AMSI SUMMER RESEARCH SCHOLARSHIPS 2025–26

Get a taste for Research this Summer

Optimal Control of Air Conditioning with Thermal Storage

Rong Xu

Supervised by Peter Pudney

Adelaide University

Jane Street

AMSI AUSTRALIAN MATHEMATICAL SCIENCES INSTITUTE

As part of the AMSI Summer Research program, I conducted research in collaboration with my supervisor, Professor Peter Pudney, and co-authored a paper based on the findings. This paper outlines the primary methodology, results, and contributions of the research. The submitted paper is attached.

Highlights

Using Pontryagin's principle to determine the structure of an optimal control strategy for heat-pump air conditioning with thermal storage

Peter Pudney,Rong Xu

- The optimal control of a heat-pump air-conditioning system with thermal storage must balance heat-pump efficiency against time-of-use electricity costs
- Pontryagin's principle can be used to derive control modes and switching times for the optimal control of the heat pump, and for heat transfer to or from the conditioned space
- The optimal control operates the heat pump when the electricity price is low and the heat-pump efficiency is high
- The optimal control incorporates pre-heating or pre-cooling the conditioned space.

Using Pontryagin's principle to determine the structure of an optimal control strategy for heat-pump air conditioning with thermal storage

Peter Pudney[a,*], Rong Xu[a]

[a] Adelaide Unversity, South Australia, Australia

ARTICLE INFO

Keywords:
optimal control
Pontryagin's principle
air conditioning
heat pump
thermal storage
demand flexibility
time-of-use tariff

ABSTRACT

Heat-pump air-conditioning systems are efficient but often contribute to peak demand on electricity grids. By incorporating thermal storage, the operation of the heat pump can be separated in time from the heating or cooling of the conditioned space, so that the electrical load occurs when electricity is clean or cheap.

We use an energy-based model of a heat-pump air-conditioning system to formulate and solve the problem of how to control the heat pump and the transfer of heat between the store and the conditioned space so that the operating cost is minimised. We take into account the efficiency of the heat pump, which varies in time with ambient temperature. We also consider thermal losses from the store and the conditioned space, target temperatures in the conditioned space, and a time-of-use electricity tariff.

Using Pontryagin's principle to derive necessary conditions for the optimal control gives important insights into the structure of the optimal control strategy. The heat pump operates at full power when the price of electricity is low and the efficiency of the heat pump is high, and at partial power when the store is empty or full. Heat transfer occurs at full rate when active heating or cooling is required, or at partial rate when maintaining the target temperature in the conditioned space. The necessary conditions can be used to determine the sequence of controls and precise switching times.

1. Introduction

As the world becomes more reliant on electricity generated from variable renewable energy sources such as wind and solar, it is becoming increasingly important for households to align their energy use with the available supply of renewable energy. In many climates, air conditioning accounts for a significant proportion of residential electricity use, with peak demand for heating or cooling in the evening as solar generation falls to zero and the price of electricity from the grid is highest.

One way to help households match their air-conditioning load to the available supply of renewable energy is to incorporate thermal storage into the air-conditioning system so that the production of heat or cooling energy can be done when the price of electricity is low, independently of the required heat transfer to or from the conditioned space.

We consider the use of a heat-pump air-conditioning system with thermal storage for heating and cooling. We start by modelling the heat flows in the system when heating, including heat losses from the store and conditioned space, heat flow from the heat-pump to the store, heat flow from the store to the conditioned space, and temperatures in the store and conditioned space (Section 3). We then formulate an optimal control problem—manage the power to the heat-pump and the heat flow from the store to the conditioned space to minimise the cost of satisfying given temperature constraints in the conditioned space (Section 4). The necessary conditions for an optimal control can be found using Pontragin's principle (Section 5). Although Pontryagin's principle is sometimes regarded as being difficult to use, it often leads to insights into the structure of an optimal control strategy that can be used to implement simple controls. We show that an optimal control has:

- three simple modes for heating the store: power off, maximum power, and intermediate power used when the store is empty or full
- three simple modes for transferring heat from the store to the conditioned space: zero transfer, maximum heat transfer, and intermediate heat transfer to maintain the minimum desired temperature.

In Section 6 we show how these control modes can be combined and sequenced to meet heating requirements in different scenarios. Section 7 describes the changes to the model and necessary conditions when the system is used to cool the conditioned space, and gives an example. Section 8 summarises the rules for constructing optimal control sequences.

2. Related work

Our aim is to use Pontryagin's principle to derive the necessary conditions that define the structure of an optimal control strategy for an air-conditioning system with storage.

Several authors consider the problem of controlling air-conditioning systems incorporating thermal storage, but use heuristic control methods (Arteconi, Hewitt and Polonara, 2013; Dar, Sartori, Georges and Novakovic, 2014; Schibuola, Scarpa and Tambani, 2015; Jarvinen, Goldsworthy, Pudney, White, Cirocco and Bruno, 2023) or 'black box' optimisation methods including Sequential Quadratic Programming

*Corresponding author

peter.pudney@adelaide.edu.au (P. Pudney); rong.xu@adelaide.edu.au (R. Xu)

ORCID(s): 0000-0001-5596-5473 (P. Pudney); 0009-0002-7303-2102 (R. Xu)

(Kintner-Meyer and Emery, 1995), Dynamic Programming (Henze, Dodier and Krarti, 1997)), Mixed Integer Linear Programming (Renaldi, Kiprakis and Friedrich, 2017; Risbeck, Maravelias, Rawlings and Turney, 2017), genetic algorithms (Liu, Zhu, Zhao, Ren, Groll and Cai, 2017), collocation solvers (Salameh, Schalbart and Peuportier, 2021) or neural networks (Muñoz, Garcia-Hernandez, Ruelas, Palomares-Ruiz and Álvarez Macías, 2022). These methods do not give direct insight into the structure of an optimal control strategy.

Zhang, Stockar and Canova (2016) use Pontryagin's principle to find a control strategy for car air conditioning, but do not consider the thermal dynamics of the cabin and do not consider thermal storage.

Zargari, Levron and Belikov (2019) and Chowdhury, Ofir, Zargari, Baimel, Belikov and Levron (2021) use Pontryagin's principle and dynamic programming to show that the optimal control of an energy system with storage is equivalent to finding the shortest path through a tube defined by bounds between the load profile and the load profile plus storage capacity. Gates and Westcott (1996) showed a similar relationship for solar racing cars. Chowdhury et al. (2021) consider a grid-connected storage device with losses that depend on the control. In our case, losses depend on the state of the store.

Cirocco, Pudney, Riahi, Liddle, Semsarilar, Hudson and Bruno (2022) use Pontryagin's principle to determine control strategies that minimise the cost of operating an industrial cooling system with ice storage, PV electricity generation and real-time electricity prices. The optimal strategy has four operating modes, depending on the import and export prices for electricity:

- maximum charging of the thermal store when the import and export prices are low
- limited charging of the thermal store when the import price and export prices are low but there is a cost associated with the demand for electricity
- use of solar energy to charge the store when import price is high but the export price is low
- discharging the store when the import and export prices are high.

The critical price that determines the operating mode depends on the predicted future loads, PV generation and electricity prices. Our problem is different—we consider simpler cyclic time-of-use pricing, but allow the temperature of the conditioned space to vary within limits.

Fleming, Barbour and Urquhart (2025) use Pontryagin's principle to find a control strategy for a home heating system with a heat pump and thermal energy storage, to minimise the cost of running the heat pump with a dynamic electricity price. They assume the heating demand for the conditioned space is a known function of time, so they do not consider pre-heating of the space as part of the control. They find solutions using Dynamic Programming on 30-minute intervals, and using Pontryagin's principle. The optimal control is a function of a single co-state variable, which they estimate to keep the state-of-charge stable over long time intervals. Their store had a capacity of 100 kWh, which was sufficient to mean that it never became full.

3. System model

Figure 1 shows a schematic diagram of an air-conditioning system with storage. Our model, and the diagram, depict heating. Cooling would reverses the heat flows. For the rest of this paper we focus on heating.

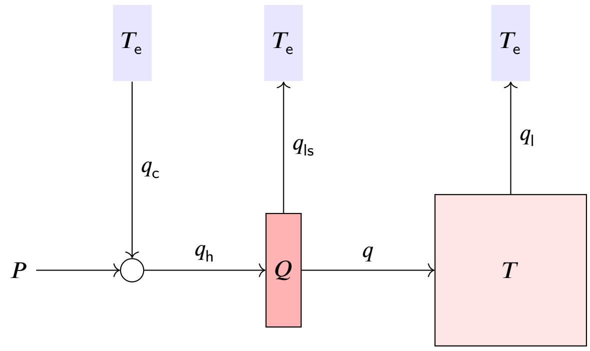

Figure 1: Schematic diagram of an air-conditioning system with thermal storage

In this diagram:

- T_e is the external temperature
- P is the time-varying electrical power to the heat pump
- q_c is the rate of thermal energy flow into the heat pump
- $q_h = \eta P$ is the rate of thermal energy flow from the heat pump to the thermal store
- $Q \in [0, Q_{max}]$ is the stored thermal energy
- q is the rate of thermal energy flow from the thermal store to the conditioned space
- T is the temperature in the conditioned space
- q_{ls} is the rate of thermal energy flow from the store to the atmosphere.
- q_l is the rate of thermal energy flow from the conditioned space to the atmosphere.

