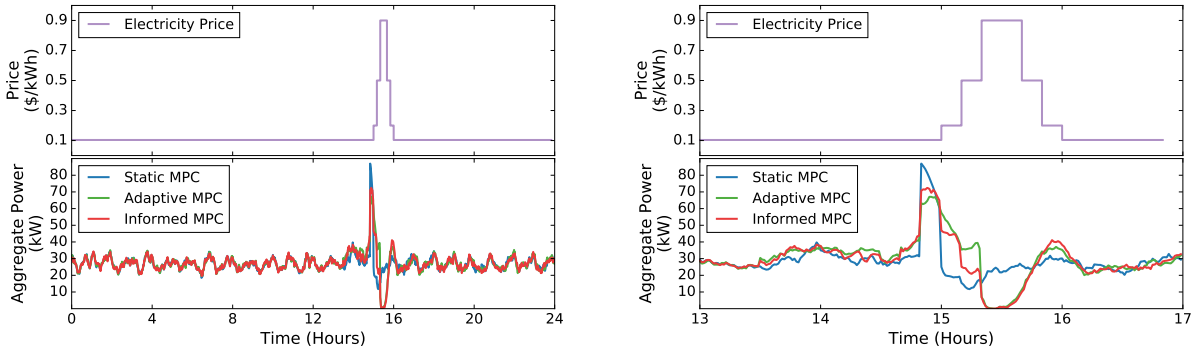

(a) Demand Response Simulation (Full Day)

(b) Demand Response Simulation (Price Event)

Figure 17: Demand Response of Simulated TCL Population

For each of the 1000 TCLs in the population, we simulate the receding horizon model predictive control (MPC) over one day using the static, adaptive, and informed controllers. The results of this study are illustrated in Figure 17. The top subplots show the demand response event pricing and the bottom subplots show the aggregate power demand of the refrigerator population. In Figure 17(b), we can observe how the controllers respond to the demand response pricing event and how the model accuracy impacts the capability of the controllers to optimally schedule the temperature setpoints. In other words, the static, adaptive, and informed controllers each produce an optimal temperature setpoint schedule using their respective TCL models. However, if the models do not accurately represent the dynamics of the system, the setpoint schedule implemented by the controller is likely to be suboptimal. This is illustrated in Figure 17 by the fact that the static controller does not avoid the peak pricing as well as the adaptive and informed controllers.

The average cost of operating a TCL in the population is \$0.0769 per day using the static controller, \$0.0740 per day using the adaptive controller, and \$0.0725 per day using the informed controller. Based on these results, the adaptive controller achieves a 3.71% cost savings compared to the static controller.

Additionally, the adaptive controller is more capable of achieving the desired demand response behavior of shifting load away from the peak price period. Within the 1-hour event, the average electricity costs of operating a TCL are \$0.01175 per hour, \$0.00855 per hour, and \$0.00702 per hour using the static, adaptive, and informed controllers, respectively. Within the 1-hour event, the adaptive controller achieves a 27.2% cost savings compared to the static controller. By recursively estimating the parameters using the single UKF method, the adaptive controller better enables the TCL to participate in demand response programs and other smart grid applications.

9. Conclusions

This paper examines online parameter estimation of thermostatically controlled loads (TCLs). First, we briefly discuss the Kalman filter (KF) and unscented Kalman filter (UKF) algorithms. Next, we present four filter methods (single, joint, dual, and triple) for recursively estimating the parameters of a discrete-time thermostatically controlled load (TCL) model. We also present experimental results using real temperature data from a 500W and a 100W residential refrigerator. For each of the four filter methods, the algorithm successfully converges to comparable parameter estimates and adapts to changing TCL characteristics.

Finally, we present 2 simulation studies. The first study demonstrates the capability of the single UKF method to quickly converge to new parameter estimates in response to changes in the system dynamics. The second study presents simulation results for a population of refrigerators which optimize their power demand based on a demand response electricity price event. These studies show the advantage of using model predictive control with the single UKF method rather than employing a fixed set of model parameters.

In this paper, we have applied the model predictive control of a TCL population to the relatively simple problem of demand response load shifting. In future work, we will incorporate the recursive parameter estimation methods presented here into model predictive control algorithms for ancillary services like frequency regulation and load following. As demonstrated in the demand response study above, utilizing a recursive parameter estimation method will improve the performance of a model predictive controller and therefore improve the capability of a TCL population to perform ancillary services.

A challenge not addressed in this manuscript is the estimation of TCL electric power demand and the coefficient of performance (COP). To enable TCL populations to perform ancillary services in power systems, it is important to be able to accurately estimate the power demand and to account for power dynamics such as inrush current in compressors. In future work, we plan to research models for forecasting TCL power demand which recognize nonlinearity due to the COP.

