



Temperature prediction at critical points in district heating systems

Pinson, Pierre; Nielsen, Torben Skov; Nielsen, Henrik Aalborg; Poulsen, Niels Kjølstad; Madsen, Henrik

Published in:
European Journal of Operational Research

Link to article, DOI:
[10.1016/j.ejor.2007.11.065](https://doi.org/10.1016/j.ejor.2007.11.065)

Publication date:
2009

Document Version
Early version, also known as pre-print

[Link back to DTU Orbit](#)

Citation (APA):
Pinson, P., Nielsen, T. S., Nielsen, H. A., Poulsen, N. K., & Madsen, H. (2009). Temperature prediction at critical points in district heating systems. *European Journal of Operational Research*, 194(1), 163-176.
<https://doi.org/10.1016/j.ejor.2007.11.065>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Temperature prediction at critical points in district heating systems

P. Pinson^{a,*}, T.S. Nielsen^a, H.Aa. Nielsen^a, N.K. Poulsen^a,
H. Madsen^a,

^a*Informatics and Mathematical Modelling, Technical University of Denmark, 2800
Kgs. Lyngby, Denmark*

Abstract

Current methodologies for the optimal operation of district heating systems use model predictive control. Accurate forecasting of the water temperature at critical points is crucial for meeting constraints related to consumers while minimizing the production costs for the heat supplier. A new forecasting methodology based on conditional Finite Impulse Response (cFIR) models is introduced, for which model coefficients are replaced by coefficient functions of the water flux at the supply point and of the time of day, allowing for nonlinear variations of the time delays. Appropriate estimation methods for both are described. Results are given for the test case of the Roskilde district heating system over a period of more than 6 years. The advantages of the proposed forecasting methodology in terms of a higher forecast accuracy, its use for simulation purposes, or alternatively for better understanding transfer functions of district heating systems, are clearly shown.

Key words: forecasting, time series, finite impulse response, nonlinear time delay, district heating systems.

* Corresponding author. Pierre Pinson, Informatics and Mathematical Modelling, Technical University of Denmark, bg. 321 (020), 2800 Kgs. Lyngby, Denmark. Tel: +45-4525 3428; fax: +45-4588 2673

Email address: pp@imm.dtu.dk (P. Pinson).

1 Introduction

District heating systems consist of centralized heat production facilities with associated distribution networks. They play an important role in Nordic countries, where they are used to meet the demand for space heating and hot tap water. Owing to this centralized production combined with complex network architectures, decisions made from the supply point of view have highly significant economic impacts. In order to operate district heating systems optimally from an economic point of view, control strategies are implemented with some restrictions e.g. a minimum guaranteed inlet temperature at the consumers. The aim of these control strategies is to meet these restrictions while minimizing the supply temperature, and thus the production costs for the heat supplier. The general problem of flow and temperature control at the supply point is described in (Madsen et al. 1994). Alternative methodologies based on predictive control have been introduced by Nielsen (2002) and Sandou et al. (2004).

If considering a single heat supplier, the decision variables are the magnitude of the water flux and the supply temperature. The magnitude of the water flux is directly imposed by the load. Thus, load forecasts serve as a basis for making decisions on the flux. The accuracy of load forecasts has been discussed by Nielsen and Madsen (2000) or Dotzauer (2002). Similarly, predictive control methods applied to supply temperature necessitate a statistical model permitting forecasts of water temperature at critical points of a network. The relevant forecast horizons may be up to 12 hours ahead. Increasing the accuracy of these temperature forecasts is expected to significantly lower the production costs for the heat suppliers, as a consequence of them making more efficient control decisions. The aim of the present paper is to contribute to reaching a higher forecasting accuracy by proposing a new forecasting methodology.

The models in the literature either derive from a physical description of the heat and mass transfers in the network (Sandou et al. 2005), or they are based on a statistical description of the transfer function from the supply point to the critical point considered (Søgaard 1993). Detailed physical modelling may lead to the assumption that time delay in the network is fixed, owing to potentially large computational costs of making it time-varying. The proposed forecast methodology in (Søgaard 1993) is to set an ensemble of ARMAX (Auto-Regressive Moving Average with Exogenous input) models with different fixed time delays, and to switch between models depending on some estimated current time delay. In contrast, the forecasting methodology introduced here makes it possible to account for a varying time delay within the statistical model itself. This model is a Finite Impulse Response (FIR) for which the model coefficients are replaced by nonparametric coefficient functions of influential variables. Owing to this consideration of the nonlinear influence of external factors on the FIR, the model is described as conditional Finite Impulse Response (cFIR). The proposed cFIR models are used here to capture the nonlinear influence of the water flux at the supply point on the transfer function of the district heating system. Another interesting point about cFIR models is that they can account for the influence of the social behaviour of the consumers on the temperature at the critical point. This may be done either by having the time of the day as an influential variable of the cFIR model, or more classically, by having an offset term in the form of a diurnal harmonic. This way, the impact of the consumers' social behaviour on the transfer function of district heating

systems will be discussed.

The problem of predicting the water temperature at critical points of a district heating system is described in section 2, as well as the proposed forecasting methodology. Then, section 3 introduces the method for model coefficient estimation. Particularly, it allows for a recursive estimation of the coefficients so that it accommodates long-term variations. In addition, regularization of the recursive estimation method is proposed to enhance its generalization ability and to control its multi-step ahead accuracy. The case-study of the district heating system of Roskilde in Denmark is considered in section 4 in order to illustrate the benefits of this new forecasting methodology. The original dataset includes temperature and flux measurements at the supply point, as well as temperature measurements at 3 critical points of the distribution network, for a period of more than 6 years. In addition to demonstrating its significantly higher prediction performance, the possibility offered by this new methodology to better understand time delays in networks is discussed. Section 5 ends the paper by summarizing the main conclusions and gathering perspectives regarding future developments.

2 Description of the proposed forecasting methodology

A district heating system often consists of a complex network. It may not be necessary to model it as a whole, as one is mainly interested in what occurs at some specific points of this network. These specific points serve as references for designing and optimizing the control strategies, and are thus referred to as critical points. Focusing on a single critical point, the overall network is conceptually simplified: it is considered that there is a unique simple pipe between the supply and critical points. The district heating system operator injects continuous quantities of warm water (at a controlled temperature), and is interested in knowing what will be the temperature at this critical point depending on his operation strategy at the supply point. Denote by x_t and u_t the value of the flux and of the water temperature at the supply point at time step t . Let y_t be the water temperature at the critical point considered at that same time. The problem is here discretized. Since district heating systems are slowly-varying systems, and since a very fine temporal resolution would not yield substantial benefits (while increasing computational costs), the suitable time step is typically one hour. The available data consist of the time-series $\{x_t\}$, $\{u_t\}$ for the supply point and $\{y_t\}$ for the critical point, all including n observations. The state-of-the-art statistical approach is introduced in a first part, followed by the description of our new forecasting methodology based on a cFIR model.

2.1 The state-of-the-art statistical approach

The model traditionally used for predicting the temperature at critical points, initially proposed by Sogaard (1993), takes the form of a linear transfer function model with a first-order autoregressive component

$$y_t = a_1 y_{t-1} + b_0 (h_t - \tau) u_{t-\tau} + b_1 (h_t - \tau - 1) u_{t-\tau-1} + b_2 (h_t - \tau - 2) u_{t-\tau-2} + \varepsilon_t, \quad \forall t \quad (1)$$

where h_t is the hour of the day corresponding to the time step t , $\{\varepsilon_t\}$ is the noise sequence, such that $\mathbb{E}[\varepsilon] = 0$ and $\sigma_\varepsilon^2 < \infty$. While a_1 is not conditional on any variable, the coefficients b_j , $j = 0, 1, 2$, of the transfer function are made a function of the time of the day and of the time delay τ in the system, by using a Fourier harmonics with a period of 24 hours

$$b_j(h_t - \tau - j) = b_j^0 + b_j^1 \sin\left(\frac{\pi(h_t - \tau - j)}{12}\right) + b_j^2 \cos\left(\frac{\pi(h_t - \tau - j)}{12}\right), \quad \forall t, \quad j = 0, 1, 2 \quad (2)$$

i.e. they account for the diurnal variations in the system behaviour, owing to the social behaviour of the consumers. This model structure requires that τ is fixed. When using this model for predictive control, [Nielsen \(2002\)](#) proposed choosing τ as the time delay that maximizes the correlation between the time series $\{y_t\}$ and $\{u_{t-\tau}\}$. It can be allowed to change over time by using a sliding window, thus yielding different time delays in the system depending on the time of the year.

