

Institute of Architecture of Application Systems

University of Stuttgart  
Universitätsstraße 38  
D-70569 Stuttgart

Bachelorarbeit

# **Domain-Independent AI Planning for Demand-Side Management in Office Buildings**

Lukas Heiland

**Course of Study:** Computer Science  
**Examiner:** Prof. Dr. Marco Aiello  
**Supervisor:** Dr. Ilche Georgievski

**Commenced:** May 4, 2020  
**Completed:** November 4, 2020



---

## Abstract

Commercial buildings are characterised by high operational costs and high CO<sub>2</sub> emission levels due to the high demand for energy. The management of energy demand, also known as Demand-side Management (DSM), allows buildings to make informed decisions about their energy consumption. However, current building management systems control heating, ventilation and air conditioning, lighting and other aspects only by basic scheduling functions. At the same time, Internet of Things (IoT) devices have the capability to transform building management systems into intelligent tools for DSM. This means that buildings equipped with a large variety of sensors and actuators have the potential to become intelligent and significantly reduce energy consumption and, as a consequence, operational cost. However, a small amount of installed suitable infrastructure in offices and a lack of understanding of the benefits of DSM hinders this development.

This calls for developing approaches that can efficiently find sequences of actions which, upon execution, reduce energy consumption in an office building. In this context, Artificial Intelligence (AI) planning provides powerful techniques to intelligently plan an office building's demand by computing effective plans or schedules of device actions.

This thesis presents an approach for demand-side management by using AI planning. It involves defining and modeling scenarios from demand-side management as a domain-independent planning problem using the Planning Domain Definition Language (PDDL) and solving the planning problem using an existing AI planner. Additionally, the approach is evaluated regarding performance and potential for savings, both energy-wise and financially. The evaluation provides evidence in favor of DSM being beneficial for saving energy and operational cost.



# Contents

<b>1</b>	<b>Introduction</b>	<b>7</b>
1.1	Motivation . . . . .	7
1.2	Objective of This Thesis . . . . .	8
1.3	Structure of This Thesis . . . . .	8
<b>2</b>	<b>Related Works</b>	<b>9</b>
<b>3</b>	<b>Background</b>	<b>11</b>
3.1	Demand-side Management . . . . .	11
3.2	Artificial Intelligence Planning . . . . .	14
3.3	Planning Domain Definition Language . . . . .	16
<b>4</b>	<b>Approach</b>	<b>19</b>
4.1	Classification . . . . .	19
4.2	Scenarios . . . . .	20
4.3	Modeling with PDDL . . . . .	22
<b>5</b>	<b>Evaluation</b>	<b>27</b>
5.1	Experiment Design . . . . .	27
5.2	Results . . . . .	27
5.3	Discussion . . . . .	28
<b>6</b>	<b>Conclusion</b>	<b>31</b>
6.1	Summary of the Thesis . . . . .	31
6.2	Future Work . . . . .	31
	<b>Bibliography</b>	<b>33</b>



# 1 Introduction

## 1.1 Motivation

Commercial buildings are characterised by high operational costs and high CO<sub>2</sub> emission levels due to the high demand for energy. Office buildings have a special role in this category; in the European Green Building Programme, they are the most common building category while having the lowest average energy savings [DCB17].

The European Union (EU) made it a core component of their energy policy to improve energy efficiency in its member states [Com14]. The EU recognizes the importance of buildings for achieving their energy goals: Among other things, the EU's Energy Performance of Buildings Directive requires countries to set cost-optimal minimum energy performance requirements for new buildings, for existing buildings undergoing major renovation, and for the replacement or retrofit of building elements [Com19].

Demand-side Management (DSM) is a powerful tool for realizing this potential. DSM includes all measures to improve energy efficiency that can be realized on the demand side of the energy system. This helps both consumers and producers: buildings are empowered to make informed decisions regarding their energy consumption while producers can shape load demand profiles [PD11].

However, automated DSM measures and their coordination still leave a lot to be desired, which is the result of a lack of suitable infrastructure in offices and lack of understanding of the benefits of DSM [PD11]. A simple example of a suboptimal DSM measure would be a thermostat which turns on the AC in a room as soon as the temperature rises above a certain threshold and turns it off when the temperature is under the threshold again. However, what is not considered in this example is that it would be more energy efficient to lower the blinds in order to reduce the influence of sunlight or to open windows and doors. Moreover, there could be different energy prices depending on the time of day [PD11], and this fact can also influence the decision making process.

Using smart Internet of Things (IoT) devices, buildings have the potential to turn into systems employing intelligent DSM [GNN+17]. In the context of the previous example, this could mean using a multitude of sensors in combination with a controlling unit in order to devise a sequence of actions that both reduces the temperature and minimizes energy consumption.

One possible way to unleash this potential is Artificial Intelligence (AI) planning, which offers powerful techniques for automated and dynamic selection and organisation of actions that, when executed, achieve a given objective [GNT04]. Using AI planning to intelligently plan the demand of (office) buildings could be helpful in achieving the EU's energy efficiency goals.

## **1.2 Objective of This Thesis**

The aim of this thesis is to develop an approach for DSM in office buildings. Since this thesis is focused on AI planning, the approach is centered around reducing energy consumption by controlling the devices in an office and disregards measures related to building structure, e.g. installing better insulation.

Another goal of the development of the approach is to gain insight into the benefits of DSM regarding energy savings and reduction of operational cost, since a lack thereof was previously identified as one of the main obstacles for a wider acceptance and adaptation of DSM measures.

The approach includes creating fictional but realistic DSM problem settings, translating them into the Planning Domain Definition Language (PDDL), and using domain-independent AI planning to compute plans for these problems. A short study will show the efficiency of the developed approach, regarding performance and potential for savings, both ecologically and financially.

## **1.3 Structure of This Thesis**

Chapter 2 shortly describes the current state of related research. In Chapter 3, important background knowledge is introduced. After that, the actual approach is presented in Chapter 4, Chapter 5 contains the experiment which was designed to test the approach and the results thereof. The thesis is concluded with Chapter 6.



## 2 Related Works

**Overview of Demand-side Management (DSM)** Palensky and Dietrich [PD11] give an overview and a taxonomy of DSM, analyzes the many types of DSM, and gives an outlook on major DSM-related projects. It also briefly mentions benefits and challenges for DSM and identifies a lack of suitable Information and Communication Technology (ICT) as one of the main challenges. This thesis will use the terminology introduced in this paper.

**DSM: benefits and challenges** Strbac's paper [Str08] goes more into detail about the benefits and challenges of DSM. It confirms lack of ICT as one of the major challenges for DSM, as well as the lack of understanding of the benefits of DSM. This thesis provides evidence supporting the importance and potential of DSM and therefore the benefits of employing ICT infrastructure, such as intelligent building management systems and IoT devices.