10. Appendix

10.1. Augmented Unscented Kalman Filter Algorithm

To estimate systems with non-additive noise, it is necessary to employ the *augmented* unscented Kalman filter algorithm. To present the augmented UKF algorithm, we will consider the state estimation of a discrete-time nonlinear dynamic system with *non-additive* noise given by the state-space model

$$x^k = F(x^{k-1}, u^k, v^k) \tag{28a}$$

$$y^k = H(x^k, u^k, w^k) \tag{28b}$$

where variable $x^k \in \mathbf{R}^X$ is the state of the system, $u^k \in \mathbf{R}^U$ is the known exogenous input, and $y^k \in \mathbf{R}^Y$ is the observed measurement signal. Function F is the transition model and the process noise $v^k \in \mathbf{R}^X$ has covariance $Q_v \in \mathbf{R}^{X\times X}$, $v^k \sim N(0, Q_v)$. Function H is the observation model and the measurement noise $w^k \in \mathbf{R}^Y$ has covariance $Q_w \in \mathbf{R}^{Y\times Y}$, $w^k \sim N(0, Q_w)$.

The UKF will model x^k as a GRV with estimated mean $\hat{x}^k$ and covariance Q_x^k. At each time step k, the UKF will generate sigma points for the previous state estimate, $\hat{x}^{k-1|k-1}$. For systems with non-additive noise, the state estimate and covariance is *augmented* with the process and measurement noise, as given by

$$\bar{s}^{k-1} = \begin{bmatrix} \hat{x}^{k-1|k-1} \\ \bar{v}^k \\ \bar{w}^k \end{bmatrix} \tag{29a}$$

$$Q_s^{k-1} = \begin{bmatrix} Q_x^{k-1|k-1} & 0 & 0 \\ 0 & Q_v & 0 \\ 0 & 0 & Q_w \end{bmatrix} \tag{29b}$$

where $\bar{v}^k$ and $\bar{w}^k$ are the mean of the process and measurement noises, respectively. In other words, if Gaussian, $\bar{v}^k \in \{0\}^X$ and $\bar{w}^k \in \{0\}^Y$. The dimensionality of $\bar{s}$ is therefore $L = 2X + Y$.

Next, the UKF will generate the sigma points. Because we are using the augmented state, we will introduce $\mathcal{S}_x$, $\mathcal{S}_v$, and $\mathcal{S}_w$, the sigma points associated with the state estimate, process noise, and measurement noise, respectively.

$$\mathcal{S}^{k-1} = UT_s(\bar{s}^{k-1}, Q_s^{k-1}) \tag{30a}$$

$$\mathcal{S}_{x,j}^{k-1} = \mathcal{S}_j^{k-1} \quad j = 0, \ldots, X-1 \tag{30b}$$

$$\mathcal{S}_{v,j}^{k-1} = \mathcal{S}_j^{k-1} \quad j = X, \ldots, 2X-1 \tag{30c}$$

$$\mathcal{S}_{w,j}^{k-1} = \mathcal{S}_j^{k-1} \quad j = 2X, \ldots, 2X+Y-1 \tag{30d}$$

where j refers to the rows of the L by $2L+1$ matrix $\mathcal{S}$.

In the prediction step, the UKF will propagate the sigma points through the process model and generate the a priori sigma points $\mathcal{X}_i^{k|k-1}$, state estimate $\hat{x}^{k|k-1}$, and covariance $Q_x^{k|k-1}$ as follows,

$$\mathcal{X}_i^{k|k-1} = F(\mathcal{S}_{x,i}^{k-1}, u^k, \mathcal{S}_{v,i}^{k-1}) \quad i = 0, \ldots, 2L \tag{31a}$$

$$\hat{x}^{k|k-1} = UT_m(\mathcal{X}^{k|k-1}) \tag{31b}$$

$$Q_x^{k|k-1} = UT_c(\mathcal{X}^{k|k-1}) \tag{31c}$$

In the correction step, the UKF will propagate the a priori sigma points through the measurement model to generate the measurement sigma points $\mathcal{Y}_i^k$, estimate $\hat{y}^k$, and covariance Q_y as follows,

$$\mathcal{Y}_i^k = H(\mathcal{X}_i^{k|k-1}, u^k, \mathcal{S}_{w,i}^{k-1}) \quad i = 0, \ldots, 2L \tag{32a}$$

$$\hat{y}^k = UT_m(\mathcal{Y}^k) \tag{32b}$$

$$Q_y = UT_c(\mathcal{Y}^k) \tag{32c}$$

These are used to calculate the cross-covariance Q_{xy}, the Kalman gain $\mathcal{K}$, and the observation error r^k. Finally, the state estimate and covariance are corrected, producing the a posterior estimate $\hat{x}^{k|k}$ and covariance $Q_x^{k|k}$.