The above model may provide an acceptable description of $\{y_t\}$. Even though the autoregressive term in (1) may be justified from physical consideration of the heat capacity and resultant inertia of the system, such term has the effect of dramatically lowering forecast accuracy as the lead time progresses. Indeed, when issuing at time t a k -step ahead forecast $\hat{y}_{t+k|t}$, the model is fed with the $(k-1)$ -step ahead forecast $\hat{y}_{t+k-1|t}$. A consequence is that forecasting errors directly sum up as k increases. This will be illustrated in section 4, where model (1) will be used as a benchmark. Another drawback of the model, which significantly affects its performance when used for forecasting purposes, is the fixed time delay τ , which is not realistic. In practice, τ not only varies depending on the season: it is also a function of the flow in the network ([Arvatson 2001](#)).

2.2 The proposed forecasting methodology

2.2.1 Modeling the transfer function of the network

Owing to the drawbacks of model (1), it is proposed to introduce a new model with a varying time delay, and without any autoregressive component. For modelling the transfer function of the network, the conditional Finite Impulse Response (cFIR) model initially described by [Nielsen \(2000\)](#) is used as a basis, i.e.

$$y_t = \sum_{j \in S_j} \beta_j(x_{t-1}) u_{t-j} + \varepsilon_t, \quad \forall t \quad (3)$$

where y_t is the temperature at the considered critical point at time t , x_{t-1} is the value of the flux at the supply point at time $t-1$ and u_{t-j} are the lagged values of the temperature at the supply point. S_j corresponds to the finite set of indexes related to the lagged values for the cFIR model. $\{\varepsilon_t\}$ is a white noise sequence, for which $\mathbb{E}[\varepsilon_t] = 0$ and $\sigma_{\varepsilon_t}^2 < \infty$.

The advantage of model (3) is that the cFIR is conditional on the flux at the supply point, since the coefficients β_j are coefficient functions of x_{t-1} . This way, the time delay in the system is also made a (nonlinear) function of the flux. However, since the model does not integrate any component describing the autocorrelation of the $\{y_t\}$ time-series, it may

require the cardinal of S_j to be large. In addition, even though potential heat losses may be accounted for in the stationary gain of the cFIR model, it is preferred here to include such effects in an offset term (also function of the flux), so that model (3) becomes

$$y_t = \beta_0(x_{t-1}) + \sum_{j \in S_j} \beta_j(x_{t-1})u_{t-j} + \varepsilon_t, \quad \forall t \quad (4)$$

2.2.2 Integrating the social behaviour of consumers

Model (4) is expected to provide an adequate description of the flux-dependent transfer function of the distribution network. However, it does not account for the potential influence of the social behaviour of consumers, i.e. for their consumption pattern (as a function of the hour of the day) that necessarily depends on the type of consumers connected to the critical point considered. Intuitively the consumption pattern will differ if the critical point corresponds to a residential area, to an industrial area, or to a hospital.

There may be two alternative views on how to integrate the social behaviour of consumers within model (4). On the one hand, one may consider that it does not impact the transfer function of the network. In this case, it is only necessary to modify the offset term of the model, in order for it to exhibit diurnal variations. It is proposed here to use Fourier harmonics of period 24 hours, so that the offset term β_0 in (4) is replaced by

$$\beta_0^\dagger(x_{t-1}, h_t) = \beta_{0,0}^\dagger(x_{t-1}) + \beta_{0,1}^\dagger(x_{t-1}) \sin\left(\frac{\pi h_t}{12}\right) + \beta_{0,2}^\dagger(x_{t-1}) \cos\left(\frac{\pi h_t}{12}\right), \quad \forall t \quad (5)$$

with h_t the hour of the day at time step t , while the other β_j^\dagger coefficient functions remain unchanged

$$\beta_j^\dagger(x_{t-1}) = \beta_j(x_{t-1}), \quad \forall j, j > 0 \quad (6)$$

This model will be referred to as a *rigid* cFIR in the following.

The β_j^\dagger coefficient functions can be gathered in a vector denoted by $\boldsymbol{\beta}^\dagger(x_{t-1}, h_t)$. Let \mathbf{u}_t^\dagger be the corresponding vector of ones, harmonics values, as well as lagged values of u_t . Then, the rigid cFIR model can simplify to

$$y_t = \boldsymbol{\beta}^{\dagger\top}(x_{t-1}, h_t)\mathbf{u}_t^\dagger + \varepsilon_t, \quad \forall t \quad (7)$$

On the other hand, one may consider that the social behaviour of the consumer also influences the transfer function of the network. Then, this translates to having the β_j coefficients' functions in (4) as a function of the time of the day, in addition to being a function of the flux at the supply point. This yields the alternative model

$$y_t = \beta_0^*(x_{t-1}, h_t) + \sum_{j \in S_j} \beta_j^*(x_{t-1}, h_t)u_{t-j} + \varepsilon_t, \quad \forall t \quad (8)$$

that is, with the β_j^* coefficient functions being a function of both the flux x_{t-1} at the supply point and the time of the day h_t . Owing to its more supple structure, the model

will be referred to as a *supple* cFIR in the following. In the same manner as for the rigid cFIR, denote by $\boldsymbol{\beta}^*(x_{t-1}, h_t)$ the vector of coefficient functions for this model and by \mathbf{u}_t^* the corresponding vector of ones and lagged values of u_t . In such case, model (8) becomes

$$y_t = \boldsymbol{\beta}^{*\top}(x_{t-1}, h_t)\mathbf{u}_t^* + \varepsilon_t^*, \quad \forall t \quad (9)$$

To sum up, the difference between the *supple* and rigid cFIR models is that the former has coefficient functions conditional on both the flux at the supply point and the time of the day, while for the latter they are only conditional on flux values.

2.2.3 Forecasting with the cFIR models

Models (7) and (9) describe the temporal evolution of $\{y_t\}$ from past information at the supply point, i.e. measurements of flux and supply temperature. From these models, the one-step ahead prediction at time t of the temperature at the critical point can be defined as the conditional expectation of the process at time $t+1$ given the information set Ω_t up to time t , and the chosen cFIR model. If denoting by $\hat{y}_{t+1|t}^\dagger$ and $\hat{y}_{t+1|t}^*$ the one-step ahead predictions with the rigid and *supple* cFIR models, this gives

$$\hat{y}_{t+1|t}^\dagger = \mathbb{E}[y_{t+1}|\boldsymbol{\beta}^\dagger, \Omega_t], \quad \hat{y}_{t+1|t}^* = \mathbb{E}[y_{t+1}|\boldsymbol{\beta}^*, \Omega_t], \quad \forall t \quad (10)$$

In practice, since $\mathbb{E}[\varepsilon^\dagger] = 0$ and $\mathbb{E}[\varepsilon^*] = 0$, the one-step ahead predictions are obtained with

$$\hat{y}_{t+1|t}^\dagger = \boldsymbol{\beta}^\dagger(x_t, h_{t+1})^\top \mathbf{u}_{t+1}^\dagger, \quad \hat{y}_{t+1|t}^* = \boldsymbol{\beta}^*(x_t, h_{t+1})^\top \mathbf{u}_{t+1}^*, \quad \forall t \quad (11)$$

for the rigid and *supple* cFIR models, respectively.