All of these values will vary with time. External temperature T_e and the price of electricity r are given functions of time. The desired minimum temperature T_{min} in the conditioned space is also a given function of time.

We ignore solar heat gain and heating from occupants and appliances, but these could be modelled as additional time-varying heat flows into the conditioned space.

In practice, the Coefficient of Performance (CoP) η of the heat pump will vary with the external temperature T_e, which will vary with time. We will approximate the CoP when heating water by

$$\eta(t) = \frac{aT_h}{T_h - T_e(t)}$$

where T_h is the temperature of the heated water in kelvin, T_e is the external air temperature, which varies with time, and a is a constant (Lovegrove, Alexander, Bader, Edwards, Lord, Mojiri, Rutovitz, Saddler, Stanley, Urkalan and Watt, 2019). In our examples we assume $T_h = 80\,°C$ and $a = 0.7$.

The system is controlled by varying the electrical power P to the heat pump and the heat transfer rate q from the store to the conditioned space.

4. Problem formulation

We wish to determine a heat-pump power profile P and a heat transfer rate q that will keep the conditioned space at or above the required minimum temperature T_{min} with minimum total cost. The conditioned space is not likely to overheat if we are minimising the cost of heating, so we won't specify a maximum temperature T_{max}.

Our aim is to understand the structure of an optimal control strategy rather than accurately model the thermal performance of the system, so our thermal models are deliberately simple. We assume that the temperature in the conditioned space will evolve according to the differential equation

$$\dot{T} = \frac{1}{C}(q - q_l) \tag{1}$$

where C is the heat capacity of the conditioned space in $J\,K^{-1}$ and q is the controlled rate of heat flow from the store to the conditioned space. The heat loss from the conditioned space is

$$q_l = k(T - T_e)$$

where k is a heat loss factor for the conditioned space, in $W\,K^{-1}$, which depends on the size of the conditioned space and the conductivity of the envelope.

The energy Q in the thermal store will evolve according to the differential equation

$$\dot{Q} = q_h - q_{ls} - q$$

with constraint $0 \le Q \le Q_{max}$. When heating, the heat flow into the store is

$$q_h = \eta P$$

where P is the electrical power to the heat pump and $\eta + 1$ is the coefficient of performance (CoP) of the heat pump, which varies with the external temperature.

We also assume that heat is stored as sensible heat in water. If the minimum usable mean temperature of the stored water is T_{s0} and we add heat Q then the new mean temperature of the store will be

$$T_s = T_{s0} + \frac{Q}{C_s}$$

where $C_s = 4200 m_s$ is the heat capacity of the water in $J\,K^{-1}$ and where m_s is the mass of water in the tank.

The heat loss from the store is

$$q_{ls} = k_s\left(T_s - T_e\right)$$

where k_s is a heat loss factor for the store, in $W\,K^{-1}$. The change in stored energy is therefore given by

$$\dot{Q} = \eta P - k_s\left(T_{s0} + \frac{Q}{C_s} - T_e\right) - q. \tag{2}$$

We want to minimise the total cost of heating the conditioned space

$$J = \int_0^{t_f} P(t) r(t)\, dt$$

where $r(t)$ is the cost of electricity at time t. If we want the optimal control strategy to be repeatable for a sequence of time intervals, such as days, then we also require $Q(t_f) = Q_0$ and $T(t_f) = T_0$.

5. Pontryagin analysis

We will find the optimal control using the direct adjoining approach described by (Hartl, Sethi and Vickson, 1995, Section 4). The state of the system is the temperature T of the conditioned space and the energy in the store Q. The controls are the electrical power P to the heat pump and the rate of heat transfer q from the store to the conditioned space. The constraints on the system are

$$\begin{aligned} P &\in [0, P_{max}] \\ q &\in [0, q_{max}] \\ Q &\in [0, Q_{max}] \\ T &\ge T_{min} \end{aligned}$$

where P_{max}, q_{max} and Q_{max} are constants and T_{min} is a given function of time.

5.1. Hamiltonian and Lagrangian

The Hamiltonian for our system is

$$\mathcal{H} = -\lambda_0 P r + \lambda_1\left(\eta P - k_s\left(T_{s0} + \frac{Q}{C_s} - T_e\right) - q\right) + \lambda_2 \frac{1}{C}(q - k(T - T_e)) \tag{3}$$

where $\lambda_0 \ge 0$, λ_1 and λ_2 are Lagrange multipliers, and where λ_0 is a constant and λ_1, λ_2, P, q, Q, T, T_e and r are functions of time. The Lagrangian is

$$\mathcal{L} = \mathcal{H} + \mu_1 P + \mu_2(P_{max} - P) + \mu_3 q + \mu_4(q_{max} - q) + \nu_1 Q + \nu_2(Q_{max} - Q) + \nu_3(T - T_{min}) \tag{4}$$

where μ_1, μ_2, μ_3, μ_4, ν_1, ν_2 and ν_3 are Lagrange multiplier functions of time.

5.2. Necessary conditions for an optimal control

The Pontryagin principle says that the optimal control will maximise the Hamiltonian. The Hamiltonian can be rewritten as

$$\mathcal{H} = \left(\lambda_1 \eta - \lambda_0 r\right) P + \left(\lambda_2 \frac{1}{C} - \lambda_1\right) q + \cdots \tag{5}$$

where the omitted terms do not depend on the control variables P or q. The optimal control depends on the sign of $\lambda c - \lambda_0 r$ and $\lambda_2 \frac{1}{C} - \lambda_1$. For power P:

- if $\lambda_1 \eta - \lambda_0 r < 0$ then $\mathcal{H}$ is maximised by setting $P = 0$
- if $\lambda_1 \eta - \lambda_0 r = 0$ then P can take on any value in $[0, P_{\max}]$
- if $\lambda_1 \eta - \lambda_0 r > 0$ then $\mathcal{H}$ is maximised by setting $P = P_{\max}$.

For the heat transfer rate q from the store to the conditioned space:

- if $\lambda_2 - C\lambda_1 < 0$ then $\mathcal{H}$ is maximised by setting $q = 0$
- if $\lambda_2 - C\lambda_1 = 0$ then q can take on any value in $[0, q_{\max}]$
- if $\lambda_2 - C\lambda_1 > 0$ then $\mathcal{H}$ is maximised by setting $q = q_{\max}$.

The Karush-Kuhn-Tucker conditions are

$$\frac{\partial \mathcal{L}}{\partial P} = 0 \implies \lambda_1 \eta - \lambda_0 r + \mu_1 - \mu_2 = 0 \tag{6}$$

$$\frac{\partial \mathcal{L}}{\partial q} = 0 \implies \frac{\lambda_2}{C} - \lambda_1 + \mu_3 - \mu_4 = 0. \tag{7}$$

The adjoint variables λ_1 and λ_2 evolve according to the differential equations

$$\dot{\lambda}_1 = -\frac{\partial \mathcal{L}}{\partial Q} = \frac{k_s}{C_s}\lambda_1 - \nu_1 + \nu_2 \tag{8}$$

$$\dot{\lambda}_2 = -\frac{\partial \mathcal{L}}{\partial T} = \frac{k}{C}\lambda_2 - \nu_3. \tag{9}$$

The complementary slackness conditions are

$$\begin{aligned} \mu_1 &\geq 0 & \mu_1 P &= 0 \\ \mu_2 &\geq 0 & \mu_2(P_{\max} - P) &= 0 \\ \mu_3 &\geq 0 & \mu_3 q &= 0 \\ \mu_4 &\geq 0 & \mu_4(q_{\max} - q) &= 0 \\ \nu_1 &\geq 0 & \nu_1 Q &= 0 \\ \nu_2 &\geq 0 & \nu_2(Q_{\max} - Q) &= 0 \\ \nu_3 &\geq 0 & \nu_3(T - T_{\min}) &= 0. \end{aligned}$$

At the final time t_f we have

$$\lambda(t_f^-) = \beta b_x^*[t_f] + \gamma h_x^*[t_f] \tag{10}$$

where $\lambda = (\lambda_1, \lambda_2)$, $\beta = (\beta_1, \beta_2)$ and $\gamma = (\gamma_1, \gamma_2, \gamma_3)$ and where

$$b^*(t) = \begin{pmatrix} Q(t) - Q_0 \\ T(t) - T_0 \end{pmatrix} \implies b_x^*[t_f] = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

and

$$h^*(t) = \begin{pmatrix} Q \\ Q_{\max} - Q \\ T - T_{\min} \end{pmatrix} \implies h_x^*[t_f] = \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & 1 \end{pmatrix} \tag{11}$$

and so $(\lambda_1, \lambda_2) = (\beta_1 + \gamma_1 - \gamma_2, \beta_2 + \gamma_3)$.