$$Q_{xy} = \sum_{i=0}^{2L} \mathcal{W}_{c,i}(\mathcal{X}_i^{k|k-1} - \hat{x}^{k|k-1})(\mathcal{Y}_i^k - \hat{y}^k)^T \tag{33a}$$

$$\mathcal{K} = Q_{xy} Q_y^{-1} \tag{33b}$$

$$r^k = y^k - \hat{y}^k \tag{33c}$$

$$\hat{x}^{k|k} = \hat{x}^{k|k-1} + \mathcal{K} r^k \tag{33d}$$

$$Q_x^{k|k} = Q_x^{k|k-1} - \mathcal{K} Q_y \mathcal{K}^T \tag{33e}$$

where $r^k \in \mathbf{R}$ is the error between the measurement y^k and the estimate $\hat{y}^k$ at time step k.

To simplify notation, the non-additive unscented Kalman filter algorithm with augmented state can be expressed using the following 4 operator expressions

$$\begin{bmatrix} \mathcal{S}_x^{k-1} \\ \mathcal{S}_v^{k-1} \\ \mathcal{S}_w^{k-1} \end{bmatrix} = UKF_s \left(\begin{bmatrix} \hat{x}^{k-1|k-1} \\ Q_x^{k-1|k-1} \end{bmatrix}, \begin{bmatrix} \bar{v}^k \\ Q_v \end{bmatrix}, \begin{bmatrix} \bar{w}^k \\ Q_w \end{bmatrix} \right) \tag{34a}$$

$$\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \\ \mathcal{X}^{k|k-1} \end{bmatrix} = UKF_x\left(F, \mathcal{S}_x^{k-1}, u^k, \mathcal{S}_v^{k-1}\right) \tag{34b}$$

$$\begin{bmatrix} \hat{y}^k \\ Q_y^k \\ \mathcal{Y}^k \end{bmatrix} = UKF_y(H, \mathcal{X}^{k|k-1}, u^k, \mathcal{S}_w^{k-1}) \tag{34c}$$

$$\begin{bmatrix} \hat{x}^{k|k} \\ Q_x^{k|k} \\ r^k \end{bmatrix} = UKF_c\left(\begin{bmatrix} \hat{x}^{k|k-1} \\ Q_x^{k|k-1} \\ \mathcal{X}^{k|k-1} \end{bmatrix}, \begin{bmatrix} \hat{y}^k \\ Q_y^k \\ \mathcal{Y}^k \end{bmatrix}, y^k\right) \tag{34d}$$

where (34a) corresponds to (29-30), (34b) to (31), (34c) to (32), and (34d) to (33).

The additive UKF presented in section 5.2 can be considered a common special case of the augmented UKF algorithm. Specifically, $\bar{s}^{k-1} = \hat{x}^{k-1|k-1}$, $Q_s^{k-1} = Q_x^{k-1|k-1}$, and there are no sigma points for the process and measurement noise. This reduces the computational complexity of each iteration of the algorithm from $O((2X+Y)^3)$ for the augmented UKF to $O(X^3)$ for the additive UKF where X is the dimensionality of the state space and Y the dimensionality of the observation space.

References

[1] D. S. Callaway, I. Hiskens, et al., Achieving controllability of electric loads, Proceedings of the IEEE 99 (1) (2011) 184–199.

[2] J. L. Mathieu, S. Koch, D. S. Callaway, State estimation and control of electric loads to manage real-time energy imbalance, Power Systems, IEEE Transactions on 28 (1) (2013) 430–440.

[3] E. M. Burger, S. J. Moura, Generation Following with Thermostatically Controlled Loads via Alternating Direction Method of Multipliers Sharing Algorithm (October 2015).
URL `http://escholarship.org/uc/item/2m5333xx`

[4] J. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, et al., Demand response spinning reserve demonstration. lawrence berkeley national laboratory, berkeley, ca. prepared for energy systems integration, Public Interest Energy Research Program, California Energy Commission.

[5] J. L. Mathieu, D. S. Callaway, State estimation and control of heterogeneous thermostatically controlled loads for load following, in: System Science (HICSS), 2012 45th Hawaii International Conference on, IEEE, 2012, pp. 2002–2011.

[6] R. Malhame, C.-Y. Chong, Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system, Automatic Control, IEEE Transactions on 30 (9) (1985) 854–860.