Note that for computing one-step ahead predictions with cFIR models, the current filtered flux value x_t and the lagged supply temperature values u_{t-j+1} , $j \in S_j$, are available measurements. For multi-step ahead prediction, say k -step ahead, the above equations become

$$\hat{y}_{t+k|t}^\dagger = \boldsymbol{\beta}^\dagger(\hat{x}_{t+k-1|t}, h_{t+k})^\top \mathbf{u}_{t+k}^\dagger, \quad \hat{y}_{t+k|t}^* = \boldsymbol{\beta}^*(\hat{x}_{t+k-1|t}, h_{t+k})^\top \mathbf{u}_{t+k}^*, \quad \forall t \quad (12)$$

Since u is the control variable, future values of u can be assumed to be known. In contrast, future flux values at the supply point cannot be exactly known, though they could be deduced from load predictions, since the load q_t at time t is related to x_t with

$$q_t = c_p x_t (u_t - v_t), \quad \forall t \quad (13)$$

where v_t is the return temperature at that same time and c_p is the heat capacity of the water in the pipe. It is known that the variations of v_t are very smooth, so that future values can be assumed to be known or could be accurately predicted. Predictions of future flux at the supply point could then be straightforwardly obtained and used to feed the model. Note that the weather directly impacts the load in district heating systems, and consequently the flux at the supply point. Since the flux is a variable of the cFIR coefficient

functions, the impact of the weather on the temperature at critical points of the network is implicitly considered.

Actually, both flux and temperature at the supply point are control variables in practice. Since the flux variable is closely related to the load as explained above, it can be assumed that the supply temperature is controlled separately, using flux forecasts or scenarios as input. Restrictions on the range of potential values and variations of both flux and temperature variables may also appear as constraints in this complex decision-making problem. In all cases, the cFIR models can be used for simulation purposes in order to evaluate the impact of decisions on the flux and temperature variables in the subsequent hours. A forecasting exercise will be considered in section 4 in order to illustrate the sensitivity of the prediction performance of the cFIR models to the choice of future flux values. In this exercise, the supply temperature is the unique control variable and a simple Auto-Regressive (AR) model of order p is used to describe $\{x_t\}$, and consequently for multi-step ahead prediction. This model is

$$x_t = a_0 + \sum_{i=1}^p a_i x_{t-i} + \xi_t, \quad \forall t \quad (14)$$

where $\{\xi_t\}$ is a white noise sequence, i.e. such that $\mathbb{E}[\xi] = 0$ and $\sigma_\xi^2 < \infty$. Such a modelling approach may be less appropriate than that mentioned above, so that the results given in section 4 can be considered as a lower bound on the potential performance of the proposed forecasting methodology in comparison to the case where more advanced flux predictions are available. This will allow us to perform a sensitivity analysis on the prediction performance of the cFIR models even if the information on future flux at the supply point is not perfectly accurate. After inspection of the correlogram of the model residuals, it has been decided to enhance this AR(p) model with a Fourier harmonics of period 24 hours. This makes it possible to account for periodic diurnal variations in the $\{x_t\}$ time series that cannot be captured by the autoregressive component only. Model (14) then becomes

$$x_t = a_0^0 + a_0^1 \sin\left(\frac{\pi h_t}{12}\right) + a_0^2 \cos\left(\frac{\pi h_t}{12}\right) + \sum_{i=1}^p a_i x_{t-i} + \xi_t, \quad \forall t \quad (15)$$

where h_t is the hour of the day corresponding to time step t . Multi-step ahead flux forecasts obtained from model (15) are then used to feed the cFIR model as expressed in (12).

3 Estimation of the model parameters

For a model such as (15), the model parameters can easily be estimated with a Least Squares (LS) or Recursive Least Squares (RLS) method, as described in e.g. (Ljung and Söderström 1983). In contrast, for the case of cFIR models, the chosen method for nonparametric parameter estimation is described in the following, both for offline and online applications. It combines local polynomial regression, weighted LS for the offline case – and respectively weighted RLS with exponential forgetting for the online case, as well as Tikhonov regularization (Tikhonov and Arsenin 1977). For simplicity, the method is described for a generic cFIR model whose transfer function is described by $\beta(\mathbf{r})$, with $\mathbf{r} = [r_1 \ r_2 \ \dots \ r_l]$

the vector of variables that condition the cFIR model. l should be kept to a low value, say below 3, owing to the curse of dimensionality (Hastie and Tibshirani 1990, pp. 83-84). When necessary, specific points related to the estimation of the two cFIR models introduced above will be discussed. Note that periods with non-optimal operation of the district heating system may affect estimation of model coefficients. In general when fitting statistical models, it is assumed that the data characteristics are similar over both the training and evaluation periods, thus translating to a continuous optimal operation of the district heating system. In addition, a benefit of online adaptive estimation combined with regularization proposed here is to diminish and even accommodate potential effects of periods with non-optimal operation.

3.1 Local polynomial estimates

The coefficient functions $\beta_j(\mathbf{r})$ are estimated in a nonparametric framework, i.e. without assuming a shape for these functions. This is done by using local polynomial regression (Cleveland and Devlin 1988), for which the only assumption regarding the β_j coefficient functions is that they are sufficiently smooth to be locally approximated with polynomials. The estimation problem is reduced to locally fitting linear models at a number m of fitting points $\mathbf{r}_{(i)}$. A fitting point is defined by a pair of flux and time values (i.e. in our case $l = 2$)

$$\mathbf{r}_{(i)} = [r_{(i),1} \ r_{(i),2} \ \dots \ r_{(i),l}]^\top, \quad i = 1, \dots, m \quad (16)$$

so that these m fitting points span the range of potential values on the various dimensions of \mathbf{r} . Defining these fitting points is best done by using information on the distribution of \mathbf{r} . For the case of the cFIR models introduced above, this mainly concerns the distribution of flux values, as hour values will be uniformly distributed.

Let us focus on the fitting point $\mathbf{r}_{(i)}$ only. The local polynomial approximation \mathbf{z}_t of the vector of explanatory variables \mathbf{u}_t at $\mathbf{r}_t = [r_{t,1} \ r_{t,2} \ \dots \ r_{t,l}]^\top$ is given by

$$\mathbf{z}_t = [u_{t,1}\mathbf{p}_d^\top(\mathbf{r}_t) \ \dots \ u_{t,k}\mathbf{p}_d^\top(\mathbf{r}_t) \ \dots \ u_{t,l}\mathbf{p}_d^\top(\mathbf{r}_t)]^\top \quad (17)$$

where $\mathbf{p}_d(\mathbf{r}_t)$ corresponds to the d -order polynomial evaluated at \mathbf{r}_t . For instance if $d = 2$, $\mathbf{p}_2(\mathbf{r}_t)$ can be obtained as

$$\mathbf{p}_2(\mathbf{r}_t) = [1 \ r_{t,1} \ r_{t,2} \ r_{t,1}^2 \ r_{t,1}r_{t,2} \ r_{t,2}^2]^\top \quad (18)$$

In parallel, write

$$\boldsymbol{\phi}_{(i)} = \boldsymbol{\phi}(\mathbf{r}_{(i)}) = [\boldsymbol{\phi}_{(i),1}^\top \ \dots \ \boldsymbol{\phi}_{(i),k}^\top \ \dots \ \boldsymbol{\phi}_{(i),l}^\top]^\top \quad (19)$$

the vector of local coefficients at $\mathbf{r}_{(i)}$, where the element vector $\boldsymbol{\phi}_{(i),k}$ is the vector of local coefficients related to the local polynomial approximation of the k -th explanatory variable, that is, $u_{t,k}\mathbf{p}_d(\mathbf{r}_{(i)})$.

The nonlinear cFIR is thus locally approximated at $\mathbf{r}_{(i)}$ by the linear model

$$y_t = \mathbf{z}_t^\top \boldsymbol{\phi}_{(i)}, \quad \forall t \quad (20)$$

so that the problem of fitting the nonlinear cFIR model is converted in a number m of local linear models to be fitted, that is, one for each fitting point $\mathbf{r}_{(i)}$.

3.2 Offline estimation of the local coefficients

In an offline setting, a set of n observations for each of the time series is available and one then wants to estimate the local coefficients for this set of data. In this type of setting, with the one-step ahead prediction defined as the conditional expectation (cf. section 2.2.3), the nonlinear cFIR model can be fitted by minimizing the sum of squared residuals over the set of observations, that is,

$$S(\boldsymbol{\beta}) = \sum_{t=1}^n \rho(y_t - \boldsymbol{\beta}^\top(\mathbf{r}_t)\mathbf{u}_t) \quad (21)$$

where ρ is a quadratic criterion, i.e. such that $\rho(\epsilon) = \epsilon^2/2$.