**Demand Side Management in Smart Grid** In this work [LSS12], an algorithm is presented which optimizes the overall load in several simulated Smart Grids [GSK+11]. This thesis differs in two ways: it does not specify which form of power grid is in place, and it uses AI planning to employ DSM measures.

**AI planning in dynamic environments** Dynamic Environments pose a challenge to AI planning due to unexpected events and uncertain outcomes of some actions. [KWLA13] presents an architecture for a Smart Home, which uses domain-independent AI planning to generate compositions of smart device actions. Performance and usability analyses are performed. This thesis focuses more on energy efficiency rather than performance and usability, although performance is also relevant since plans should be computed in an acceptable amount of time.

**Service coordination for Intelligent Buildings** In [GNN+17], a system is developed that uses Hierarchical Task Network (HTN) planning [GA15] – a sub-category of AI planning – to improve occupants' experiences in a building while also reducing energy consumption. While this thesis is about a closely related topic, it will use domain-independent, classical AI planners to gain more insight into energy savings through AI-planning-driven DSM. Moreover, this thesis is different from [GNN+17] because it does not focus on activity recognition, instead it assumes that e.g. the absence of an occupant has already been detected.



## 3 Background

### 3.1 Demand-side Management

Demand-side Management (DSM) includes all measures to improve the energy system at the side of consumption. This section introduces key concepts as well as a short discussion of the potential benefits and challenges that come with DSM.

#### 3.1.1 Concepts

There are several ways of categorizing DSM [FA19; KPH06; PD11]. This thesis considers three main categories, which will be described in the following sub-sections.

##### **Energy Efficiency Measures**

Energy efficiency can decrease energy consumption while maintaining the level of service offered to occupants. These measures are usually driven by a desire for conservation, environmental protection and/or utility bill savings, and they aim to reduce peak load by reducing overall consumption [KPH06].

Energy efficiency measures include, but are not limited to, changes to equipment, e.g. exchanging old light bulbs with compact fluorescent lamps or installing sensors and actuators to better identify power guzzlers.

##### **Time-Based Energy Pricing**

Time-Based Energy Pricing refers to a change in tariffs during certain periods of time. These changes can happen more or less frequently, as explained below.

*Time-Of-Use (TOU) pricing:* Prices change between TOU pricing periods. These periods can be hours, days or even seasons long [DG12]. Rates during one pricing period are guaranteed by the supply contract between energy provider and customer [PD11]; this means, a change in TOU rates would require a change in the contract, so TOU rates can be assumed to rarely change over time.

*Day-ahead pricing:* As the name suggests, electricity prices are communicated to the customer one day ahead. Suppliers could just forward wholesale prices or include network data or load data gathered from customers to optimize the pricing model [DG12].

*Real-time pricing:* Energy prices applied to customers are updated hourly based on wholesale prices [DG12].

### **Demand Response**

Demand Response (DR) refers to the modification of customer electricity usage at times of peak usage in order to help address system reliability, reflect market conditions and pricing, and support infrastructure optimization or deferral [KPH06]. According to [HP08], day-ahead and real-time pricing can also be assigned to a subcategory called *time-based DR*.

*Incentive-based DR* involves methods like direct load control, where the utility and grid operators get full access to their customers' processes, and emergency programs, where customers get emergency signals when the grid is in a critical situation [PD11]. Since these programs require the initiative of utility or grid operators, they are out of the scope of this thesis.

Often, another category by the name of *Spinning Reserve* is mentioned. This set of methods is out of the scope of this thesis and will therefore not be discussed here. The interested reader is kindly referred to [PD11].

#### **3.1.2 Benefits of DSM**

##### **Less energy consumption**

The European Commission's *GreenBuilding Programme* [E3P20], which was active from 2005 to 2014, reports that office buildings, which implemented at least one energy efficiency measure, showed average savings of 87 kWh/m<sup>2</sup>/y [DCB17]. Considering the Europe average specific energy consumption in the non-residential sector is 280 kWh/m<sup>2</sup> [EAD+11], this is a rather large improvement and can both contribute to the European Union's emission goals and a saving in energy expenses on the consumer side.

##### **Reduced generation margin**

The generation capacity of a system must exceed the maximum demand in the system; this generation margin was historically considered to be sufficient at around 20% [Str08]. This results in a lot of unused capacity during off-peak times. Since DSM measures have the power to shave off peaks, the generation margin can be lowered and therefore unused capacity is reduced.

### 3.1.3 Challenges of DSM

Strbac identifies several challenges DSM has to overcome in order to become more widespread and successful [Str08]. Since not all of them relate to this thesis, only some are described in the following.

#### **Lack of Understanding of the Benefits**

There is a need for better quantification of the costs and benefits of DSM. For example, in segments of a system with a lot of spare capacity, the value of DSM will be a lot lower than in segments of the same system where the maximum of capability is near.

#### **Lack of Infrastructure**

Sensors, measurement and control devices and metering devices are largely absent from electricity systems. With the rising popularity of Internet of Things (IoT) devices, this is becoming less of a challenge. However, a more widespread use of IoT devices does not guarantee the success of DSM.

#### **Lack of Smart Approaches**

While the presence of IoT devices is nice-to-have, it does not directly relate to DSM success. There is a need for smart approaches to find management strategies for these interconnected devices automatically, and this thesis aims to provide one.

## 3.2 Artificial Intelligence Planning

AI planning deals with the development of representation languages for planning problems and with the development of algorithms for plan construction [Sch05]. An AI planner, which is a specific implementation of such languages and algorithms, tries to find a sequence of actions that, when executed, achieve a given goal. Key definitions are given in the following, according to [ERH18].

### Definition 3.2.1

A planning domain is a tuple  $(L, O)$ , where  $L$  is defined as the language and  $O$  is a set of operators. An operator  $o$  is a triple  $(name(o), pre(o), eff(o))$ , where

- $name(o)$  is a syntactic expression needed to refer unambiguously to instances of  $o$ . It has the form  $n(x_1, \dots, x_k)$ , where  $n$  is the operator symbol unique for each operator and  $x_1, \dots, x_k$  is a list of parameters that appear in  $o$ ,
- $pre(o)$  are preconditions; predicates that must be true in order to use the operator,
- $eff(o)$  effects of an action; predicates the operator will make true in the state after its application.