[7] S. Moura, V. Ruiz, J. Bendsten, Modeling heterogeneous populations of thermostatically controlled loads using diffusion-advection PDEs, in: ASME 2013 Dynamic Systems and Control Conference, American Society of Mechanical Engineers, 2013, pp. V002T23A001–V002T23A001.

[8] A. Ghaffari, S. Moura, M. Krstić, Modeling, control, and stability analysis of heterogeneous thermostatically controlled load populations using partial differential equations, Journal of Dynamic Systems, Measurement, and Control 137 (10) (2015) 101009.

[9] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, Energy Conversion and Management 50 (5) (2009) 1389–1400.

[10] S. Bashash, H. K. Fathy, Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads, in: American Control Conference (ACC), 2011, IEEE, 2011, pp. 4546–4553.

[11] S. Koch, J. L. Mathieu, D. S. Callaway, Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services, in: Proc. PSCC, 2011, pp. 1–7.

[12] J. A. Short, D. G. Infield, L. L. Freris, Stabilization of grid frequency through dynamic demand control, Power Systems, IEEE Transactions on 22 (3) (2007) 1284–1293.

[13] Z. Xu, J. Ostergaard, M. Togeby, C. Marcus-Moller, Design and modelling of thermostatically controlled loads as frequency controlled reserve, in: Power Engineering Society General Meeting, 2007. IEEE, IEEE, 2007, pp. 1–6.

[14] D. Angeli, P.-A. Kountouriotis, A stochastic approach to dynamic-demand refrigerator control, Control Systems Technology, IEEE Transactions on 20 (3) (2012) 581–592.

[15] S. H. Tindemans, V. Trovato, G. Strbac, Decentralized control of thermostatic loads for flexible demand response, Control Systems Technology, IEEE Transactions on 23 (5) (2015) 1685–1700.

[16] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, Distributed optimization and statistical learning via the alternating direction method of multipliers, Foundations and Trends in Machine Learning 3 (1) (2011) 1–122.

[17] A. Mercurio, A. Di Giorgio, F. Purificato, Optimal fully electric vehicle load balancing with an admm algorithm in smartgrids, in: Control & Automation (MED), 2013 21st Mediterranean Conference on, IEEE, 2013, pp. 119–124.

[18] H. Xing, Y. Mou, Z. Lin, M. Fu, Fast distributed power regulation method via networked thermostatically controlled loads, Vol. 19, Cape Town, South Africa, 2014, pp. 5439 – 5444.

[19] M. Liu, Y. Shi, Distributed model predictive control of thermostatically controlled appliances for providing balancing service, in: 2014 IEEE 53rd Annual Conference on Decision and Control (CDC), 2014, pp. 4850–4855.

[20] G. T. Costanzo, F. Sossan, M. Marinelli, P. Bacher, H. Madsen, Grey-box modeling for system identification of household refrigerators: A step toward smart appliances, in: Energy (IYCE), 2013 4th International Youth Conference on, IEEE, 2013, pp. 1–5.

[21] F. Sossan, V. Lakshmanan, G. T. Costanzo, M. Marinelli, P. J. Douglass, H. Bindner, Grey-box modelling of a household refrigeration unit using time series data in application to demand side management, Sustainable Energy, Grids and Networks 5 (2016) 1–12.

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

[23] E. M. Burger, H. E. Perez, S. J. Moura, Piecewise Linear Thermal Model and Recursive Parameter Estimation of a Residential Heating System (January 2015).
URL `http://escholarship.org/uc/item/8kx450mg`

[24] R. E. Mortensen, K. P. Haggerty, A stochastic computer model for heating and cooling loads., IEEE Transactions on Power Systems 3 (3) (1998) 1213–1219.

[25] S. J. Julier, J. K. Uhlmann, New extension of the kalman filter to nonlinear systems, in: AeroSense'97, International Society for Optics and Photonics, 1997, pp. 182–193.

[26] E. A. Wan, R. Van Der Merwe, A. T. Nelson, Dual Estimation and the Unscented Transformation, in: NIPS, Citeseer, 1999, pp. 666–672.

[27] E. A. Wan, R. V. D. Merwe, The Unscented Kalman Filter, in: Kalman Filtering and Neural Networks, Wiley, 2001, pp.

221–280.

[28] E. Wan, R. Van Der Merwe, et al., The unscented Kalman filter for nonlinear estimation, in: Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000, IEEE, 2000, pp. 153–158.

[29] J. L. Mathieu, M. Dyson, D. S. Callaway, Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations, ACEEE Summer Study on Energy Efficiency in Buildings.