Then, if focusing on a given fitting $\mathbf{r}_{(i)}$, one can estimate the vector of corresponding local coefficients, that is, $\boldsymbol{\phi}_{(i)}$, by using weighted least-squares. The estimate $\hat{\boldsymbol{\phi}}_{(i)}$ is then given by

$$\hat{\boldsymbol{\phi}}_{(i)} = \arg \min_{\boldsymbol{\phi}_{(i)}} S(\boldsymbol{\phi}_{(i)}) = \arg \min_{\boldsymbol{\phi}_{(i)}} \sum_{t=1}^n w_{t,(i)} \rho(y_t - \mathbf{z}_t^\top \boldsymbol{\phi}_{(i)}) \quad (22)$$

where the weights $w_{t,(i)}$ are assigned by a Kernel function of the following form

$$w_{t,(i)} = T \left(\prod_{k=1}^l \frac{|r_{t,k} - r_{(i),k}|_k}{\hat{h}_{(i),k}(\alpha_k)} \right) \quad (23)$$

In the above, $|\cdot|_k$ denotes the chosen distance on the k^{th} dimension of \mathbf{r} . For the cFIR models introduced in section 2.2, one would choose an Euclidian distance on the dimension of the flux values and a polar distance on the dimension of hour values.

In (23), $\hat{h}_{(i),k}$ is the bandwidth for that particular fitting point $\mathbf{r}_{(i)}$ and for the k^{th} dimension of \mathbf{r} . Whatever the dimension considered, $\hat{h}_{(i),k}$ is determined by using a nearest-neighbour principle (Nielsen et al. 2000). For a chosen proportion α_k , the bandwidth $\hat{h}_{(i),k}(\alpha_k)$ is such that

$$\alpha_k = \int_{\mathbb{D}_{(i),k}} f_{r_k}(v) dv \quad (24)$$

where $\mathbb{D}_{(i),k} = \{v \in \mathbb{R} \mid |v - r_{(i),k}|_2 < \hat{h}_{(i),k}\}$ defines the neighborhood of $\mathbf{r}_{(i)}$ on the k^{th} dimension of \mathbf{r} , while f_{r_k} denotes the density function of the r_k values. In practice, f_{r_k} is replaced by the empirical distribution function of the available data.

Finally in (23), T is defined as the tricube function, i.e.

$$T : v \in \mathbb{R}^+ \rightarrow T(v) \in [0, 1], \quad T(v) = \begin{cases} (1 - v^3)^3, & v \in [0, 1] \\ 0 & , v > 1 \end{cases} \quad (25)$$

as introduced and discussed by e.g. [Cleveland and Devlin \(1988\)](#).

Denote by \mathbf{A} the data matrix such that its t^{th} row is \mathbf{z}_t^\top , i.e. the transpose of the local polynomial approximation of \mathbf{u}_t . Write \mathbf{W} the diagonal matrix for which the t^{th} element on the diagonal corresponds to the weight $w_{t,(i)}$ to assign to the data point \mathbf{r}_t . The solution of (22) is then straightforwardly given by

$$\hat{\phi}_{(i)} = (\mathbf{A}^\top \mathbf{W} \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{W} \mathbf{y} \quad (26)$$

where \mathbf{y} is the vector of observations for the time-series $\{y_t\}$.

The elements of $\beta_{(i)}$ are finally obtained with

$$\hat{\beta}_{(i)} = \hat{\beta}(\mathbf{r}_{(i)}) = \mathbf{p}_d^\top(\mathbf{r}_{(i)}) \hat{\phi}_{(i)}, \quad i = 1, \dots, m \quad (27)$$

The vector of coefficient function values $\hat{\beta}(\mathbf{r}_t)$ evaluated at a given data point \mathbf{r}_t is obtained by linear-type interpolation. For instance, if the number of variables in coefficient functions is such that $l = 2$, it is done by using bilinear interpolation of the coefficient function values at the four fitting points located around \mathbf{r}_t .

3.3 Online estimation of the local coefficients

For real-world applications, one does not want to consider the whole set of available observations for estimating the local coefficients every time new observations become available. Instead, for this online setting, one aims at tracking the local coefficients by using a recursive formulation of the estimation method. In addition, a recursive formulation can allow for an exponential forgetting of old observations, which leads to the model being adaptive with respect to long-term variations in the process characteristics. From here on, it is considered that at time n a set of n past observations is available for each time-series, and thus that the dataset grows as time increases.

First, let us introduce the objective function to be minimized at each time n , which consists of a modified version of that given by (21)

$$S_n(\beta) = \sum_{t=1}^n \lambda^{n-t} \rho(y_t - \beta^\top(\mathbf{r}_t) \mathbf{u}_t) \quad (28)$$

where λ is the forgetting factor that permits an exponential forgetting of past observations.

For a given λ , $\lambda \in [0, 1)$, the effective number of observations n_λ is

$$n_\lambda = 1 + \lambda + \lambda^2 + \dots = \frac{1}{1 - \lambda} \quad (29)$$

Denote by $\hat{\phi}_{n,(i)}$ the estimate of the local coefficients for the fitting point $\mathbf{r}_{(i)}$ at time n . Then, the objective function to be minimized for estimating the local coefficients at $\mathbf{r}_{(i)}$ and at time n is

$$S_n(\phi_{n,(i)}) = \sum_{t=1}^n \Lambda_{n,(i)}(t) w_{t,(i)} \rho(y_t - \mathbf{z}_t^\top \phi_{n,(i)}) \quad (30)$$

where $\phi_{n,(i)}$ is related to $\beta_n(\mathbf{r}_{(i)})$ following a relation equivalent to (27). In parallel, $\Lambda_{n,(i)}$ is the function that permits exponential forgetting of past observations, i.e.

$$\Lambda_{n,(i)}(t) = \begin{cases} \lambda_{n,(i)}^{\text{eff}} \Lambda_{n-1,(i)}(t-1), & 1 \leq t \leq n-1 \\ 1 & , \quad i = n \end{cases} \quad (31)$$

In the above definition, $\lambda_{n,(i)}^{\text{eff}}$ is the effective forgetting factor for the fitting point $\mathbf{r}_{(i)}$ which makes it possible to account for the weighting in the formulation of (30). This follows the definition given by [Nielsen et al. \(2000\)](#), which tells that $\lambda_{n,(i)}^{\text{eff}}$ is a function of $w_{n,(i)}$ such that

$$\lambda_{n,(i)}^{\text{eff}} = 1 - (1 - \lambda)w_{n,(i)} \quad (32)$$

where λ is the classical user-defined forgetting factor, $\lambda \in [0, 1)$. This effective forgetting factor ensures that old observations are downweighted only when new information is available. By using this exponential forgetting scheme, n_λ as given by (29) consists of a lower bound on the effective number of observations ([Nielsen et al. 2000](#)).

The local coefficients $\hat{\phi}_{n,(i)}$ at time n for model (20) are then given by

$$\hat{\phi}_{n,(i)} = \arg \min_{\phi_{(i)}} S_n(\phi_{(i)}) = \arg \min_{\phi_{(i)}} \sum_{t=1}^n \Lambda_{n,(i)}(t) w_{t,(i)} \rho(y_t - \mathbf{z}_t^\top \phi_{(i)}) \quad (33)$$

The recursive formulation for an adaptive estimation of the local coefficients $\hat{\phi}_{n,(i)}$ (and therefore of $\hat{\beta}_{n,(i)}$, by using Equation (27) at each time-step) leads to the following three-step updating procedure at time n :

$$\epsilon_{n,(i)} = y_n - \mathbf{u}_n^\top \hat{\beta}_{n-1,(i)} \quad (34)$$

$$\hat{\phi}_{n,(i)} = \hat{\phi}_{n-1,(i)} + \epsilon_{n,(i)} w_{n,(i)} \left(\mathbf{R}_{n,(i)} \right)^{-1} \mathbf{z}_n^\top \quad (35)$$

$$\mathbf{R}_{n,(i)} = \lambda_{n,(i)}^{\text{eff}} \mathbf{R}_{n-1,(i)} + w_{n,(i)} \mathbf{z}_n \mathbf{z}_n^\top \quad (36)$$

where $\lambda_{n,(i)}^{\text{eff}}$ is again the effective forgetting factor. One sees that when the weight $w_{n,(i)}$

equals 0 (thus meaning that the local estimates should not be affected by the new information), then one has $\hat{\phi}_{n,(i)} = \hat{\phi}_{n-1,(i)}$ and $\mathbf{R}_{n,(i)} = \mathbf{R}_{n-1,(i)}$. This confirms the role of the effective forgetting factor, i.e. downweight old observations, but only when new information is available.