An example of an operator *pick-up* which lets a robotic arm pick up an item could look like this:  
 $(pick-up(arm, item), (empty(arm), clear(item)), (\neg(empty(arm)), holding(arm, item)))$

### Definition 3.2.2

A planning problem is a tuple  $\mathcal{P} = (P, A, I, G)$ , where

- $P$  is a finite set of predicates,
- $A$  is a finite set of actions, that is, in many cases, instances of operators of the domain
- $I \subseteq P$  is the initial state and
- $G \subseteq P$  is the goal state

Additionally, a state  $s \subseteq P$  is defined by all predicates that evaluate to true in  $s$  [ERH18]. An action  $a$  is called *applicable* to a state  $s$  if all predicates in  $pre(o)$  of  $a$ 's ground operator  $o$  are true in  $s$ . For example, the action  $pick-up(arm_1, item_1)$  is applicable in state  $(free(arm_1), clear(item_1))$ .

Planners can be classified regarding their specificity towards domains:

1. Domain-specific: Tuned well to the specifics of a certain domain; will not work well (if at all) in other domains.
2. Domain-configurable: Contains a domain-independent planning engine, will need input including information about the specific domain of the problem to solve.
3. Domain-independent: Works (in principle) for any domain, but is not as efficient as domain-specific planners in any given domain.

## Planning Algorithm

Planning can be seen as a search through the space of possible states. The general framework of planning algorithms is divided into three main steps. In the *refinement step*, additional information that could have influence over the next steps is computed. The nature of this information depends on the planning technique. The following *branching step* divides the current set of (partial) solutions into several sets to be explored individually. After that, unpromising members of the individual sets are removed in the *pruning step*. *Algorithm 3.1* gives a pseudocode example for this approach.

---

**Algorithm 3.1** Abstract planning algorithm in a state  $u$  [GNT04]

---

```
procedure ABSTRACTSEARCH( $u$  : state of a planning problem)
  if Terminal( $u$ ) then
    return  $u$ 
  end if
   $u \leftarrow$  Refine( $u$ )
   $B \leftarrow$  Branch( $u$ )
   $B' \leftarrow$  Prune( $B$ )
  if  $B' = \emptyset$  then return failure
  end if
  Nondeterministically choose  $v \in B'$ 
  return AbstractSearch( $v$ )
end procedure
```

---

The differences between AI planners stem from the implementation of the aforementioned steps; i.e., Metric-FF [Hofb], the planner used in this thesis, uses Enforced Hill-Climbing [WKG08], an algorithm which refines and prunes in a way that minimizes memory usage. An overview over different approaches to implementing this algorithm can be found in [GNT04].

### 3.3 Planning Domain Definition Language

The Planning Domain Definition Language (PDDL) is an AI planning language created by Drew McDermott and his colleagues for the 1998 International Planning Competition [McD00]. It has since become a de-facto standard in the planning community for representation and exchange of planning models [FL03].

PDDL separates the model of a planning problem into two parts, the domain description and the problem description. Domains and problems are usually maintained in separate files [AHK+98]. This enables one to pair the same domain with different problem files in order to solve multiple problems in the same domain.

Much of the structure of PDDL files is a Lisp-like list of parenthesised expressions [Lon03]. Strings preceded by a colon are treated as keywords. Variables begin with a question mark. Single-line comments can be inserted with a semicolon. Example domains and problems can be found at [Vel02].

With PDDL 2.1, there came a classification of the different language features into five levels with increasing expressivity [Lon03]. Level 1 consists of the STRIPS-fragment, the numeric extensions to PDDL are counted towards level 2, discretised durative actions comprise level 3, continuous durative actions appear in level 4 and the final level, level 5, features all extensions of PDDL 2.1 and additional components. For this thesis, all features up to and including level 3 are relevant.

The following two sub-sections briefly introduce basic knowledge about PDDL domains and problem definitions. For a more detailed reference of PDDL features, the interested reader is referred to [AHK+98] and [FL03].

#### 3.3.1 Domain Description

Every domain description starts with the domain name, followed by the requirements needed by this domain. To determine which PDDL language level the domain, and therefore the level of problems within it, one can look at this requirements section. Since not all planners support all language features [AHK+98], this section influences the choice of planner. For example, type hierarchies can be introduced by including the `:typing` requirement. For a full list of permitted arguments and an Extended Backus-Naur Form (EBNF) of a domain description, see [AHK+98].

After that, the domain designer can define predicates, which correspond to the set  $P$  of predicates from **Definition 3.2.2**. Action descriptions in PDDL refer to operators as defined in **Definition 3.2.1**. Every action has parameters, a precondition which has to be true before execution and an effect which takes place after execution. The actions from set  $A$  in **Definition 3.2.2** would be instances of these operators, where the parameters are actual objects in a planning problem.



### 3.3.2 Problem Description

A problem in PDDL, defined according to a domain, specifies the objects which are involved in the problem, an initial state, and the goal state to be reached [AHK+98].

Objects are the agents of a problem statement, they can perform actions and actions have effects on them.

The initial state is a set of predicates that are true at the beginning of the problem. At this point, it is worth mentioning that PDDL supports the idea of the *closed-world assumption*, which means that every predicate not explicitly listed as true is assumed to be false [AG19]. Similarly, a goal state is a set of predicates that are required to be true. Initial state and goal state of the problem description correspond to the sets  $I$  and  $G$  from **Definition 3.2.2**.

#### **Definition 3.3.1**

*A sequence of actions is called a solution for the problem if*

- 1. the action sequence can be executed starting in the provided starting state*
- 2. the predicates in the provided goal description are true at the end of the action sequence*

*assuming the goal description does not contain a so-called action expansion [AHK+98].*

Since the introduction of PDDL 2.1, there can be an optional `:metrics` section, which aims at helping planners evaluate their plans post hoc [FL03].



## 4 Approach

### 4.1 Classification

To test the efficiency of the approach developed in this thesis, two different scenarios are considered. These are formally defined using the classification for ubiquitous computing environments introduced in [GA16].

*Spatial properties* qualify the relations among entities and their environment. This subdimension itself can be divided into object locations and human locations. Spatial properties can be encoded by combining objects, constants and predicates in a PDDL domain.

*Temporal properties* characterize the relationship of entities with time. These relationships are expressed through time points and intervals, which can be represented by durative actions in PDDL.

*Behavioral inputs* are defined as the information expressing one's desires according to which a ubiquitous computing environment should behave. Requests provide models of desired results issued for the purpose of achieving a mandatory behaviour, adaptation, or organisation of environments. Preferences, on the other hand, include individual attitudes towards the environment behaviour (or organisation). Contrary to requests, preferences do not enforce mandatory satisfaction. In PDDL, requests can be mapped directly to goals in planning problems. The specification of preferences is supported since PDDL3 [GL05].

*Behavioral outputs* are actions performed by devices and robots, software components, and humans whose outcomes modify the states of the environments. All behavioral outputs are represented by actions in planning domains.