Our problem is a 'normal' problem (the linearised version is completely controllable and satisfies the Slater constraint qualifications) so we can set $\lambda_0 = 1$. The conditions that determine the optimal control become:

- if $\lambda_1 \eta - r < 0$ then $\mathcal{H}$ is maximised by setting $P = 0$
- if $\lambda_1 \eta - r = 0$ then P can take on any value in $[0, P_{\max}]$
- if $\lambda_1 \eta - r > 0$ then $\mathcal{H}$ is maximised by setting $P = P_{\max}$.
- if $\lambda_2 - C\lambda_1 < 0$ then $\mathcal{H}$ is maximised by setting $q = 0$
- if $\lambda_2 - C\lambda_1 = 0$ then q can take on any value in $[0, q_{\max}]$
- if $\lambda_2 - C\lambda_1 > 0$ then $\mathcal{H}$ is maximised by setting $q = q_{\max}$.

The inequalities give regular control modes; the equalities give singular control modes.

Hartl et al. (1995, equations (4.12) and (4.14)) give the conditions for jump discontinuities in the adjoint values—the adjoint variable $\lambda = (\lambda_1, \lambda_2)$ can have a jump discontinuity at time τ if

$$\begin{aligned} \lambda(\tau^-) &= \lambda(\tau^+) + \zeta(\tau) h_x^*[\tau] \\ &= \lambda(\tau^+) + (\zeta_1, \zeta_2, \zeta_3)\begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & 1 \end{pmatrix} \\ \mathcal{H}^*[\tau^-] &= \mathcal{H}^*[\tau^+] - \zeta(\tau) h_t^*[\tau] = \mathcal{H}^*[\tau^+] \end{aligned}$$

where $\zeta(\tau) > 0$ and $\zeta(\tau)h^*[\tau] = 0$. The jump discontinuities in λ can be written as

$$\begin{aligned} \lambda_1[\tau^-] &= \lambda_1[\tau^+] + \zeta_1 - \zeta_2 \\ \lambda_2[\tau^-] &= \lambda_2[\tau^+] + \zeta_3. \end{aligned}$$

and so

- λ_1 can jump down if $Q = 0$
- λ_1 can jump up if $Q = Q_{\max}$
- λ_2 can jump down if $T = T_{\min}$.

5.3. Singular control modes

There are two singular control modes:

- $\lambda_1 \eta = r$, which allows partial power $P \in (0, P_{\max})$
- $\lambda_2 = C\lambda_1$, which allows partial heat transfer $q \in (0, q_{\max})$.

Further analysis of these singular modes gives additional necessary conditions for an optimal solution.

5.3.1. First singular control mode: partial power

If $\lambda_1\eta = r$ on any interior interval $[\tau_0, \tau_1]$ with $\tau_0 < \tau_1$, and the price r is constant on that interval then differentiating $\lambda_1\eta = r$ with respect to t gives

$$\frac{d}{dt}\left[\lambda_1\eta\right] = \eta\dot{\lambda}_1 + \lambda_1\dot{\eta} = 0 \implies \dot{\lambda}_1 = -\frac{\lambda_1\dot{\eta}}{\eta}.$$

From (8) we also have

$$\dot{\lambda}_1 = \frac{k_s}{C_s}\lambda_1 - \nu_1 + \nu_2,$$

so on a singular interval these two expressions for $\dot{\lambda}_1$ must coincide, and hence

$$\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) - \nu_1 + \nu_2 = 0. \tag{12}$$

Here ν_1 and ν_2 must satisfy the complementary slackness conditions

$$\nu_1 Q = 0, \quad \nu_2(Q_{\max} - Q) = 0, \quad \nu_1, \nu_2 \geq 0.$$

Hence

$$\nu_1 > 0 \implies Q = 0, \quad \nu_2 > 0 \implies Q = Q_{\max}.$$

The factor $k_s/C_s + \dot{\eta}/\eta$ will be positive unless the CoP η is decreasing at a rate $\dot{\eta} < -\eta k_s/C_s$. Our examples show that this is possible in practice.

Now consider cases where (12) is satisfied, if $k_s > 0$:

- If $0 < Q < Q_{\max}$ then $\nu_1 = \nu_2 = 0$ and (12) reduces to

  $$\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) = 0.$$

  This can occur if $\lambda_1 = 0 \implies r = 0$, i.e. when the price is zero, or if $k_s/C_s + \dot{\eta}/\eta = 0$, which requires η to be reducing at the precise rate $\dot{\eta} = -\eta k_s/C_s$, which is unlikely.

- If the store is empty ($Q = 0$) then $\nu_2 = 0$ and $\nu_1 \geq 0$, and (12) gives

  $$\nu_1 = \lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) \geq 0.$$

  If the price $r >= 0$ then $\lambda_1 = r/\eta > 0$ so we must have

  $$\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta} > 0. \tag{13}$$

  That is, you cannot maintain partial power with $Q = 0$ if the CoP η is reducing too rapidly.

- If the store is full ($Q = Q_{\max}$) then $\nu_1 = 0$ and $\nu_2 \geq 0$, and (12) gives

  $$\nu_2 = -\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) \geq 0.$$

  For $r > 0$ we have $\lambda_1 = r/\eta > 0$, so we require

  $$\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta} < 0. \tag{14}$$

  That is, the CoP η must be decreasing at a sufficient rate.

5.3.2. Second singular control mode: partial heat transfer

If $\lambda_2 = C\lambda_1$ on any interior interval $[\tau_0, \tau_1]$ with $\tau_0 < \tau_1$ then we have $\dot{\lambda}_2 = C\dot{\lambda}_1$ and so, from (8) and (9),

$$C\left(\frac{k_s}{C_s}\lambda_1 - \nu_1 + \nu_2\right) = k\lambda_1 - \nu_3$$
$$\implies \left(\frac{k}{C} - \frac{k_s}{C_s}\right)\lambda_1 + \nu_1 - \nu_2 - \frac{\nu_3}{C} = 0.$$

We assume that $k_s/C_s < k/C$ in practice. The following conditions are necessary for this singular mode to apply:

- if $\nu_1 = \nu_2 = \nu_3 = 0$, because $T > T_{\min}$ and $q \neq 0$ and $q \neq Q_{\max}$, then we must have $\lambda_1 = 0$ and hence $\lambda_2 = 0$
- if $\lambda_1 > 0$ then we require one or both of $\nu_2 > 0$ or $\nu_3 > 0$, so the store must be full or $T = T_{\min}$
- if $\lambda_1 < 0$ then we require $\nu_1 > 0$, so the store must be empty.

6. Heating examples

In the following examples we assume that every day has the same external temperature profile with a period of 1 day, shown in Figure 2. This is a typical temperature profile for a winter day in South Australia. We will express temperatures in °C, but use kelvin as required when doing calculations.

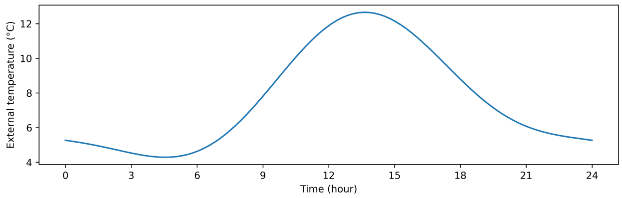

Figure 2: External winter temperature profile

The examples have the same daily price profile, based on residential time-of-use tariffs in South Australia:

$$r(t) = \begin{cases} 0.30\,\$/\text{kWh} & t \in [00{:}00, 06{:}00) \\ 0.50\,\$/\text{kWh} & t \in [06{:}00, 10{:}00) \\ 0.20\,\$/\text{kWh} & t \in [10{:}00, 16{:}00) \\ 0.50\,\$/\text{kWh} & t \in [16{:}00, 00{:}00). \end{cases}$$

For the example problems we require the solution to repeat daily, so $Q(24) = Q(0)$, $T(24) = T(0)$, $\lambda_1(24) = \lambda_1(0)$, and $\lambda_2(24) = \lambda_2(0)$ where the times are in hours.

The maximum electrical power into the heat pump is $P_{max} = 1.5\,\mathrm{kW}$ and the maximum heat transfer rate from the store to the conditioned space is $q_{max} = 4\,\mathrm{kW}$. The energy capacity of the store for a 400 litre tank with a usable temperature difference of 40 °C is $Q_{max} = 67.2\,\mathrm{MJ}$. We assume that the mean temperature of the tank varies between 40 °C and 80 °C, with $T_s = 40 + Q_{max}/C_s$ where $C_s = Q_{max}/40$. The heat loss factor for the store is $k_s = 2\,\mathrm{W\,K^{-1}}$, the heat loss factor for the conditioned space is $k = 100\,\mathrm{W\,K^{-1}}$, and the heat capacity of the conditioned space is $C = 50\,\mathrm{MJ\,K^{-1}}$.