For initializing the recursive process, the matrices $\mathbf{R}_{0,(i)}$, $i = 1, \dots, m$, can be chosen as

$$\mathbf{R}_{0,(i)} = \delta \mathbf{I}, \quad \forall i \quad (37)$$

where δ is a small positive number and \mathbf{I} is an identity matrix of appropriate size. Note that this size, which corresponds to the number of coefficients to be estimated, is equal to the order of the chosen model in Equation (20) times the order of the polynomials used for local approximation. In parallel, the coefficient functions are initialized with a vector of zeros, or alternatively from a best guess of the target regression.

3.4 Regularization for a better generalization ability

The cFIR models (7) and (9) were originally designed to perform one-step ahead prediction. However, they are used here for multi-step ahead forecasting purposes with flux predictions as input. Therefore one should try not to amplify the error in flux forecasts when passed through the model. This can be done by ensuring that one works with a ‘smooth’ model. In addition, using recent data for fitting the model does not ensure an optimal performance when subsequently used for predicting on new and independent data. This ability to perform well with new and independent data is referred to as the generalization ability of the model (see e.g. Stone (1974)). For these two reasons, a regularized version of the estimation method described in section 3.3 is proposed here.

Several approaches may be considered for regularization in recursive least squares methods. They are widely used for ill-conditioned numerical problems, for avoiding overfitting when training neural networks (Leung et al. 1999; Sjöberg and Ljung 1995) or more generally for estimation in nonlinear systems (Bishop 1995; Johansen 1997). Following (Tikhonov and Arsenin 1977), it is chosen to apply the so-called Tikhonov regularization, which consists of adding a penalty term related to the norm of the coefficients (or of their derivatives) in the loss function to be minimized for model fitting.

For the case of the offline estimation of the model coefficients, it is well known that adding Tikhonov regularization means that equation (26) becomes

$$\hat{\phi}_{(i)} = \left(\mathbf{A}^\top \mathbf{W} \mathbf{A} + \mu \mathbf{I} \right)^{-1} \mathbf{A}^\top \mathbf{W} \mathbf{y} \quad (38)$$

where \mathbf{I} is an identity matrix of appropriate size, and where μ is the regularization parameter that makes it possible to control the trade-off between the minimization of the fitting errors and the norm of the model estimates. One thus sees that Tikhonov regularization consists of penalizing the diagonal elements of the inverse covariance matrix, equivalently to Ridge regression.

For the case of online estimation of the cFIR model coefficients, the loss function to be

minimized at each time t can be reformulated as

$$\tilde{S}_n(\boldsymbol{\phi}_{n,(i)}) = \frac{\mu}{1-\lambda} \boldsymbol{\phi}_{n,(i)}^\top \boldsymbol{\phi}_{n,(i)} + \sum_{i=1}^n \Lambda_{n,(i)}(t) w_{n,(i)} \rho(y_t - \mathbf{z}_t^\top \boldsymbol{\phi}_{n,(i)}) \quad (39)$$

where $\boldsymbol{\phi}_{n,(i)}^\top \boldsymbol{\phi}_{n,(i)}$ represents the quadratic norm of the model estimates. The forgetting factor λ is such that $\lambda < 1$. The regularization parameter μ is here multiplied by $(1-\lambda)^{-1}$, which corresponds to the effective number of observations for this loss function formulation. By doing so, μ represents the regularization load to be added to each observation accounted for in the loss function. Thus, the regularization is independent of the size of the virtual sliding window considered, which in turn is controlled by λ . If no exponential forgetting was used, a single parameter μ would multiply the norm of the model estimates. Choosing μ in such a case would be an issue, as the result of the sum in (39) would increase as n increases, while μ is not a function of n . Note that when regularization in RLS methods is considered in the literature, see e.g. (Hubing and Alexander 1991; Ismail and Principe 1997), it is always for the case $\lambda = 1$ and with effect of the regularization fading as the size of the dataset increases. The aim of such regularization is mainly to control first adaptation steps after model initialization.

For the adaptive formulation of the loss function in (39), the recursive procedure for updating the cFIR model coefficients at each time step can be rewritten in order to account for regularization. In fact, the main difference with the classical updating procedure described above relates to the updating formula for the inverse covariance matrix $\mathbf{R}_{n,(i)}$. One then works with a regularized inverse covariance matrix $\tilde{\mathbf{R}}_{n,(i)}^*$, which replaces $\mathbf{R}_{n,(i)}$ in (35), and which is updated with

$$\tilde{\mathbf{R}}_{n,(i)} = \lambda_{n,(i)}^{\text{eff}} \tilde{\mathbf{R}}_{n-1,(i)} + w_{n,(i)} \mathbf{z}_n \mathbf{z}_n^\top + \frac{1 - \lambda_{n,(i)}^{\text{eff}}}{1 - \lambda} \mu \mathbf{I} \quad (40)$$

where \mathbf{I} is an identity matrix of appropriate size. The inverse covariance matrix $\tilde{\mathbf{R}}_{(i),0}$ can be initialized with $\delta = \mu$ in equation (37). Such a recursive scheme for the updating of the covariance matrices ensures that its diagonal elements are always penalized with the same regularization parameter $\mu(1-\lambda)^{-1}$ (as is the case in (38)). The model estimates are still updated by applying (35) at each time step.

4 Case-studies and results

4.1 Description of the case-studies

The models and forecasting methodology are applied to the test case of the Roskilde district heating system. Roskilde is a city of around 80,000 inhabitants located 30 kilometres west of Copenhagen in Denmark. The district heating system investigated has a single heat production facility, and focus is given to three critical points on the distribution network. They correspond to a Viking museum, the local hospital (Sankt Hans) and a retirement home. They are hereafter referred to as critical points 1, 2 and 3, respectively.

The available data consist of measurements of the water flux and temperature at the supply point, as well as measurements of the water temperature at the three critical points. They have a time resolution of 5 minutes. Since the temporal resolution of the temperature controller installed at Roskilde is one hour, this is also the temporal resolution chosen for temperature forecasts. Hourly data are obtained by averaging the 5-minute measurements so that, for instance, the hourly temperature value at 01:00 is the average of all values between 00:05 and 01:00. If more than two measurements are missing for a given hour, the corresponding hourly value is considered as erroneous. The period for which measurements are available goes from August 16th, 2000, to December 12th, 2006. This translates to 55423 hourly values for each variable. For the specific case of critical point 2, the last 11000 data points are not considered owing to suspicious behaviour of the time-series, which is in turn due to the application of local control strategies of the water temperature at Roskilde hospital since 2005. The overall percentage of valid data for the three critical points are 88.49%, 86.28% and 94.56%, respectively.

The aim of the present exercise is to demonstrate the significantly higher performance of the proposed forecasting methodology in comparison with the state-of-the-art approach described in section 2.1. But also, the structure of the proposed models, and the two rival approaches with the integration of the social behaviour of the consumers, will allow us to discuss the time delays in the network and the assumption that social behaviour impacts (or does not impact) the transfer function of the distribution network.

4.2 Setup and optimal fitting of the cFIR models

The cFIR models (7) and (9) comprise the central part of the forecasting methodology. For both models the set of indices that defines the past values of the water temperature at the supply point is such that $S_j = \{1, 2, \dots, 10\}$, i.e. the cFIR models rely on the last 10 hourly values of supply temperature.