*Uncertainty* refers to situations of environments in which the information describing the current state is ambiguous and/or unreliable due to the inherited dynamism of the environments. This dynamism can come from unexpected events, from non-deterministic behavior of devices and from incomplete information about environment states (partial observability).

Since dealing with uncertainty requires a lot of work regarding orchestration, i.e. dealing with service failure [GNN+17], it is considered out of the scope of this thesis and will therefore not be dealt with.

### 4.2 Scenarios

In this section, the scenarios used to test the approach are introduced. These scenarios are designed for the purpose of this thesis. The assumption is that they take place inside an office building which has installed a variety of sensors and actuators and uses a building management system of some sort to gather data and control devices and actuators.

The first scenario is about lowering the temperature in an office in order to meet regulatory standards. It focuses on the usage of Energy Efficiency Measures from Section 3.1.1. The second scenario concerns itself with an occupant of an office leaving for lunch break. However, the occupant has left their work station, as well as ceiling and desk lamps and the heater, on. This scenario additionally includes time-based pricing as introduced in Section 3.1.1.

#### Reduce the Temperature in an Office to meet Regulatory Standards

In this scenario, the outside temperature is quite high and the office temperature rises above a critical threshold. The building management system recognizes that the office temperature is passing that threshold, pulls the outside temperature and subsequently triggers the planning process in order to find actions to reduce the office temperature.


The German Safety and Health at Work Act (ArbSchG) defines rules for the temperature in different kinds of workplaces. According to its Technical Rules section [ASR17], office temperatures of 26°C and above pose health risks for certain groups of the population, i.e. elderly workers or pregnant women. Table 4.1 summarizes the temperatures and consequences for the occupants according to the ArbSchG. It is worth noting that these rules only apply if the outside temperature exceeds 26°C.

Figure 4.1 shows a ground plan of this fictional office, which defines the object locations. The office has two doors on the northern side of the office, and three windows facing south for maximized natural lighting. There are four desks with work stations in the office, positioned as two groups of two work stations each. Two Air Conditioning (AC) units are installed on the ceiling. The doors, windows, work stations and AC units are numbered from left to right in Figure 4.1. The human locations are marked in Figure 4.1; the first occupant sits at work station 1, the leftmost work station in the office, while the second occupant works at work station 3, the second work station from the right.

The request of this scenario is 'reduce the temperature to below 26°C'. This is modeled directly through the predicates in the `:goal` section of the problem file. The preference would be to prioritize opening windows and doors over turning on AC units, since the former consumes less energy. In this thesis, this is realized by only allowing the use of AC units when the temperature is much higher than 26°C and declaring that opening windows and doors reduces the temperature by a lesser amount than AC units.

The behavioral outputs are trivial: AC units can be turned on and off, doors and windows can be opened and closed, blinds can be raised and lowered, and so on.

Temporal properties are facts like 'the doors have to be open as long as the windows' (to improve air flow and cooling) or 'before a window is opened, all AC units have to be turned off'.



**Figure 4.1:** Office with two ACs and four workstations

**Table 4.1:** Office temperature zones and consequences according to the ArbSchG

Temperature	Consequences
<26°C	–
26°C - 30°C	Additional measures needed, examples in [ASR17]
30°C - 35°C	Advanced measures needed to reduce strain on employees; technical measures are to be favored before person-related measures
>35°C	The room is not suitable for working anymore, unless technical/organisational measures or protective gear is used

### Reduce Energy Consumption While Occupant Is Not Using Their Devices


An occupant has left their work station for his lunch break without turning off their PC, lamps and space heater. The building management system recognizes this (e.g. through pressure and movement sensors) and triggers the planning process. The expected behavioral outputs are: turning off the lamps, putting the PC in sleep mode by making use of software tools such as *Sleepy* [SNN+14] and, depending on the current TOU tariff, turning off the space heater. This means that this scenario is split in half: one where the TOU rates are too high, and one where the rates are acceptable.

A ground plan of the office in the lunch break scenario is provided in Figure 4.2. The objects are similar to the previous scenario, but there is only one desk with a work station and a desk lamp on it.

Human locations can be ignored in this scenario, since the occupant has left the office for his lunch break. Object locations are as follows: the two ceiling lamps are on the ceiling, the window is facing south, the door is facing north. The workstation and desk lamp on the desk. The space heater is on the western side of the office.

The request of this scenario is that unused devices shall be turned off, since the employee has left his office for lunch break. The sole preference is leaving the heater running as long as the TOU rates are acceptable.

The temporal properties are the same as the ones from the previous scenario, substituting AC units for heaters.



**Figure 4.2:** Office two workstations, a heater and several ceiling lamps

### 4.3 Modeling with PDDL

The aforementioned scenarios, classified using the taxonomy from Section 4.1, are modelled using PDDL2.1.

Types of objects involved in the scenario are introduced in the `:types` section of the planning domain. A type hierarchy is introduced as well, e.g. workstations, AC units, etc. are subtypes of the device type:

```
(:types
  device person - object
  workstation hvac lamp - device
  ac heater - hvac
)
```

The domain file defines three numeric functions. These functions are:

- `total-cost`: The expenses created by the actions in the plan

- energy-price: The current energy price in ct/kWh
- office-temp: The current temperature in the office in °C

Behavioral outputs are translated to PDDL actions, where the objects on which actions are performed are included as parameters. Preconditions are usually only the negated form of the desired object state (if a door has to be opened, it cannot already be open), and effects include, aside from the desired object state, the action's consequences for energy consumption and, in some cases, the office temperature. An example for an action is given in the following listing. The only parameter in this action is an AC unit, the precondition is that the unit cannot be turned on already. The effect section is a conjunction of three separate effects: after execution of this action, the unit will be turned on, the total cost will have increased by a certain amount and the office temperature will have gotten down by a certain estimated number of °C (lines 5-8).

```

1  ;turn on an AC unit
2  (:action turn_on_ac
3      :parameters (?unit - ac)
4      :precondition (not (turned-on ?unit))
5      :effect (and (turned-on ?unit)
6                  (increase (total-cost) (* (energy-price) 3))
7                  (decrease (office-temp) 2)
8          )
9  )

```

**Listing 4.1:** PDDL action which represents turning on an AC unit

The aforementioned parts of the taxonomy presented in Section 4.1 are modeled in the PDDL domain file. Spatial properties and behavioral inputs, however, reside in the problem file.