Example 1: Constant temperature

We start with a simple example where we maintain temperature $T(t) = T_{min}(t) = 20$ °C throughout the day. The heat transfer rate required to maintain this constant temperature is

$$q = k(20 - T_e).$$

We assume that it is feasible to maintain the required temperature, so $q \in [0, q_{max}]$.

Over the period of a day we need to generate and store enough heat to meet the heating requirements of the conditioned space plus the losses from the store. The best time to generate this heat will be when the price of electricity is low and the efficiency of the heat-pump is high, during the interval [10:00, 16:00). We expect the optimal solution to fill the store during the day when the price of electricity is lowest. If the stored energy is not sufficient to keep the conditioned space warm overnight then further heating may be required during the off-peak interval [06:00, 10:00). If possible, we should avoid generating heat during the peak periods [06:00, 10:00) and [16:00, 00:00).

Figure 3 shows the optimal solution for this scenario. The top graph shows the energy in the store as a function of time, with the control $P = 0$ indicated by blue lines, the control $P = P_{max}$ indicated by red lines, and the control $P \in (0, P_{max})$ on the boundary $Q = 0$ or $Q = Q_{max}$ indicated by orange lines. The background shading indicates the price of electricity. The bottom graph shows the price profile in grey and the value of $\lambda_1\eta$ using the same colours as in the top graph. The critical points are shown in the table below the graphs. The red values in the table are values after λ_1 has jumped.

To find an optimal solution we must solve the differential equations that govern the energy Q in the store, the temperature T in the conditioned space, and the adjoint variable λ_1 that determines the power mode. The process of finding the optimal solution is easiest to understand if you start at t = 10.

1. We want to get as much energy into the store while the price is low, so we start with Q(10) = 0. Filling the store does not take 6 hours, so we search for the most efficient time in the interval [10, 16) to start filling the store. The time interval [11.95, 16) is the interval $[a, b]$ with $Q(a) = 0$, $Q(b) = Q_{max}$, and $\lambda_1(t)\eta(t) > r(t)$ on $t \in (a, b)$.
2. We want to switch to $P = 0$ when the price increases at $t = 16$. We will have $P = 0$ as long as $\lambda_1\eta < 0.5$. Solving the differential equations for Q gives $Q(24)$ and hence $Q(0)$. The adjoint value λ_1 can jump up at $t = 16$ since the store is full. We choose $\lambda_1(16)$ so that $\lambda_1(0) = \lambda_1(24)$ gives the right behaviour on the next interval $t \in [0, 6)$.
3. We require additional heating on the interval $t \in [0, 6)$. We switch from $P = 0$ to $P = P_{max}$ at a time a when $\lambda_1(a)\eta(a) = 0.3$. The switching time $a = 4.40$ puts enough energy into the store by time $t = 6$ so that the store is empty again at $t = 10$. We know λ_1 at the switching point, so can work back to calculate λ_1 at times $t = 0$, $t = 24$ and $= 16$. The adjoint value λ_1 can jump down at $t = 10$ because we have $Q(10) = 0$.

The optimal solution delays overnight charging until the end of the off-peak price period because this maximises the efficiency of the heat-pump in this interval and minimises the losses from the store.

Example 2: Evening heating

Our second example has a minimum temperature requirement of $T_{min} = 20$ °C between 17:00 and 21:00, and $T_{min} = 16$ °C at all other times.

Figure 4 shows the optimal solution for this example. The line colour in the top two graphs indicates the value of P: blue indicates $P = 0$, red indicates $P = P_{max}$ and orange indicates an intermediate power. The line colour in the bottom two graphs indicates the value of q: blue indicates $q = 0$, red indicates $q = q_{max}$ and orange indicates an intermediate heat transfer. The critical points are shown in the table.

In this example we find a control that meets the necessary conditions by following these steps:

1. Create a temperature profile that follows the boundary $T = T_{min}$ with $q \in (0, q_{max})$ on the interval $[17, 21)$, then apply $q = 0$ on $[21, 24)$, set $T(0) = T(24)$, then apply $q = 0$ until time $a = 11.15$, so that applying $q = q_{max}$ on $[a, 17)$ gives $T(17) = T_{min}$.
2. Set $Q(0) = 2.54$ so that the store becomes empty exactly at $h = 10$. Keep the store empty on $[10, 10.49)$, heat with $P = P_{max}$ during the cheapest price period until $t = 16$, then set $P = 0$ on $[16, 24)$. The time $b = 10.49$ is chosen so that $Q(0) = Q(24)$.
3. If we set $\lambda_1(10) = r(10)/\eta(10)$ and $\lambda_1(a) = r(a)/\eta(a)$ then λ_1 evolves so that $\lambda_1\eta = r$ on $t \in [(0, b)$, $\lambda_1\eta \geq r$ on $t \in [b, 16)$, $\lambda_1\eta < r$ on $t \in [16, 24) \cup [0, 10)$, and λ_1 jumps down at $t = 10$.
4. To calculate λ_2, find a value λ_2 so that $\lambda_2 < C\lambda_1(h)$ until $t = 11.15$. The value of λ_2 can jump down at $t = 17$ to follow $\lambda_2 = C\lambda_1$ on the interval $[17, 21)$, then jump down again at $t = 21$ so that $\lambda_2(24) = \lambda_2(0)$.

Example 3: Home all day

Example 3 has $T_{min} = 20$ °C from 07:00 until 21:00 each day. The optimal solution is shown in Figure 5. In

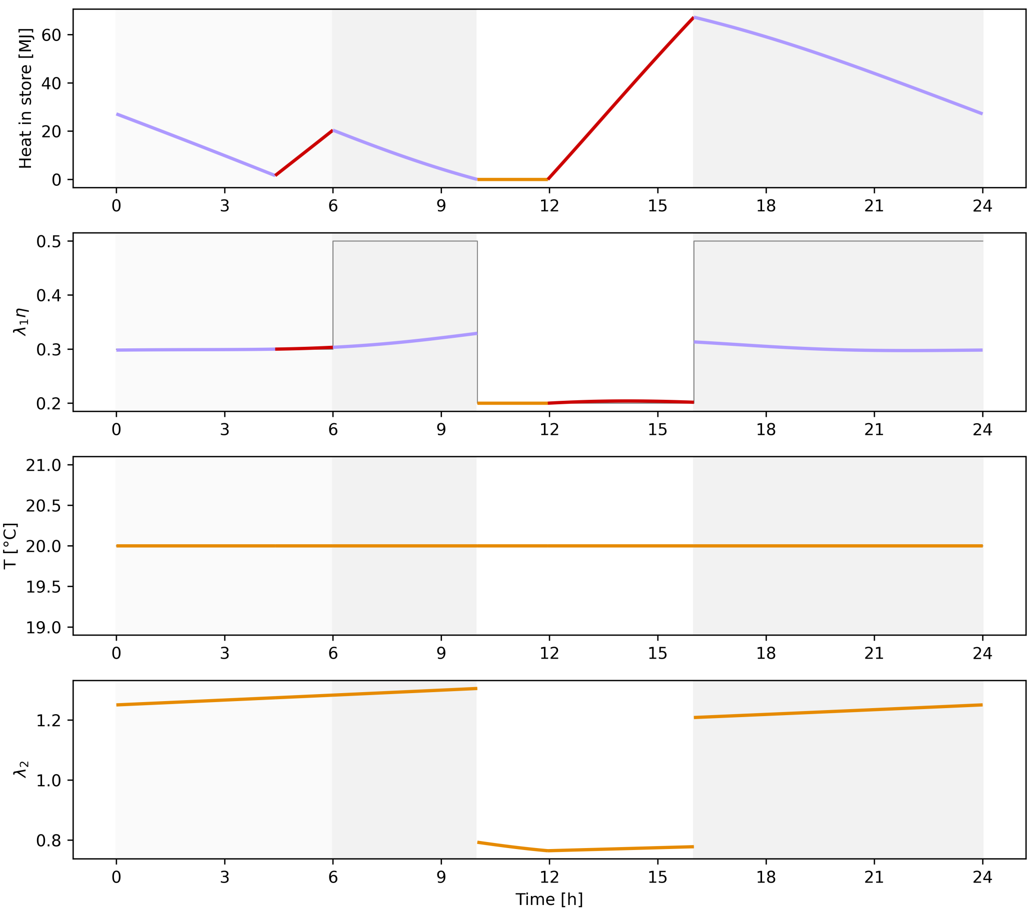

t	Q	$\lambda_1\eta$	$r(t)$	P
0.00	27.14	0.2982	0.50	0
4.40	1.60	0.3000	0.20	$P_{\max}$
6.00	20.41	0.3034	0.50	0
10.00	0.00	0.2000	0.20	$(0, P_{\max})$
11.95	0.00	0.2000	0.20	$P_{\max}$
16.00	67.20	0.3133	0.50	0
24.00	27.14	0.2982	0.50	0

Figure 3: Optimal stored energy and adjoint value profiles for Example 1, with constant temperature

this example, heating the store during the cheapest period only does not generate enough heat, so the optimal solution requires the store to be filled during the cheapest period $t \in [10, 16)$ and additional heating during $t \in [0, 6)$. The critical features of the solution are:

- heat transfer to the conditioned space starts at $t = 3.32$ so that $T(24) = T(0)$
- the store must be empty at $t = 10$ and be completely filled during the interval $t \in [10, 16]$
- $\lambda_1\eta \le r$ on $t \in [0, 4.42]$ and $\lambda_1\eta \ge r$ on $t \in [(4.42, 6)$ so that extra heating can occur overnight
- λ_1 can jump down at $t = 10$ and can jump up at $t = 16$.