Local polynomial approximation is restricted to first order polynomials, in order to limit the number of local coefficients to be estimated, and thus the related computational costs. Also, because of the yearly cycle in the flux values, one year of flux values is used for obtaining a representative empirical distribution \hat{f}_x . This year serves to set the fitting points and the proportions that define the nearest-neighbour bandwidths. The fitting points along the flux dimension $x_{(i)}$, $i = 1, \dots, m_x$, are then defined such that there is the same proportion of flux values in each of the intervals defined by two successive fitting points. This writes

$$\int_0^{x_{(i)}} f_{\hat{x}}(v)dv = \frac{i-1}{m_x-1}, \quad i = 1, \dots, m_x \quad (41)$$

The number of fitting points has an impact on the computational costs for model estimation, provided that m_x is set to a sufficient value so that local polynomial approximation is suitable. It has been found that the model fitting was not significantly improved by having m above 11, and so m is set to this value. The related bandwidth values $h_{(i),x}$, $i = 1, \dots, m_x$ are obtained by applying the nearest-neighbour principle introduced in (24), parameterized by α_x . Similarly, it has been witnessed that the improvement in the model

Table 1

Results from the cross-validation procedure for the choice of the optimal forgetting factor λ and regularization parameter μ . These results are for the rigid cFIR model (a) and supple cFIR model (b).

(a) rigid cFIR model			
critical point number	λ	μ	mean RMSE [$^{\circ}\text{C}$]
1	0.98	0.00075	0.4851
2	0.98	0.0008	0.4902
3	0.9	0.1	0.8210

(b) supple cFIR model			
critical point number	λ	μ	mean RMSE [$^{\circ}\text{C}$]
1	0.975	0.001	0.5014
2	0.98	0.0006	0.5304
3	0.9	0.1	0.8108

fitting was negligible when having α_x higher than 0.4, and therefore this value is chosen in the following. For the specific case of the supple cFIR model, four fitting points are uniformly distributed over the range of daily hours. The local coefficients for the supple cFIR models will hence be estimated for 00:00, 6:00, 12:00 and 18:00. In order to have very smooth variations along this dimension of the supple cFIR, the bandwidth is set to a large value ($\alpha_h = 0.6$).

With regard to the recursive estimation method itself, one has to set the value of the forgetting factor λ , which defines the rate of forgetting of old observations, and controls the ability of the method to account for long-term variations in the process characteristics. However, if a too low value for λ is chosen, the fitting of the model will be very poor. Finally, a last parameter to consider is the regularization parameter μ , defining the trade-off between model fitting and the smoothness of the model estimates. Our methodology for defining optimal values for these two parameters is to use one-fold cross validation. The first year of the available dataset is used as a training period, while the second year is seen as the validation set. Since it is considered that cFIR models will be used in practice for generating temperature predictions up to 12 hours ahead, the criterion to be minimized on the validation set is the Root Mean Square Error (RMSE), averaged over this range of forecast horizons. For the purpose of model fitting, only flux measures are used.

The results from the cross-validation procedure are collated in Table 1. For the three critical points, the optimal parameters λ and μ are very similar for the rigid or supple cFIR models. For two out of the three critical points, the mean RMSE related to the optimal parameters is slightly higher for the supple cFIR model than for the rigid one. This is not the case for the third critical point. However, the value of the optimal forgetting factor is low (0.9), while the regularization parameter value is very high (0.1) in comparison with the other critical points. The mean RMSE also reaches a surprisingly high level

in comparison with the two other critical points. This might indicate some abnormal behaviour in the data, owing to e.g. the quality of the temperature measurements. One should note though that this critical point corresponds to Roskilde hospital, for which local temperature control strategies may sometimes be applied. This may thus affect the temperature measurements used for modelling the transfer function between the supply point and Roskilde hospital.

A nice feature of the cFIR models is the low sensitivity to the choice of the two parameters λ and μ . For instance, Figure 1 shows the variation of the average RMSE on the validation set for the supple cFIR model applied to critical point 1, as a function of λ (y -axis) and μ (x -axis). These variations are described by a contour plot, with 30 level lines corresponding to uniformly distributed mean RMSE values. The minimum average RMSE is obtained for $(\lambda, \mu) = (0.975, 0.001)$. One sees that the 3-dimensional surface describing the variations is a smooth and convex surface. The convex nature of this surface means that there is a unique (λ, μ) combination that minimizes the mean RMSE value on the cross-validation set. Also, the fact that this surface is smooth demonstrates the low sensitivity of the performance of the cFIR models to the choice of the parameters λ and μ .




Figure 1. Mean Root Mean Square Error [°C] over the validation set as a function of both the forgetting factor λ (y -axis) and the regularization parameter μ (x -axis). These results are for the supple cFIR model and for critical point 1.

In a last stage, the estimated cFIR model for the transfer function between the supply point and critical point 1 is shown in Figure 2. This gives the amplitude of the coefficient functions in the cFIR model as a function of both the water flux and the lag at the supply point. For high flux values, the transfer function is highly concentrated: the water temperature at the critical point 1 is almost exclusively determined by that at the supply point 2 hours before. A similar behaviour can be observed for flux values down to $1000\text{m}^3\text{h}^{-1}$. For lower values, the transfer function is less concentrated and more lagged values of the water temperature at the supply point contribute to determining this at critical point 1. This may be explained by a different mixing of the water in the pipe owing to this lower

flux. In addition, it has been noticed (for both cFIR models and for the three critical points) that the offset term in the cFIR tends to decrease as the flux values become lower, indicating higher temperature losses.

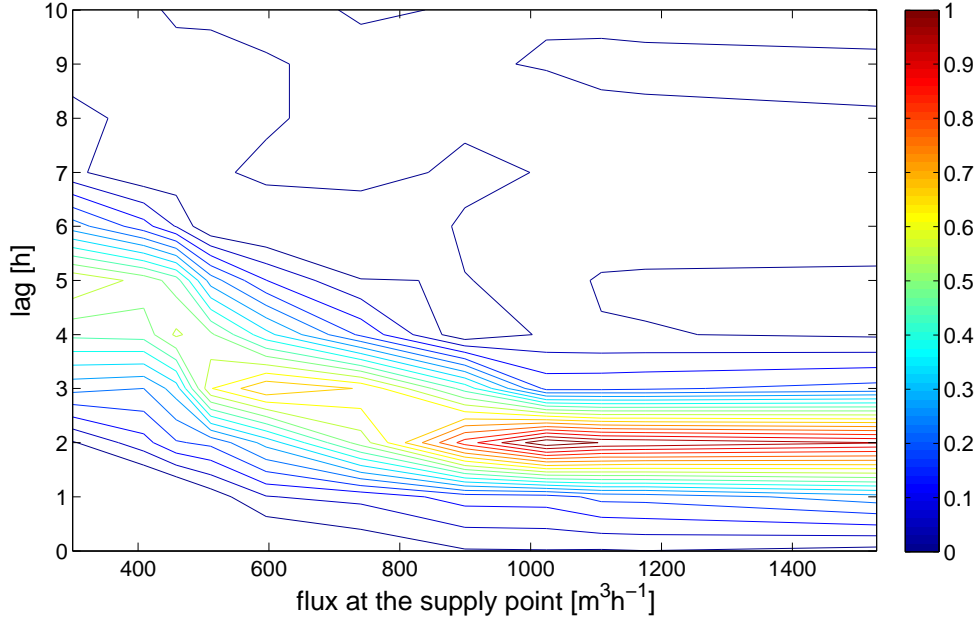


Figure 2. Contour plot of the amplitude of the coefficient functions of the rigid cFIR model at the end of the validation set for critical point 1. The x -axis gives the flux at the supply point (in m^3h^{-1}) while the y -axis gives the lag of the transfer function in hours.


Figure 2 enables better appraisal of the physical behaviour of the district heating system between the supply point and a given critical point. The influence of the flux on the losses has already been mentioned above. In addition to this, the figure illuminates the influence of the flux at the supply point on the time delays in the network. Here, this varies between 2 and 5 hours when going from the highest to lowest flux values. The time delay variations appear to be nonlinear, though this may come from some numerical artefact. Indeed, the same regularization parameter μ is used for all fitting points. However, since an optimal regularization parameter (for a linear model) is related to the variance and the norm of model estimates, see e.g. (Golub et al. 2000; Wang and Chow 1989), the optimal μ for the estimation of the coefficient functions in the cFIR should also be a function of the fitting point considered. This is because both the variance and the norm of the coefficient functions will vary depending on the fitting point. The optimal and local tuning of the regularization parameter μ should be further investigated in the future. Note that for the case of the supple cFIR models, one could also visualize contour plots such as in Figure 2 for different hours of the day. Since variations in the shape of the transfer functions through the day have not been found to be highly significant, this point is not further discussed here.