Spatial properties – human and object locations – are modeled using the `:objects` and `:init` sections of a PDDL problem file. Listings 4.2 and 4.3 show how this could be done using the scenarios from Section 4.2. Adjacencies between objects are not considered important for this scenario, e.g. it does not make sense to declare a work station 'near a window' or 'far away from a door' since, according to Figure 4.1, the distances from the work stations to the next window/door are all equal to one another.

```
1 (:objects
2   ceil_lamp1 ceil_lamp2 desk_lamp - lamp
3   ws1 ws2 ws3 ws4 - workstation
4   ac1 ac2 - ac
5   office_door - door
6   office_window - window
7   heater1 - heater
8   p1 p2 - person
9 )
```

**Listing 4.2:** Example for a PDDL objects section

In Listing 4.2, the objects and humans which are involved in the scenario are defined within the `:objects` section. At the end of each line, objects defined at an earlier position in the line are assigned a type, which is specified after a dash (-).

Listing 4.3 shows an excerpt of the `:init` section, in which the spatial properties at the beginning of the scenario are defined: Workstation `ws1` is occupied by person `p1`, the AC unit `ac1` is turned on and the office door is open. Because of the closed-world assumption mentioned in sub-section 3.3.2, predicates which are false do not have to be listed, e.g. the fact that the second AC unit is not turned on is already assumed.

```
1 (:init
2   ...
3   (occupied ws1 p1)
4   (turned-on ac1)
5   (open office_door)
6   ...
7 )
```

**Listing 4.3:** Example for a PDDL init section

Behavioral inputs are modeled using the `:goal` and `:metric` sections of a PDDL file. Requests are translated to a logical formula in the `:goal` section, while preferences are modeled using the `:metric` feature introduced with PDDL2.1. In Listing 4.4, an example for such a modeling is provided. Lines 1-6 show the `:goal` section, which, in this case, consists of a conjunction of three predicates: the work station `ws1` as well as the ceiling and desk lamps have to be turned off at the end of plan execution.

The preference in this scenario is to minimize the `total-cost` function, which is introduced by defining the plan metric accordingly. As of PDDL3, preferences can also be defined, but since this thesis only uses features from PDDL2.1 and below, these PDDL3 preferences are not considered.



```
1 (:goal (and
2     (not (turned-on ws1))
3     (not (turned-on ceil_lamps))
4     (not (turned-on desk_lamp))
5   )
6 )
7
8 (:metric minimize (total-cost))
```

**Listing 4.4:** Example for a PDDL goal and metric section

Energy prices are modeled using a numeric function, which gets assigned a value (in ct/kWh) in the problem definition.



## 5 Evaluation

### 5.1 Experiment Design

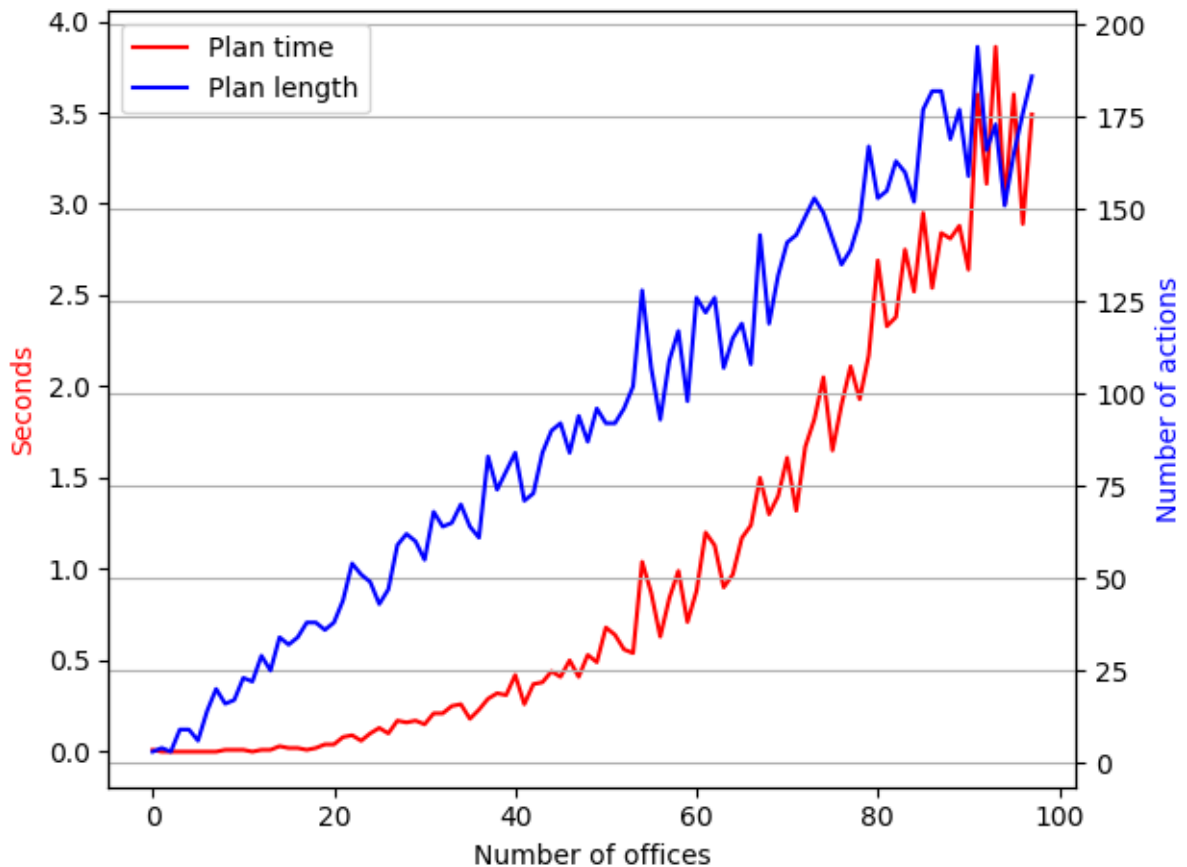
To test the efficiency of the approach, tests with an increasing number of offices were conducted. For testing, a machine with an AMD Ryzen 7 3700U (2.3 GHz) with 16GM of RAM and a 64-bit Windows 10 operating system was used. To show that the approach is efficient, tests were conducted with an increasing number of offices and therefore an increasing number of devices. This is why the test problem files contain models of the scenarios from the previous section with an increasing number of offices, starting at one and going up to 100 offices. The timeout limit was set to 5 seconds, which means that if the planning process took longer than this limit, the test would be considered a failure. Initial office temperatures were chosen randomly between 26°C and 28°C.

Finding a suitable AI planner for the PDDL domain proved to be difficult. Reasons for this are lack of documentation, infrequent updates and lack of support of certain requirements. The AI planner used in this thesis is Metric-FF [Hofb], a variant of the popular Fast Forward (FF) planner [Hofa], a top performer of the Second and Third International Planning Competitions [00; 02]. It extends FF by supporting PDDL's ADL features and numerical state variables.