Example 4: A leaky house

In this example we consider a house with a larger heat loss coefficient, $k = 200\,\mathrm{W\,K^{-1}}$. The profiles and the optimal solution are shown in Figure 6.

A cyclic solution requires more heat generation than can be achieved by heating continuously during the cheapest period $t \in [10, 16)$ and filling the store overnight. To overcome the storage constraint and avoid heating when the price is high, some heat is transferred from the store to the conditioned space overnight.

7. Cooling

7.1. Model changes

If the system is used for cooling instead of heating then we need to make some changes to our model. We will still use the heat flow directions indicated in Figure 1, but the heat

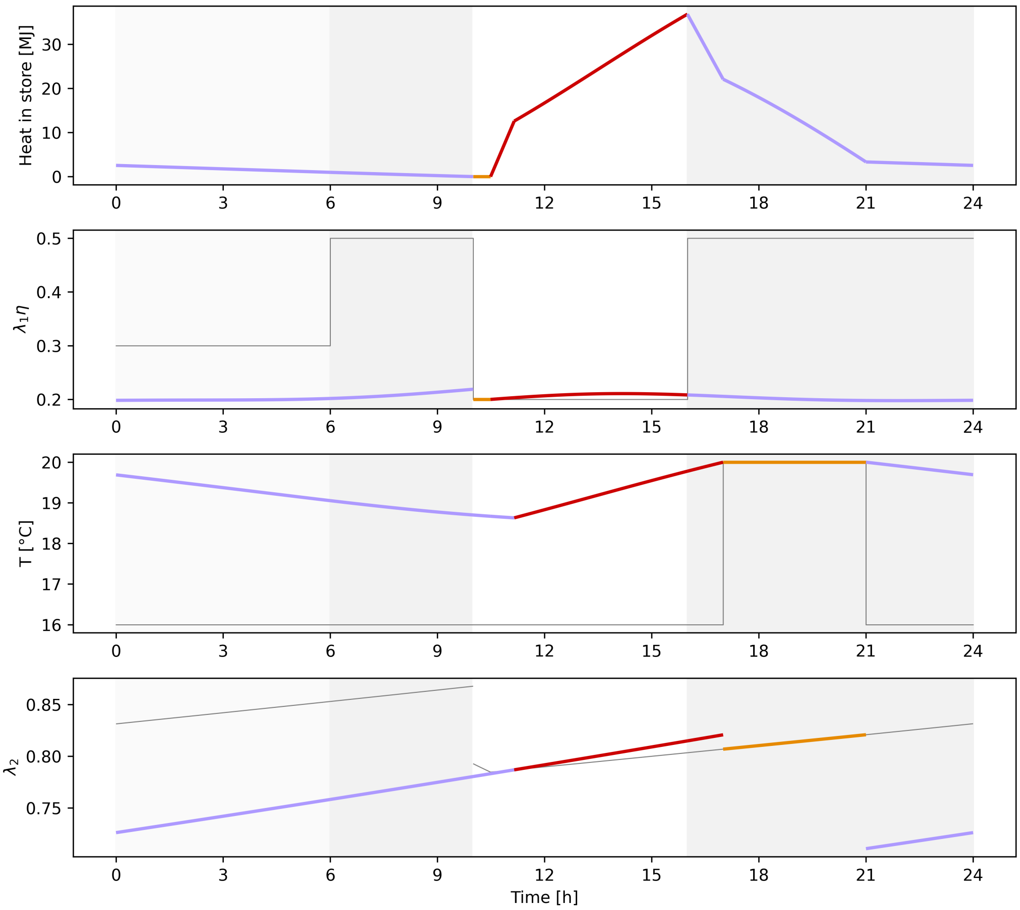

t	Q	$\lambda_1\eta$	T	λ_2	$C\lambda_1$	P	q
0.00	2.54	0.1983	19.69	0.7262	0.8313	0	0
6.00	0.97	0.1998	19.05	0.7583	0.8530	0	0
10.00	0.00	0.2000	18.70	0.7804	0.8677	$(0, P_{max})$	0
10.49	0.00	0.2000	18.66	0.7832	0.7847	P_{max}	0
11.15	12.58	0.2032	18.63	0.7869	0.7869	P_{max}	q_{max}
16.00	36.80	0.2083	19.78	0.8149	0.8034	0	q_{max}
17.00	22.07	0.2057	20.00	0.8069	0.8069	0	$(0, q_{max})$
21.00	3.33	0.1979	20.00	0.7107	0.8208	0	0
24.00	2.54	0.1983	19.69	0.7262	0.8314	0	0

Figure 4: Operation of the store and conditioned space temperature for Example 2

flow values will be negative when heat is flowing against the arrows. In particular, we will have $q_c \leq q_h \leq 0$ when the heat pump is operating, $q \in [-q_{max}, 0]$ when cooling the conditioned space, and $T_e > T_s \implies q_{ls} < 0$, $T_e > T \implies q_l < 0$.

Operating the heat pump removes heat from the store rather than pushing heat into the store, so the first term in equation (2) changes sign to give

$$\dot{Q} = -\eta P - k_s \left(T_{s0} + \frac{Q}{C_s} - T_e \right) - q. \tag{15}$$

Consequently, the Hamiltonian (3), with $\lambda_0 = 1$, changes to

$$\mathcal{H} = -Pr + \lambda_1 \left(-\eta P - k_s \left(T_{s0} + \frac{Q}{C_s} - T_e \right) - q \right)$$

$$+ \lambda_2 \frac{1}{C}(q - k(T - T_e)) \tag{16}$$

$$= -(\lambda_1 \eta + r)P + (\lambda_2 \frac{1}{C} - \lambda_1)q + \cdots . \tag{17}$$

We replace the constraint $q \in [0, q_{max}]$ with $q \in [-q_{max}, 0]$ and the constraint $T \geq T_{min}$ with $T \leq T_{max}$ to give the new Lagrangian

$$\mathcal{L} = \mathcal{H} + \mu_1 P + \mu_2 (P_{max} - P) + \mu_3 (q + q_{max}) - \mu_4 q + \nu_1 Q + \nu_2 (Q_{max} - Q) + \nu_3 (T_{max} - T). \tag{18}$$

The Karush-Kuhn-Tucker condition (6) changes to

$$\frac{\partial \mathcal{L}}{\partial P} = 0 \implies -\lambda_1 \eta - r + \mu_1 - \mu_2 = 0 \tag{19}$$