4.3 Out-of-sample evaluation of the cFIR models


The remainder of the dataset is seen as an out-of-sample evaluation set, for which the application of the forecasting methodology has to mimic operational conditions, so that observed performance is representative. The evaluation set consists of 33900 data points (~ 4 years) for critical points 1 and 3, and of 22500 data points (slightly less than three years) for critical point 2. The cFIR models estimated for the three critical points make it possible to describe the transfer function between the supply point and each of these critical points. For control purposes in operational conditions, these models can serve to simulate what would be the temperature at a given critical point for the following hours, depending on the chosen control strategy (on both flux and water temperature variables) at the supply point. Therefore, for the out-of-sample evaluation of the cFIR models, the measurements of both the flux and the water temperature at the supply point are used as input to the models.

Several criteria are considered for evaluating the prediction performance, namely bias, Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). All these criteria are calculated on a per-horizon basis, since it is expected that the prediction performance deteriorates as the lead time progresses. The RMSE criterion is used as the main criterion since it relates to the quadratic loss function used for model estimation (and parameter selection through the cross-validation procedure). The range of temperature variations over the whole dataset are 31.26°C , 28.58°C and 30.53°C for critical points 1, 2 and 3, respectively. The full set of results from the evaluation are available in (Pinson et al. 2007), while only the main results are discussed here.

For both types of cFIR models, bias values are at a low level whatever the look-ahead time, indicating that there is no systematic part in the prediction error. Figure 3 summarizes the evaluation carried out with the MAE and RMSE criteria, depicted as a function of the look-ahead time. There is a slight increase in the prediction error as the lead time progresses for the first 2 critical points, while this increase is more pronounced for the third point. In addition, the average level of prediction error is significantly higher for this last critical point, as was also the case when fitting the cFIR models with the cross-validation procedure. The MAE averaged over the forecast length ranges between 0.291 and 0.597°C depending on the cFIR models and the critical point considered. For the first 2 critical points, there is significant difference in the accuracy of predictions generated with the supple and rigid cFIR models, with a clear advantage for the latter ones. However for the third critical point their forecast accuracies are much more similar, with an advantage for the supple cFIR models for horizons further than 5-6 hours ahead. This may be explained by the fact that diurnal effects (which have more importance as the look-ahead time progresses) at Roskilde hospital cannot be properly accounted for with a simple first-order harmonic function. Hence, the nonparametric approach used in supple cFIR models becomes a relevant alternative.



(a) MAE as a function of the look-ahead time.



(b) RMSE as a function of the look-ahead time.

Figure 3. Out-of-sample evaluation of the forecast performance of both the rigid and supplc cFIR models for the three critical points. Flux measurements are used as input to the cFIR models. The forecast performance is evaluated with MAE and RMSE error measures as a function of the look-ahead time up to 12-hour ahead.

4.4 Sensitivity to the flux values used as input to the cFIR models

In this section, a sensitivity analysis on the performance of the proposed forecasting methodology is carried out. Especially, focus is given to the sensitivity of this performance to the flux values used as input to the cFIR models. It is then imagined that in operational conditions a simple model is used for predicting the flux at the supply point for the subsequent hours. The accuracy of forecasts obtained with the autoregressive model (15) is discussed in (Pinson et al. 2007). Such flux predictions are then used to feed cFIR models. The impact on the performance of the resulting temperature predictions is then discussed.


As was the case in section 4.3, predictions are here also unbiased, even though bias values are slightly larger for supple cFIR models. Figure 4 summarizes the evaluation of their prediction accuracy with the MAE and RMSE error measures calculated as a function of the look-ahead time. It can then be directly compared to the results of Figure 3 in order to quantify the loss in forecast accuracy owing to the use of flux values that are not the true values for feeding the cFIR models.

The range of variations of both MAE and RMSE criteria appears to be similar to that observed in Figure 3, i.e. in the case for which flux measurements input cFIR models. The increase in error measures as the lead time progresses is slightly sharper for the case of critical points 1 and 2, but not for the third point. The MAE averaged over the forecast length goes from 0.318 and 0.590°C depending on the cFIR models and the critical point considered. Because of the smooth nonlinear variations of the transfer function of the network (if seen as a function of the flux at the supply point), the errors in flux predictions are actually dampened when passed through the cFIR models. Obviously, using more advanced methods for flux prediction (such as that described in section 2.2.3) should make it possible to have more accurate temperature forecasts, but the impact on the gain in accuracy would be limited. Finally, note that while the loss in prediction accuracy is much more significant for the rigid than for the supple cFIR models. The rigid cFIR models still appear to be globally superior, though the difference in prediction accuracy between the two types of cFIR model is now very low.


4.5 Comparison with the state-of-the-art statistical approach

In a last part of this study, focus is given to comparing the forecasting methodology introduced with the state-of-the-art statistical approach. This approach is based on the lagged transfer function of model (1). The lag τ in this model is chosen as that which maximizes the correlation between the time series $\{y_t\}$ and $\{u_{t-\tau}\}$. A single time delay is considered over the whole dataset. The chosen lag for each of the critical points, as well as the corresponding correlation values are collated in Table 2. Significant differences in these correlation values can be noticed. Note that observing a higher correlation does not mean that one should expect a higher forecast performance of the lagged transfer function approach.

The model coefficients are adaptively estimated with a RLS method with exponential



(a) MAE as a function of the look-ahead time.



(b) RMSE as a function of the look-ahead time.

Figure 4. Out-of-sample evaluation of the forecast performance of both the rigid and supplé cFIR models for the three critical points. Flux predictions are used as input to the cFIR models. The forecast performance is evaluated with MAE and RMSE error measures as a function of the look-ahead time up to 12 hours ahead.

Table 2

Characteristics of the lagged autoregressive models used for temperature predictions. The lag of each model is that which maximizes the correlation between the lagged values of the water temperature at the supply point and the temperature measurement at each critical point. The forgetting factor is obtained from a one-fold cross-validation procedure.

critical point	lag [h]	correlation	λ
1	2	0.421	0.988
2	1	0.814	0.994
3	2	0.742	0.98

forgetting. The optimal value of the forgetting factor is obtained from a one-fold cross-validation procedure, similar to that employed for optimal tuning of the cFIR models or of the autoregressive model used for flux prediction. In order to have a fair comparison between the cFIR models and the lagged transfer function models, the second year of the data is also used here as a validation set, and the remainder of the dataset for out-of-sample evaluation of the forecast performance. The autoregressive and cFIR-based approaches are evaluated on the same period and same data. The optimal forgetting factors are also given in Table 2. The level of forgetting is again more pronounced for critical point 3.

The full evaluation of the accuracy of the multi-step ahead predictions for the 3 critical points is shown and discussed in (Pinson et al. 2007). Here, focus is given to the comparison with the prediction performance of cFIR models. For appraising the reduction in prediction error when going from the state-of-the-art approach to that based on cFIR models, the improvement with respect to the RMSE criterion is calculated for each critical point and for each cFIR type. This improvement is given by the decrease in RMSE divided by the RMSE for the lagged transfer function models. The improvement values are collated in Table 3 as a function of the look-ahead time.

Improvement values increase as the lead time progresses, indicating a sharper increase in prediction error for the lagged transfer function model. For 1-hour ahead forecasts, some of the improvement values are negative: for such a short horizon the use of the state-of-the-art approach may be beneficial owing to the autoregressive component that makes it possible to better capture the short-term persistent nature of temperature variations. However, improvement values reach a very high level rapidly, with a reduction in RMSE up to 53% for 12 hours ahead predictions for critical point 2. The improvement is slightly less for the two other critical points, but the forecasting method based on cFIR models clearly outperforms the state-of-the-art approach based on the lagged transfer function model.

5 Concluding remarks

Models for the prediction of the temperature at critical points of district heating systems are paramount for heat suppliers to make optimal decisions on the water temperature at the supply point. This is because the decision-making methodologies are based on

Table 3

Improvements with respect to the RMSE criterion achieved by the genuine predictions from cFIR models in comparison with forecasts obtained from the state-of-the-art approach. Improvements are given for each critical point and both types of cFIR model, and are expressed in %.