### 5.2 Results

Test results for the office temperature scenario can be seen in Figure 5.1, where plan generation time and plan length are plotted. Metric-FF was unable to generate a plan for the last two test problem files, 99 and 100 offices, because Metric-FF's internal `MAX_LNF_COMPS` value was exceeded. As can be seen from Figure 5.1, the timeout limit of 5 seconds was never reached, with the longest time being 3.8 seconds. However, the time needed to generate a plan grows in a polynomial fashion, while plan length growth is linear.

Figures 5.2 and 5.3 show results for the lunch break scenario. The testing process was done twice, first with TOU rates that are low and then with TOU rates that are so high that heating units should also be turned off. In this case, plan length grows perfectly linear because the request is to turn all devices off, and the number of devices always increases by five with every office. Much like with the office temperature scenario, plan time grows polinomially, but stays below the timeout threshold: the maximum amount of time for plan generation is 3.7 seconds for peak TOU pricing and 3.79 seconds for off-peak TOU pricing, respectively. In contrast to the office temperature scenario, all 100 use cases were solved.




**Figure 5.1:** Test results for the office temperature scenario.

### 5.3 Discussion

The approach developed in this thesis is able to compute plans for a set of realistic DSM-related planning problems situated in an office building. Plans were generated by a domain-independent AI planner in an acceptable amount of time, that is, in less than five seconds. Plan length growth is (roughly) linear with increasing number of devices, while the time to generate plans can be closely described with  $O(n^3)$ , with  $n$  being the number of offices.

When assessing plan quality, requests and preferences and their satisfaction are the most important indicators. According to Section 4.2, turning on AC units should only be done if the office temperature is particularly high. Looking at the output of Metric-FF, this preference is satisfied, i.e. when solving the scenario with 7 offices: Offices 4, 5 and 7 hold temperatures which are too high for opening windows and doors to be sufficient, therefore AC units should be turned on there. The plan computed by Metric-FF includes turning on AC units in these, and only these offices, since the temperature in the others is low enough to be lowered beyond  $26^{\circ}\text{C}$  by just opening doors and windows. Only turning on AC units when absolutely necessary indicates that plan quality is good.

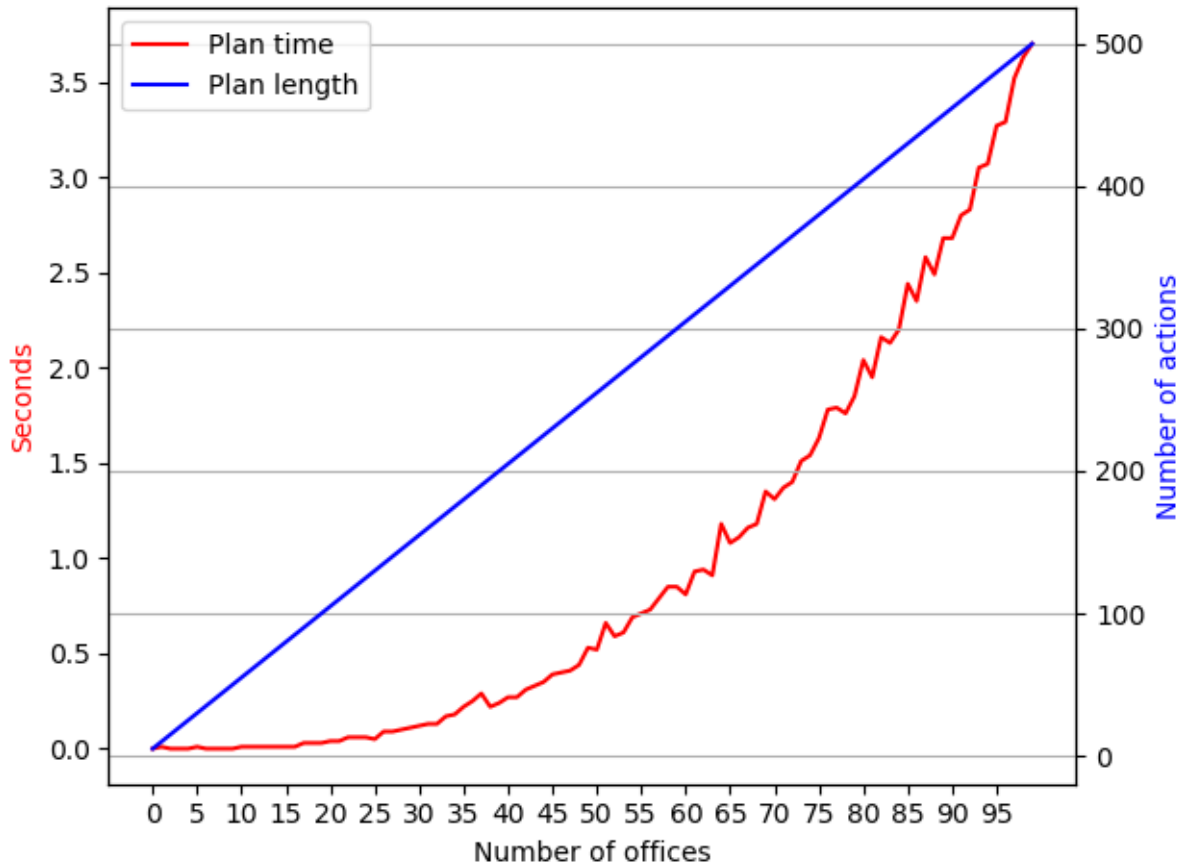


**Figure 5.2:** Test results for the lunch break scenario with non-peak TOU pricing

The potential for energy savings is evident, too, even though quantifying energy savings for the first scenario is difficult: there is, e.g., no information about for how long an AC unit would be turned on until the temperature is low enough again. Assuming a thermostat turning on a 2.5kW AC unit above 26°C and, over the course of a day, leaves it on for an hour, energy savings could be as high as 2.5kWh per day per unit. This means savings as high as 76.075 ct per day per unit, assuming an average electricity price of 30.43 ct/kWh in Germany in 2019 [ENE19] according to the German Economy and Energy Ministry (BMWi). Scaled up to the largest scenario tested – 98 offices with two AC units each –, this equals savings of up to 490kWh/day and 149.10 €/day, respectively.

Potential energy savings are easier to predict in the second scenario. By adding together the hourly energy intake of the ceiling and desk lamps, the work station, and the radiator, and then multiplying that value with the current energy price, one can easily get a precise estimate for the reduction of energy consumption as well as the financial savings. For example, turning off a set-up of

- two ceiling lamps, each with two 18W fluorescent tubes
- a 40W desk lamp
- a 300W radiator



**Figure 5.3:** Test results for the lunch break scenario with peak TOU pricing

- a work station with a 300W power adapter

for an hour in all 100 offices would lead to the office consuming up to 71.2 kWh less than if the devices would have been left running, which in turn would yield financial savings of up to 21.67 € during that hour, assuming the aforementioned average electricity price of 30.43 ct/kWh. With the number of working days in Germany averaging around 250, the savings could be as high as 17,800 kWh which in turn corresponds to 5,416.54 €.