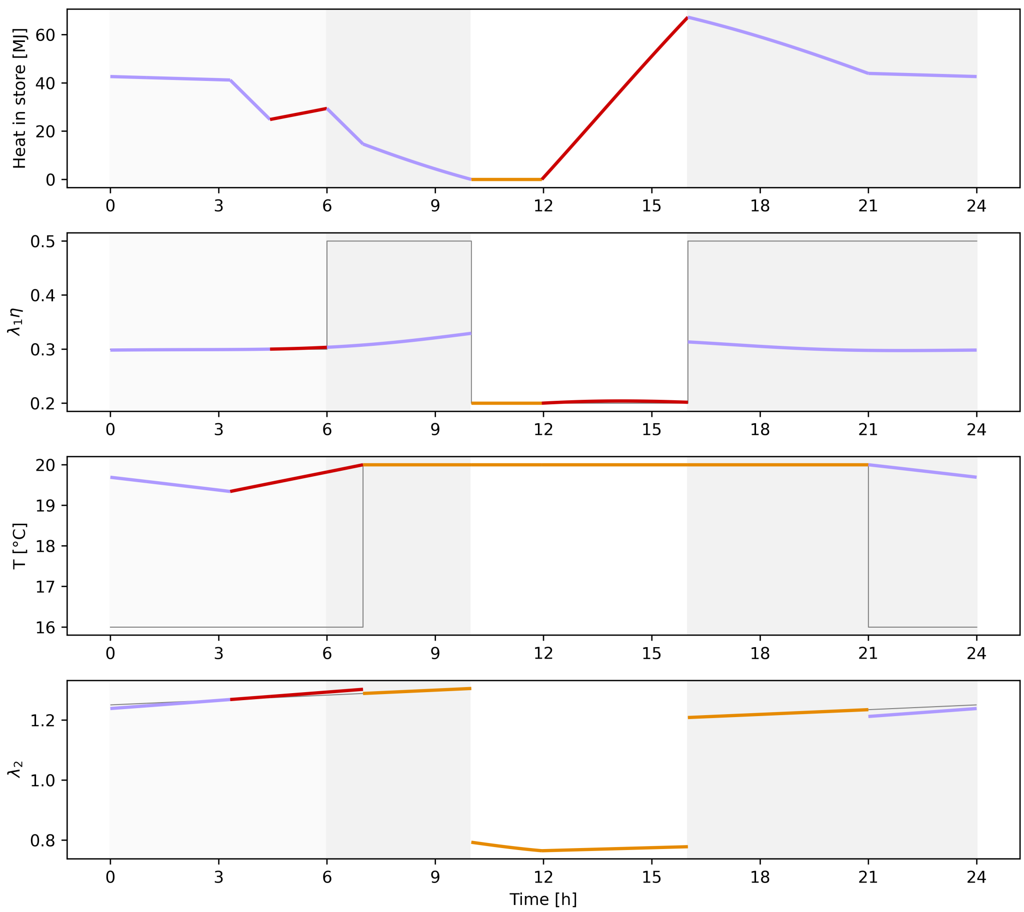

t	Q	$\lambda_1\eta$	T	λ_2	$C\lambda_1$	P	q
0.00	42.60	0.2982	19.69	1.2381	1.2501	0	0
3.32	41.16	0.2992	19.34	1.2680	1.2680	0	q_{max}
4.42	24.82	0.3000	19.54	1.2781	1.2740	P_{max}	q_{max}
6.00	29.40	0.3034	19.82	1.2927	1.2827	0	q_{max}
7.00	14.66	0.3075	20.00	1.3020	1.2882	0	$(0, q_{max})$
10.00	0.00	0.2000	20.00	1.1969	1.1969	$(0, P_{max})$	$(0, q_{max})$
11.95	0.00	0.2000	20.00	0.7645	0.7645	P_{max}	$(0, q_{max})$
16.00	67.20	0.3133	20.00	0.7779	1.2081	0	$(0, q_{max})$
21.00	43.90	0.2976	20.00	1.2116	1.2342	0	0
24.00	42.60	0.2982	19.69	1.2381	1.2502	0	0

Figure 5: Operation of the store and conditioned space temperature for Example 3

and equation (9) changes to

$$\dot{\lambda}_2 = -\frac{\partial \mathcal{L}}{\partial T} = \frac{k}{C}\lambda_2 + \nu_3. \tag{20}$$

The conditions that determine the optimal control for P become:

- if $-\lambda_1\eta > r$ then $\mathcal{H}$ is maximised by setting $P = P_{max}$
- if $-\lambda_1\eta = r$ then P can take on any value in $[0, P_{max}]$
- if $-\lambda_1\eta < r$ then $\mathcal{H}$ is maximised by setting $P = 0$.

The conditions that determine the optimal control for q change to:

- if $\lambda_2 - C\lambda_1 > 0$ then $\mathcal{H}$ is maximised by setting $q = 0$
- if $\lambda_2 - C\lambda_1 = 0$ then $q \in [-q_{max}, 0]$
- if $\lambda_2 - C\lambda_1 < 0$ then $\mathcal{H}$ is maximised by setting $q = -q_{max}$.

The first singular control mode, for P, changes to $-\lambda_1\eta = r$. The second singular control mode, for q, does not change. Equation (11) changes to

$$h^*(t) = \begin{pmatrix} Q \\ Q_{max} - Q \\ T_{max} - T \end{pmatrix} \implies h^*_x[t_f] = \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

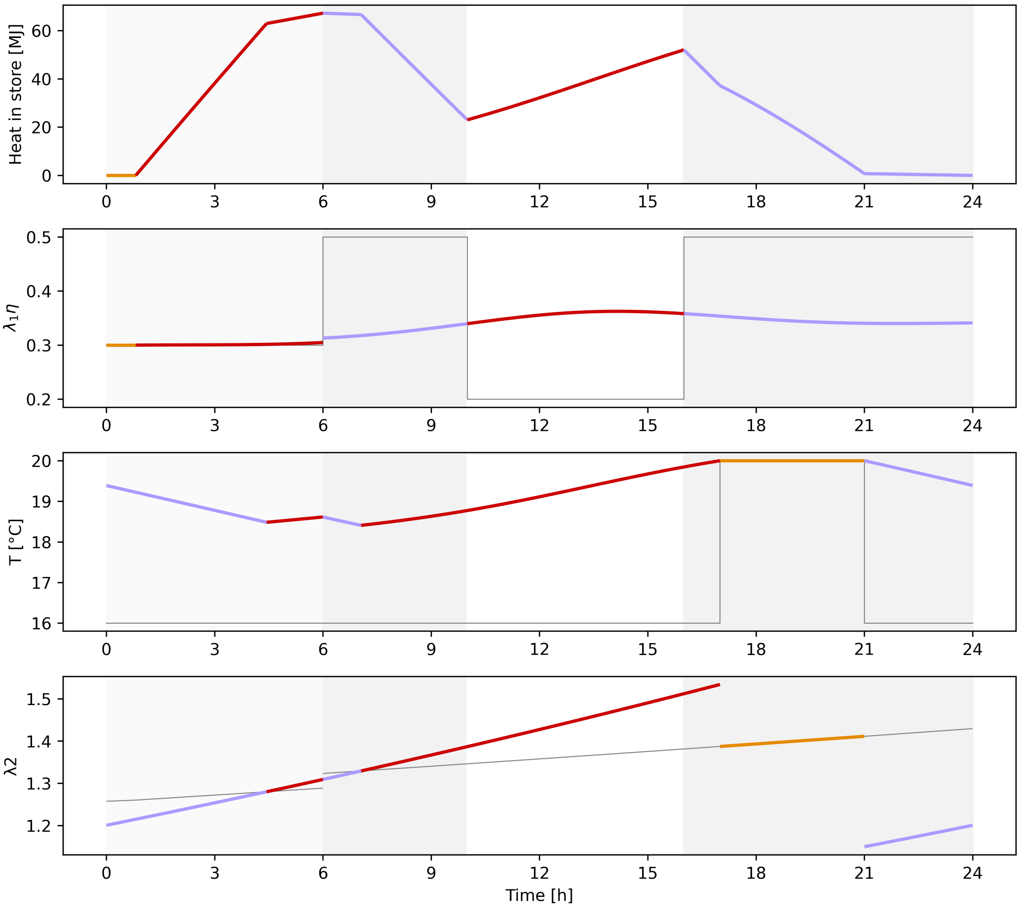

t	Q	$\lambda_1\eta$	T	λ_2	$C\lambda_1$	P	q
0.00	0.00	0.3000	19.39	1.2008	1.2576	$(0, P_{max})$	0
0.81	0.00	0.3000	19.22	1.2150	1.2603	P_{max}	0
4.44	62.92	0.3014	18.48	1.2801	1.2801	P_{max}	q_{max}
6.00	67.20	0.3130	18.61	1.2801	1.3233	0	0
7.06	66.63	0.3176	18.41	1.3293	1.3294	0	q_{max}
10.00	23.01	0.3396	18.77	1.3868	1.3462	P_{max}	q_{max}
16.00	52.00	0.3582	19.85	1.5119	1.3813	0	q_{max}
17.00	37.19	0.3536	20.00	1.3872	1.3872	0	$(0, q_{max})$
21.00	0.75	0.3403	20.00	1.1500	1.4112	0	0
24.00	0.00	0.3000	19.39	1.2008	1.4295	0	0

Figure 6: Operation of the store and conditioned space temperature for Example 4

The adjoint variable $\lambda = (\lambda_1, \lambda_2)$ can have a jump discontinuity at time τ if

$$\lambda(\tau^-) = \lambda(\tau^+) + \zeta(\tau) h_x^*[\tau]$$
$$= \lambda(\tau^+) + (\zeta_1, \zeta_2, \zeta_3) \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ 0 & -1 \end{pmatrix}$$
$$\mathcal{H}^*[\tau^-] = \mathcal{H}^*[\tau^+] - \zeta(\tau) h_t^*[\tau] = \mathcal{H}^*[\tau^+]$$

where $\zeta(\tau) > 0$ and $\zeta(\tau) h^*[\tau] = 0$. The jump discontinuities in λ can be written as

$$\lambda_1[\tau^-] = \lambda_1[\tau^+] + \zeta_1 - \zeta_2$$
$$\lambda_2[\tau^-] = \lambda_2[\tau^+] - \zeta_3.$$

and so

- λ_1 can jump down if $Q = 0$
- λ_1 can jump up if $Q = Q_{max}$
- λ_2 can jump up if $T = T_{max}$.

The efficiency of the heat pump becomes

$$\eta(t) = \frac{aT_c}{T_e(t) - T_c} \tag{21}$$

where, typically, T_c is 5 °C and $a = 0.4$.

The energy capacity of the store is smaller than when heating because of the smaller operating temperature range.

If the water temperature in the store will vary between 5 °C and 15 °C and the store contains 400 kg of water then the thermal energy storage capacity will be $Q_{\max} = 16.8$ MJ.

The temperature constraint for the conditioned space changes from $T \geq T_{\min}$ to $T \leq T_{\max}$, which changes the last term in the Lagrangian (4) to $\nu_3(T_{\max} - T)$, with complementary slackness conditions $\nu_3 \geq 0$ and $\nu_3(T_{\max} - T) = 0$.