(a) rigid cFIR model

horizon [h]	critical point 1	critical point 2	critical point 3
1	-5.5146	26.9938	5.0830
2	10.5332	36.7955	14.1353
3	17.9085	39.8307	16.1290
4	21.5780	40.3498	17.1729
5	25.3272	41.0283	18.5601
6	28.8176	42.7063	20.5886
7	31.8774	44.6379	22.9004
8	34.8676	46.5684	25.2898
9	37.7255	48.5652	27.6495
10	40.6136	50.2731	29.7891
11	43.1811	51.7465	31.4742
12	45.5966	53.0583	32.8710

(b) supple cFIR model

horizon [h]	critical point 1	critical point 2	critical point 3
1	-12.6672	20.6373	-8.9477
2	3.6866	30.6035	5.2472
3	13.2943	34.7397	11.4181
4	18.3019	35.9590	14.9800
5	23.0563	37.2037	18.1995
6	27.1536	39.3316	21.3865
7	30.5186	41.6330	24.4759
8	33.6480	43.8192	27.3243
9	36.5062	45.8661	29.8913
10	39.2566	47.4513	32.0293
11	41.7219	48.7801	33.5817
12	44.0663	50.1574	34.7064

model predictive control. It is thus expected that improvements in the accuracy of temperature forecasts at critical points will significantly improve the control decisions on supply temperature, and thus lower the operational production costs of heat suppliers. The forecasting methodology introduced in the present paper contributes to achieving a higher accuracy for such temperature forecasts. An evaluation of the performance of this methodology has been carried out by considering the case of the Roskilde district heating system over a period of several years.

It has been explained that the methodology proposed makes it possible to account for varying time delays in the distribution network within the statistical model itself. Indeed, the described cFIR model belongs to the family of FIR models, but the model coefficients are replaced by smooth coefficient functions. It is obvious to choose the flux at the supply point as a variable that conditions the transfer function of the network. In contrast, the way to account for the social behaviour of the consumers is not straightforward. This is why it has been chosen to evaluate whether this should be accounted for within the cFIR model, i.e. by considering the time of the day as a second variable conditioning the cFIR model, or by keeping it outside the cFIR model by just allowing the offset term to have diurnal variations. The former alternative is computationally less expensive (owing to the lower number of coefficients to be estimated) and actually leads to higher forecast accuracy for 2 out of the 3 critical points considered. Whatever the chosen approach, the prediction accuracy of the proposed methodology is dramatically higher than that of the state-of-the-art statistical approaches to the present problem.

The proposed class of models may be used for simulation or prediction purposes for other types of problems that involve nonlinearly varying time delays in a transfer function. This could be the case in e.g. signal processing or network modelling. In addition, the described methods for the estimation of the coefficient functions in cFIR models can be improved in the future, both for offline and online applications. Mainly, the issue of optimally tuning the regularization parameter at each fitting point should be further investigated. This may involve the use of numerical methods, but one may also envisage searching for some analytical solutions similar to that existing for the choice of optimal regularization parameters for least squares estimators in linear models, see e.g. (Hoerl et al. 1975).

It would be interesting to study the alternative possibility of considering cFIR models with the load (instead of the flux) variable as that which conditions the transfer function of the district heating system. The main interest of such an approach would then be that forecasts of the temperature at critical points could not be affected by the interdependence of the flux and temperature variables at the supply point. Broader perspectives of this work include the use of the described forecasting methodology within the model predictive control-based methods employed today for decision-making in district heating systems. This will allow us to quantify the operational benefits of a higher accuracy of temperature forecasts.

Acknowledgments

The methods and results in the present report have been generated as part of the Centre for Model-Based Control (CMBC) project, partly supported by the Danish Agency for Science Technology and Innovation (under the contract number 07-000795), which is hereby greatly acknowledged. The authors also gratefully acknowledge Roskilde Forsyning, a company supplying district heating, electricity and power to the citizens of Roskilde Municipality, for providing the data used as input. Finally, three anonymous reviewers are acknowledged for their valuable comments and suggestions.

References

- Arvatson L. Stochastic modeling and operational optimization in district heating systems. Ph.D. dissertation, Lund Institute of Technology, Mathematical Statistics; 2001.
- Bishop CM. Training with noise is equivalent to Tikhonov regularization. *Neural Computation* 1995;7; 108-116.
- Cleveland WS, Devlin SJ. Locally weighted regression: an approach to regression analysis by local fitting. *Journal of the American Statistical Association* 1988;83; 596-610.
- Dotzauer E. Simple model for prediction of loads in district-heating systems. *Applied Energy* 2002;73; 277-284.
- Golub GH, Hansen PC, O'Leary DP. Tikhonov regularization and total least squares. *SIAM Journal of Matrix Analysis and Applications* 2000;21; 185-194.
- Hastie T, Tibshirani R. *Generalized additive models*. London: Chapman & Hall/CRC; 1990.
- Hubing NE, Alexander ST. Statistical analysis of initialization methods for RLS adaptive filters. *IEEE Transactions on Signal Processing* 1991;39; 1793-1804.
- Ismail MY, Principe JC. Equivalence between RLS algorithms and the Ridge regression technique. in: *Proc. 13th Asilomar Conf. Signals, Systems and Computers (IEEE)* 1997; 1083-1087.
- Johansen TA. On Tikhonov regularization, bias and variance in nonlinear system identification. *Automatica* 1997;33; 441-446.
- Hoerl AE, Kennard RW, Baldwin KF. Ridge regression : some simulations. *Communications in Statistics* 1975;4; 105-123.
- Leung CS, Young GH, Sum J, Wing-kay K. On the regularization of forgetting recursive least squares. *IEEE Transactions on Neural Networks* 1999;10; 1482-1486.
- Ljung L, Söderström T. *Theory and Practice of Recursive Identification*. Cambridge: MIT Press; 1983.
- Madsen H, Sejling K, Søggaard HT, Palsson OP. On flow and supply temperature control in district heating systems. *Heat Recovery Systems & CHP* 1994;14;613-620.
- Nielsen HAa. Parametric and non-parametric system modeling. Ph.D. dissertation, Technical University of Denmark, Informatics and Mathematical modeling; 2000.
- Nielsen HAa, Madsen H. Predicting the heat consumption in district heating systems us-

- ing meteorological forecasts. Technical report, Informatics and Mathematical modeling, Technical University of Denmark, ENS.J.Nr. 1323/98-025; 2000.
- Nielsen HAa, Nielsen TS, Joensen AK, Madsen H, Holst J. Tracking time-varying coefficient functions. *International Journal of Adaptive Control and Signal Processing* 2000;14;813-828.
- Nielsen TS. Online prediction and control in nonlinear stochastic systems. Ph.D. dissertation, Technical University of Denmark, Informatics and Mathematical modeling; 2002.
- Pinson P, Nielsen TS, Nielsen HAa, Poulsen NK, Madsen H. Temperature prediction in district heating systems with cFIR models. Technical report, Technical University of Denmark, Informatics and Mathematical modeling, IMM-2007-10; 2007.
- Sandou G, Font S, Tebbani S, Huret A, Mondon C. Global modeling and simulation of a district heating network. in: *Proc. of the 8th International Symposium on District Heating and Cooling*, Espoo, Finland; 2004.
- Sandou G, Font S, Tebbani S, Huret A, Mondon C. Predictive control of a complex district heating system. in: *Proc. of the 44th IEEE Conference on Decision and Control*, Seville, Spain; 2005.
- Sjöberg J, Ljung L. Overtraining, regularization and searching for a minimum, with application to neural networks. *International Journal of Control* 1995;62; 1391-1407.
- Stone M. Cross-validation and assessment of statistical predictions (with discussion). *Journal of the Royal Statistical Society B* 1974;36; 111-147.
- Søgaard HT. Stochastic systems with embedded parameter variations - Applications to district heating. Ph.D. dissertation, Technical University of Denmark, Institute of Mathematical Statistics and Operations Research; 1993.
- Tikhonov AN, Arsenin VY. *Solutions of Ill-posed Problems*. Washington DC: Wiscon; 1977.
- Wang S-G, Chow S-C. A note on adaptive generalized ridge regression estimator. *Statistics & Probability Letters* 1989;10; 17-21.