However, the cost of setting up and running a system employing this approach has to be taken into account, too. Such a system would not only include the actuators that would, e.g., open doors and turn on AC units, but also a large number of sensors to accurately recognize events that should trigger the planning process. Moreover, such a system would need to run one or multiple servers which would also increase the overall cost. And all these devices need power of their own, which would reduce net energy savings. If and when the costs of such a building management system can be amortized remains unknown and is an interesting direction for future research.

Code quality is another limiting factor to the relevance of the results: flaws during the modeling process could have resulted in poor PDDL code that prevents optimal plans from being computed.

## 6 Conclusion

### 6.1 Summary of the Thesis

Commercial buildings – office buildings in particular – are characterised by high operational costs and high CO<sub>2</sub> emission levels due to their high energy demand. Intelligent Demand-side Management (DSM) has the potential to reduce both emissions and operational cost. This includes efficiently coordinating the devices inside office buildings, for which Artificial Intelligence (AI) planning is a helpful tool.

This thesis defines an approach for utilizing AI planning for DSM in office buildings. The approach includes formally specifying the state of an office according to a pre-defined classification used in ubiquitous computing environments and then translating this specification into PDDL code, which in turn is used as input to an AI planner.

A study was conducted to test the efficiency of the approach. Results indicated that an AI planner can compute valid and optimal plans in an acceptable amount of time for PDDL domain and problem files produced by applying the developed approach. Additionally, potential for energy savings is evident. When following the actions suggested by the computed plans, energy consumption can be reduced significantly, reducing both the CO<sub>2</sub> emissions the office building is responsible for and the building's operational cost. This means, that the thesis also provides evidence for the potential of DSM regarding energy savings, further improving the understanding of DSM's benefits.

However, it is important to keep in mind that the savings mentioned in this thesis are based on estimates, which is one of its shortcomings: the approach was evaluated via testing with realistic, but fictional scenarios. It has not been tested in a real office building.

### 6.2 Future Work

One direction of future research is the evaluation of the approach in a real office environment in order to further investigate the feasibility of it. This could contribute to a better understanding of the benefits of DSM, and subsequently to a wider adaptation of DSM measures.

Further investigation of the approach is another interesting topic for the future. This includes generalization from the test scenarios to any buildings, analyzing dependencies between offices, proving formally that the generated plans are optimal, or extending the approach to satisfy more than one objective, e.g. both energy savings and comfort for occupants.

Moreover, the trade-off between the potential for energy savings and the cost of installing and maintaining a building management system implementing the developed approach has not yet been explored. Future work could focus on the time needed until the savings such a system provides outweigh the cost of installation and maintenance.

Another interesting direction of future research would be the development of an AI planner specifically for the domain of demand-side management. Since domain-specific planners perform better in their domain than domain-independent planners, this could further improve the efficiency of the approach and contribute to an even better understanding of demand-side management's benefits.



## Bibliography

- [00] 2000. URL: <http://ipc00.icaps-conference.org/> (cit. on p. 27).
- [02] *The Third International Planning Competition*. 2002. URL: <http://ipc02.icaps-conference.org/> (cit. on p. 27).
- [AG19] M. Aiello, I. Georgievski. *Lecture notes: Planning for Pervasive Systems*. Institute for Architecture of Application Systems, University of Stuttgart, May 2019 (cit. on p. 17).
- [AHK+98] C. Aeronautiques, A. Howe, C. Knoblock, I. D. McDermott, A. Ram, M. Veloso, D. Weld, D. W. SRI, A. Barrett, D. Christianson, et al. “PDDL| The Planning Domain Definition Language”. In: (1998) (cit. on pp. 16, 17).
- [ASR17] A. ASR. “Bekanntmachung von Technischen Regeln; ASR V3 „Gefährdungsbeurteilung””. In: *Das Gemeinsame Ministerialblatt (GMBI)* (July 2017) (cit. on pp. 20, 21).
- [Com14] E. Commission. *EU 2020 target for energy efficiency – European Commission*. 2014. URL: [https://ec.europa.eu/energy/topics/energy-efficiency/targets-directive-and-rules/eu-targets-energy-efficiency\\_en](https://ec.europa.eu/energy/topics/energy-efficiency/targets-directive-and-rules/eu-targets-energy-efficiency_en) (cit. on p. 7).
- [Com19] E. Commission. *Energy performance of buildings directive – European Commission*. 2019. URL: [https://ec.europa.eu/energy/topics/energy-efficiency/energy-efficient-buildings/energy-performance-buildings-directive\\_en#energy-performance-of-buildings-standards](https://ec.europa.eu/energy/topics/energy-efficiency/energy-efficient-buildings/energy-performance-buildings-directive_en#energy-performance-of-buildings-standards) (cit. on p. 7).
- [DCB17] D. D’Agostino, B. Cuniberti, P. Bertoldi. “Energy consumption and efficiency technology measures in European non-residential buildings”. In: *Energy and Buildings* 153 (2017), pp. 72–86 (cit. on pp. 7, 12).
- [DG12] M. Doostizadeh, H. Ghasemi. “A day-ahead electricity pricing model based on smart metering and demand-side management”. In: *Energy* 46.1 (2012), pp. 221–230 (cit. on p. 11).
- [E3P20] E. E. E. P. (E3P). *GreenBuilding Programme*. 2020. URL: <https://e3p.jrc.ec.europa.eu/node/192> (cit. on p. 12).
- [EAD+11] M. Economidou, B. Atanasiu, C. Despret, J. Maio, I. Nolte, O. Rapf, J. Laustsen, P. Ruyssevelt, D. Staniaszek, D. Strong, et al. “Europe’s buildings under the microscope. A country-by-country review of the energy performance of buildings”. In: *Buildings Performance Institute Europe (BPIE)* (2011) (cit. on p. 12).
- [ENE19] B. F. W. U. ENERGIE. *Durchschnittlicher Strompreis für einen Haushalt*. 2019. URL: <https://www.bmwi.de/Redaktion/DE/Infografiken/Energie/durchschnittlicher-strompreis-fuer-einen-haushalt.html> (cit. on p. 29).