7.2. Singular control modes

7.2.1. First singular control mode: partial power

The analysis in Section 5.3.1 of the first singular control mode changes. If $-\lambda_1\eta = r$ on any interior interval $[\tau_0, \tau_1]$ with $\tau_0 < \tau_1$, and the price r is constant on that interval then differentiating $-\lambda_1\eta = r$ with respect to t gives

$$\frac{d}{dt}\left[-\lambda_1\eta\right] = -\eta\dot{\lambda}_1 - \lambda_1\dot{\eta} = 0 \implies \dot{\lambda}_1 = -\frac{\lambda_1\dot{\eta}}{\eta}.$$

From (8) we also have

$$\dot{\lambda}_1 = \frac{k_s}{C_s}\lambda_1 - \nu_1 + \nu_2,$$

so on a singular interval these two expressions for $\dot{\lambda}_1$ must coincide, and hence

$$\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) - \nu_1 + \nu_2 = 0. \tag{22}$$

Here ν_1 and ν_2 must satisfy the complementary slackness conditions

$$\nu_1 Q = 0, \quad \nu_2(Q_{\max} - Q) = 0, \quad \nu_1, \nu_2 \geq 0.$$

Hence

$$\nu_1 > 0 \implies Q = 0, \quad \nu_2 > 0 \implies Q = Q_{\max}.$$

The factor $k_s/C_s + \dot{\eta}/\eta$ will be positive unless the CoP η is decreasing at a rate $\dot{\eta} < -\eta k_s/C_s$. Our examples show that this is possible in practice.

Now consider cases where (22) is satisfied, if $k_s > 0$:

- If $0 < Q < Q_{\max}$ then $\nu_1 = \nu_2 = 0$ and (22) reduces to
$$\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) = 0.$$
This can occur if $\lambda_1 = 0 \implies r = 0$, i.e. when the price is zero, or if $k_s/C_s + \dot{\eta}/\eta = 0$, which requires η to be reducing at the precise rate $\dot{\eta} = -\eta k_s/C_s$, which is unlikely.

- If $Q = 0$ then $\nu_2 = 0$ and $\nu_1 \geq 0$, and (22) gives
$$\nu_1 = \lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) \geq 0.$$
If the price $r >= 0$ then $\lambda_1 = -r/\eta < 0$ so we must have
$$\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta} < 0. \tag{23}$$
That is, you can maintain partial power with $Q = 0$ only if the CoP η is reducing rapidly.

- If $Q = Q_{\max}$ then $\nu_1 = 0$ and $\nu_2 \geq 0$, and (22) gives
$$\nu_2 = -\lambda_1\left(\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta}\right) \geq 0.$$
For $r > 0$ we have $-\lambda_1 = r/\eta > 0$, so we require
$$\frac{k_s}{C_s} + \frac{\dot{\eta}}{\eta} > 0. \tag{24}$$
That is, the CoP η must not be decreasing too rapidly.

7.2.2. Second singular control mode: partial heat transfer

The analysis in Section 5.3.2 also changes. If $\lambda_2 = C\lambda_1$ on any interior interval $[\tau_0, \tau_1]$ with $\tau_0 < \tau_1$ then we have $\dot{\lambda}_2 = C\dot{\lambda}_1$ and so, from (8) and (20),

$$C\left(\frac{k_s}{C_s}\lambda_1 - \nu_1 + \nu_2\right) = k\lambda_1 + \nu_3$$
$$\implies \left(\frac{k}{C} - \frac{k_s}{C_s}\right)\lambda_1 + \nu_1 - \nu_2 + \frac{\nu_3}{C} = 0.$$

Once again, we assume that $k_s/C_s < k/C$ in practice. The following conditions are necessary for this singular mode to apply:

- if $\nu_1 = \nu_2 = \nu_3 = 0$, because $T < T_{\max}$ and $q \neq 0$ and $q \neq Q_{\max}$, then we must have $\lambda_1 = 0$ and hence $\lambda_2 = 0$
- if $\lambda_1 > 0$ then we require $\nu_2 > 0$, so $Q = Q_{\max}$
- if $\lambda_1 < 0$ then we require $\nu_1 > 0$ or $\nu_3 > 0$, so $Q = 0$ or $T = T_{\max}$.

7.3. Example

Figure 7 shows a typical outdoor temperature profile for a hot day in South Australia. With this temperature profile the heat-pump CoP varies between 3.2 during the middle of the day and 5.7 overnight.

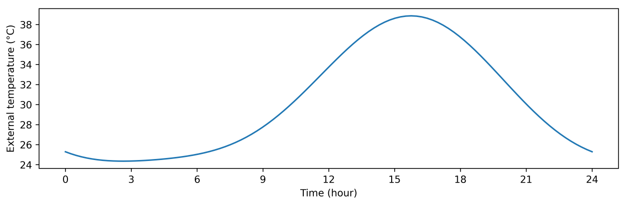

Figure 7: Summer external temperature profile

The indoor temperature requirement for this example is

$$T_{\max} = \begin{cases} 23\,°\mathrm{C} & t \in [06{:}00, 09{:}00) \cup [17{:}00, 22{:}00) \\ 26\,°\mathrm{C} & \text{otherwise.} \end{cases}$$

Figure 8 shows a control strategy satisfying the necessary conditions. In this example, the store capacity $Q_{\max} = 16.8$ MJ is not enough to maintain $T = T_{\max}$ during the interval $t \in [17, 22]$, and so the conditioned space is pre-cooled while the store is cooled during the cheap daytime period. However, the store does have enough capacity to maintain $T = T_{\max}$ during the interval $t \in [6, 10]$.

The process for finding a strategy that meets the necessary conditions is similar to the process for heating, but in this case involves more phases. The key features of the strategy are:

- The store must be held at $Q = 0$ during the middle of the day. This gives the steep downward portion of the $C\lambda_1$ curve that determines when the pre-cooling of the building finishes.
- The store must be at $Q = Q_{\max}$ at time $t = 24$ so that $-\lambda_1$ can jump down at $t = 0$.
- The store must also be at $Q = Q_{\max}$ at time $t = 10$ so that $-\lambda_1$ can jump down, which causes $C\lambda_1$ to jump up, which gives a 'gap' for the λ_2 curve to pass through at $t = 10$.

To help find the strategy, we formulated the strategy with parameters including switching times and adjoint jumps, and used an optimisation procedure to adjust these parameters to minimise temperature errors and ensure adjoint continuity.

8. Conclusion

We have used Pontryagin's principle to derive necessary conditions for the optimal control of a heat-pump air-conditioning system with thermal storage. These conditions lead to three distinct control modes for the heat pump, and three distinct control modes for the transfer of heat to or from the conditioned space. The heat pump:

- operates at maximum power when the price of electricity is sufficiently low and the efficiency of the heat-pump is sufficiently high
- operates at intermediate power only if the thermal store is empty or full
- does not operate if the price of electricity is high or the efficiency of the heat pump is low.

The transfer of heat between the store and the conditioned space:

- occurs at the maximum possible rate when heating or cooling the the conditioned space
- occurs at an intermediate rate only when maintaining the target temperature in the conditioned space
- is zero when active heating or cooling is not required.

In cases where the store does not have sufficient capacity maintain to maintain the target temperature during a period when the price of electricity is high, the optimal strategy will pre-heat or pre-cool the conditioned space while the price of electricity is low.

When heating, the optimal strategy avoids storing more heat than necessary and avoids storing heat for longer than necessary, to minimise heat losses from the store. When cooling, the optimal strategy avoids removing more heat from the store than necessary and avoids cooling the store for longer than necessary, to minimise heat flow into the store from the environment.

In each of our examples, we found a control strategy by piecing together control phases for the heat pump and for the heat transfer to or from the conditioned space to:

- satisfy the dynamics and continuity conditions for the temperature in the conditioned space and the heat in the store
- satisfy the dynamics and continuity conditions for the two adjoint variables determining the control.

In each example the solution we found appeared to be unique. However, we have not proved this.

Our analysis shows the structure of an optimal control, which can be used to develop simple, practical controllers for heat-pump air-conditioning systems with thermal storage.

Acknowledgements

Rong Xu was supported by a University of South Australia Vacation Research Scholarship and an Australian Mathematical Sciences Institute Summer Research Scholarship. We also thank Dr Michael Evans from Adelaide University for advice on our thermal models.

CRediT authorship contribution statement

Peter Pudney: Conceptualisation, methodology, formal analysis, software, writing, visualisation. **Rong Xu:** Formal analysis, software, writing, visualisation.

References

Arteconi, A., Hewitt, N., Polonara, F., 2013. Domestic demand-side management (DSM): Role of heat pumps and thermal energy storage (TES) systems. Applied Thermal Engineering 51, 155–165. doi:10.1016/j.applthermaleng.2012.09.023.