- [ERH18] S. Eriksson, G. Röger, M. Helmert. “Inductive certificates of unsolvability for domain-independent planning”. In: International Joint Conferences on Artificial Intelligence. 2018 (cit. on p. 14).
- [FA19] L. Fiorini, M. Aiello. “Energy management for user’s thermal and power needs: A survey”. In: *Energy Reports* 5 (2019), pp. 1048–1076 (cit. on p. 11).
- [FL03] M. Fox, D. Long. “PDDL2. 1: An extension to PDDL for expressing temporal planning domains”. In: *Journal of artificial intelligence research* 20 (2003), pp. 61–124 (cit. on pp. 16, 17).
- [GA15] I. Georgievski, M. Aiello. “HTN planning: Overview, comparison, and beyond”. In: *Artificial Intelligence* 222 (2015), pp. 124–156 (cit. on p. 9).
- [GA16] I. Georgievski, M. Aiello. “Automated planning for ubiquitous computing”. In: *ACM Computing Surveys (CSUR)* 49.4 (2016), pp. 1–46 (cit. on p. 19).
- [GL05] A. Gerevini, D. Long. *Plan constraints and preferences in PDDL3*. Tech. rep. Technical Report 2005-08-07, Department of Electronics for Automation ..., 2005 (cit. on p. 19).
- [GNN+17] I. Georgievski, T. A. Nguyen, F. Nizamic, B. Setz, A. Lazovik, M. Aiello. “Planning meets activity recognition: Service coordination for intelligent buildings”. In: *Pervasive and Mobile Computing* 38 (2017), pp. 110–139 (cit. on pp. 7, 9, 19).
- [GNT04] M. Ghallab, D. Nau, P. Traverso. *Automated Planning: theory and practice*. Elsevier, 2004 (cit. on pp. 7, 15).
- [GSK+11] V. C. Gungor, D. Sahin, T. Kocak, S. Ergut, C. Buccella, C. Cecati, G. P. Hancke. “Smart grid technologies: Communication technologies and standards”. In: *IEEE transactions on Industrial informatics* 7.4 (2011), pp. 529–539 (cit. on p. 9).
- [Hofa] J. Hoffmann. *Fast Forward webpage*. URL: <https://fai.cs.uni-saarland.de/hoffmann/ff.html> (cit. on p. 27).
- [Hofb] J. Hoffmann. *Metric-FF webpage*. URL: <http://fai.cs.uni-saarland.de/hoffmann/metric-ff.html> (cit. on pp. 15, 27).
- [HP08] J. Han, M. A. Piette. “Solutions for summer electric power shortages: Demand response and its applications in air conditioning and refrigerating systems”. In: (2008) (cit. on p. 12).
- [KPH06] S. Kiliccote, M. A. Piette, D. Hansen. “Advanced controls and communications for demand response and energy efficiency in commercial buildings”. In: (2006) (cit. on pp. 11, 12).
- [KWLA13] E. Kaldeli, E. U. Warriach, A. Lazovik, M. Aiello. “Coordinating the web of services for a smart home”. In: *ACM Transactions on the Web (TWEB)* 7.2 (2013), pp. 1–40 (cit. on p. 9).
- [Lon03] D. Long. *PDDL Background*. 2003. URL: <https://www.cs.cmu.edu/afs/cs/project/jair/pub/volume20/fox03a-html/node2.html> (cit. on p. 16).
- [LSS12] T. Logenthiran, D. Srinivasan, T.Z. Shun. “Demand side management in smart grid using heuristic optimization”. In: *IEEE transactions on smart grid* 3.3 (2012), pp. 1244–1252 (cit. on p. 9).

- [McD00] D. M. McDermott. “The 1998 AI planning systems competition”. In: *AI magazine* 21.2 (2000), pp. 35–35 (cit. on p. 16).
- [PD11] P. Palensky, D. Dietrich. “Demand side management: Demand response, intelligent energy systems, and smart loads”. In: *IEEE transactions on industrial informatics* 7.3 (2011), pp. 381–388 (cit. on pp. 7, 9, 11, 12).
- [Sch05] U. Schmid. *An Introduction to AI Planning*. 2005. URL: [https://cogsys.uni-bamberg.de/teaching/ws0506/s\\_planning/slides/hsintro.pdf](https://cogsys.uni-bamberg.de/teaching/ws0506/s_planning/slides/hsintro.pdf) (cit. on p. 14).
- [SNN+14] B. Setz, F. Nizamic, T. A. Nguyen, A. Lazovik, M. Aiello. “Sleepy, the Smart Activity Monitoring and Energy Saving Solution for Personal Computers”. In: *Proceedings of the Erasmus Energy Forum 2014*. 2014 (cit. on p. 21).
- [Str08] G. Strbac. “Demand side management: Benefits and challenges”. In: *Energy policy* 36.12 (2008), pp. 4419–4426. DOI: [10.1016/j.enpol.2008.09.030](https://doi.org/10.1016/j.enpol.2008.09.030) (cit. on pp. 9, 12, 13).
- [Vel02] M. Veloso. *PDDL by Example*. 2002. URL: [http://www.cs.toronto.edu/~sheila/384/w11/Assignments/A3/veloso-PDDL\\_by\\_Example.pdf](http://www.cs.toronto.edu/~sheila/384/w11/Assignments/A3/veloso-PDDL_by_Example.pdf) (cit. on p. 16).
- [WKG08] J.-H. Wu, R. Kalyanam, R. Givan. “Stochastic Enforced Hill-Climbing.” In: *ICAPS*. 2008, pp. 396–403 (cit. on p. 15).

All webpages were last visited November 3, 2020.



# Acronyms

**AC** Air Conditioning. 20, 22, 23, 24, 28, 30

**AI** Artificial Intelligence. 3, 5, 7, 8, 9, 14, 15, 16, 27, 28, 31

**ArbSchG** Safety and Health at Work Act. 20, 21

**BMWi** Economy and Energy Ministry. 29

**DR** Demand Response. 12

**DSM** Demand-side Management. 3, 7, 8, 9, 11, 12, 28, 31

**EBNF** Extended Backus-Naur Form. 16

**EU** European Union. 7

**FF** Fast Forward. 27

**HTN** Hierarchical Task Network. 9

**ICT** Information and Communication Technology. 9

**IoT** Internet of Things. 3, 7, 9, 13

**PDDL** Planning Domain Definition Language. 3, 5, 8, 16, 19, 22, 23, 24, 25, 27, 30, 31, 39

**TOU** Time-Of-Use. 11, 21, 27



### **Declaration**

I hereby declare that the work presented in this thesis is entirely my own and that I did not use any other sources and references than the listed ones. I have marked all direct or indirect statements from other sources contained therein as quotations. Neither this work nor significant parts of it were part of another examination procedure. I have not published this work in whole or in part before. The electronic copy is consistent with all submitted copies.

Stuttgart, 03.11.2020

A handwritten signature in black ink, appearing to read 'L. Heilbrunn', written over a horizontal line.

place, date, signature