Chowdhury, N.R., Ofir, R., Zargari, N., Baimel, D., Belikov, J., Levron, Y., 2021. Optimal control of lossy energy storage systems with nonlinear efficiency based on dynamic programming and Pontryagin's minimum principle. IEEE Transactions on Energy Conversion 36, 524–533. doi:10.1109/TEC.2020.3004191.

Cirocco, L., Pudney, P., Riahi, S., Liddle, R., Semsarilar, H., Hudson, J., Bruno, F., 2022. Thermal energy storage for industrial thermal loads and electricity demand side management. Energy Conversion and Management 270, 116190. doi:10.1016/j.enconman.2022.116190.

Dar, U.I., Sartori, I., Georges, L., Novakovic, V., 2014. Advanced control of heat pumps for improved flexibility of net-ZEB towards the grid. Energy and Buildings 69, 74–84. doi:10.1016/j.enbuild.2013.10.019.

Fleming, J., Barbour, E., Urquhart, A., 2025. Optimal control of heat pumps with thermal storage under time-of-use tariffs, in: 2025 IEEE 64th Conference on Decision and Control (CDC), pp. 4738–4743. doi:10.1109/CDC57313.2025.11312972.

Gates, D.J., Westcott, M., 1996. Solar cars and variational problems equivalent to shortest paths. SIAM Journal on Control and Optimization 34, 428–436. doi:10.1137/S0363012993260276.

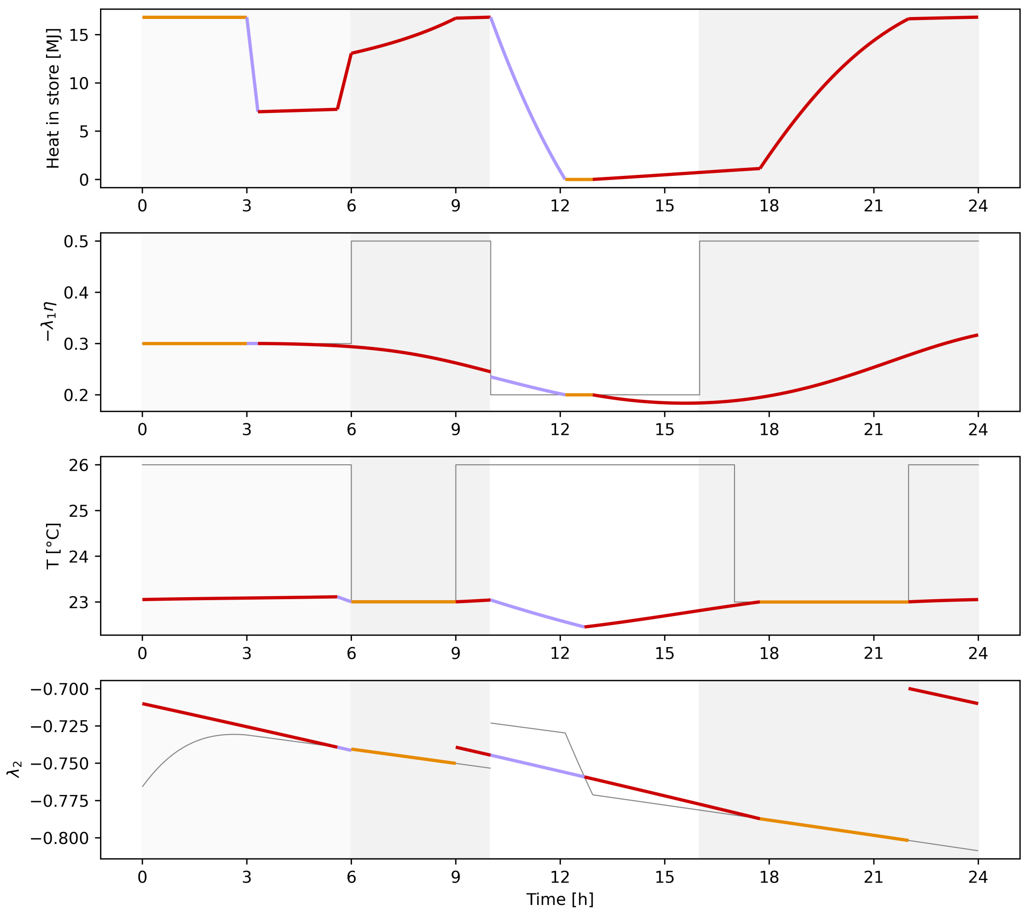

t	Q	$-\lambda_1\eta$	T	λ_2	$C\lambda_1$	P	q
0.00	16.80	0.3000	23.05	-0.7101	-0.7659	$(0, P_{max})$	0
3.00	16.80	0.3000	23.08	-0.7256	-0.7311	P_{max}	0
3.32	7.01	0.3000	23.09	-0.7273	-0.7321	0	0
5.60	7.26	0.2956	23.11	-0.7393	-0.7393	0	$-q_{max}$
6.00	13.06	0.2938	23.00	-0.7406	-0.7406	0	$(-q_{max}, 0)$
9.00	16.70	0.2621	23.00	-0.7393	-0.7502	0	0
10.00	16.80	0.2350	23.04	-0.7446	-0.7231	P_{max}	$-q_{max}$
12.14	0.00	0.2000	22.56	-0.7562	-0.7298	$(0, P_{max})$	$-q_{max}$
12.70	0.00	0.2000	22.45	-0.7592	-0.7592	$(0, P_{max})$	0
12.93	0.00	0.2000	22.47	-0.7605	-0.7713	0	0
17.73	1.13	0.1949	23.00	-0.7873	-0.7873	0	$(-q_{max}, 0)$
22.00	16.63	0.2773	23.00	-0.7000	-0.8018	0	0
24.00	16.80	0.3168	23.05	-0.7101	-0.8087	$(0, P_{max})$	0

Figure 8: Operation of the store and conditioned space temperature for a cooling example

Hartl, R.F., Sethi, S.P., Vickson, R.G., 1995. A survey of the maximum principles for optimal control problems with state constraints. SIAM Review 37, 181–218.

Henze, G.P., Dodier, R.H., Krarti, M., 1997. Development of a predictive optimal controller for thermal energy storage systems. HVAC&R Research 3, 233–264. doi:10.1080/10789669.1997.10391376.

Jarvinen, J., Goldsworthy, M., Pudney, P., White, S., Cirocco, L., Bruno, F., 2023. Using building cooling for load shifting on a 5-minute wholesale energy market. Journal of Building Engineering 73, 106777. doi:10.1016/j.jobe.2023.106777.

Kintner-Meyer, M., Emery, A., 1995. Optimal control of an HVAC system using cold storage and building thermal capacitance. Energy and Buildings 23, 19–31. doi:doi.org/10.1016/0378-7788(95)00917-M.

Liu, F., Zhu, W., Zhao, J., Ren, J., Groll, E.A., Cai, Y., 2017. A new method for optimal control of a dual-mode CO2 heat pump with thermal storage. Applied Thermal Engineering 125, 1123–1132. doi:10.1016/j.applthermaleng.2017.07.068.

Lovegrove, K., Alexander, D., Bader, R., Edwards, S., Lord, M., Mojiri, A., Rutovitz, J., Saddler, H., Stanley, C., Urkalan, K., Watt, M., 2019. Renewable energy options for industrial process heat, Appendix E. Technical Report. Australian Renewable Energy Agency (ARENA).

Muñoz, F., Garcia-Hernandez, R., Ruelas, J., Palomares-Ruiz, J., Álvarez Macías, 2022. Optimal operation for reduced energy consumption of an air conditioning system using neural inverse optimal control. Mathematics 10. doi:10.3390/math10050695.

Renaldi, R., Kiprakis, A., Friedrich, D., 2017. An optimisation framework for thermal energy storage integration in a residential heat pump heating system. Applied Energy 186, 520–529. doi:10.1016/j.apenergy.2016.02.067.

Risbeck, M.J., Maravelias, C.T., Rawlings, J.B., Turney, R.D., 2017. A mixed-integer linear programming model for real-time cost optimization of building heating, ventilation, and air conditioning equipment. Energy and Buildings 142, 220–235. doi:10.1016/j.enbuild.2017.02.053.

Salameh, C., Schalbart, P., Peuportier, B., 2021. Optimal control strategies for energy production systems using buildings thermal mass, in: Building Simulation 2021 Conference, Sep 2021, Bruges, Belgium. URL: https://hal.science/hal-03379993v1.

Schibuola, L., Scarpa, M., Tambani, C., 2015. Demand response management by means of heat pumps controlled via real time pricing. Energy and Buildings 90, 15–28. doi:10.1016/j.enbuild.2014.12.047.

Zargari, N., Levron, Y., Belikov, J., 2019. Optimal control of energy storage devices based on pontryagin's minimum principle and the shortest path method, in: 2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe), pp. 1–5. doi:10.1109/ISGTEurope.2019.8905748.

Zhang, Q., Stockar, S., Canova, M., 2016. Energy-optimal control of an automotive air conditioning system for ancillary load reduction. IEEE Transactions on Control Systems Technology 24, 67–80. doi:10.1109/TCST.2015.2418322.