POLITECNICO DI TORINO

**Master course in**
**Energy and Nuclear Engineering**

Master Thesis

# Modeling large data centers in the framework of an energy system long-term analysis

***The Danish case***

**Supervisor**
Prof. Stefano Paolo Corgnati

**Co-supervisors**
Olexandr Balyk (DTU)
Maurizio Gargiulo (E4SMA)

**Candidate**
Alessandro Colangelo

March 2018

# Preface

This Master's thesis constitutes the conclusion of the two-years course in Energy and Nuclear Engineering at Politecnico di Torino. I have conducted this thesis between October 2017 and February 2018 at the Division of Energy Systems Analysis (ESY) at the Danish Technical University, thanks to the program "thesis abroad on student proposal" financed by Politecnico di Torino.

First of all, I would like to thank Maurizio Gargiulo, head of E4SMA, for connecting me with the Energy Systems Analysis division in Copenhagen and for his precious support. Then, I would like to thank Olexandr Balyk for accepting to be my co-supervisor and Kenneth Karlsson, head of ESY, for considering me part of their working group. I also want to thank the other members of the ESY staff for their warm welcome. In particular, I would like to thank Stefan Petrović for motivating and inspiring me.

I would also like to express my gratitude to Chiara Delmastro, PhD at Politecnico di Torino, for her fundamental advice and for her valuable revision work.

Then, I must thank my family: without their exceptional support I wouldn't be writing this thesis. They gave me the opportunity to attend a full-time university course and they respected and encouraged my will to have an experience abroad. And all this flavored with love.

Last but not least, special thanks go to my dear Chiara for her essential presence and valued advice (particularly graphical) and for continuously stimulating my curiosity.

Turin, 3rd April 2018

Alessandro Colangelo

# Abstract

In the recent years the demand for internet services has experienced a severe increase, which is expected to continue. To be satisfied they need physical facilities known as data centers, whose dimensions are shifting towards large (or "hyperscale") sizes. However, few locations possess appropriate requirements for hosting them. Denmark is certainly one of these: in fact, in 2025 this country will have three fully operative large-scale data centers, and more might come in the following years. Nevertheless, these infrastructures need an elevated amount of electricity: one of the Danish energy service providers, Energinet, expects their annual power consumption to be 8% of the national power consumption registered in 2015. This might dramatically impact on the Danish energy system and delay its transition towards a low-carbon configuration. Hence, this thesis intends to assess the influence of large data centers on two long-term scenarios (the "Frozen Policy" and a target scenario) designed with the Danish bottom-up energy system model TIMES-DK. Both indicate that the Danish energy system is not able to fully resort to renewable power capacity for the additional power demand caused by large data centers until 2050, thus causing a 2% increase of $CO_2$ emissions within the time horizon. Moreover, in the target scenario, which aims at null $CO_2$ emissions in 2050, the electrification of the transport sector is delayed. However, there might be opportunities for recovering the considerable quantity of waste heat released by these facilities.

# Table of contents

# Chapter 1

# Introduction

## 1.1 Background

The need for rapidly reducing greenhouse gas emissions to mitigate the effects of climate change is probably the most relevant challenge we are facing nowadays. Europe has set its target to a reduction of 80-95% by 2050 (compared to 1990 levels) through the shift towards a low-carbon energy system [1]. Among EU countries, Denmark is definitely one of the most ambitious. In fact, its government has declared the intention to reach the independence from fossil fuels by 2050 [2].

Nevertheless, the transition towards a low-carbon society might be slowed down by the increasing demand for internet services, which require physical hyperscale data centers. In 2010 the entire data center industry was already responsible for 1.3% of the world electricity consumption [3], [4], but this value is expected to change dramatically as hyperscale facilities have almost doubled their number in the last five years [5]. Furthermore, optimal locations for large data centers are limited: only regions with a cold climate can guarantee their efficient operation.

Denmark has seized this economic opportunity by facilitating investments in this new sector, as demonstrated by the gradual abolition of the PSO tax (Public Service Obligation) on the electricity consumption [6]. Consequently, three notorious IT companies have recently decided to establish their large data centers in Western Denmark. However, this choice may result in a significant burden for the small Danish energy system, thus affecting its rapid transition towards a sustainable configuration.

## 1.2 Purpose

The purpose of this thesis is assessing the long-term impacts on the Danish energy system caused by the additional electricity demand of the recent large data centers, evaluating at the same time the possibility of reusing their waste heat in district heating networks. Therefore, relevant questions will be addressed:

- Should investments in renewable power capacity be anticipated or may it be necessary to rely on fossil fuels in the near future?
- How will $CO_2$ emissions consequently variate?
- Will the effects be limited to the power and district heating sectors alone?
- May their integration in district heating networks be profitable from a system perspective?

To answer these questions, outcomes from two well-established future scenarios for Denmark will be compared (first including and then removing the presence of large data centers):

- the "Frozen Policy" ($FP$) scenario, which simulates the current energy policies until their expiry; and
- the target scenario "Carbon Free in 2050" ($CF2050$), aiming at null $CO_2$ emissions in 2050.

The preferred tool is thus TIMES-DK, an extensive model of the Danish energy system with a broad perspective on the next 40 years.

## 1.3 State of the art

The existing literature on data centers electricity use is focused either on the estimation of the overall power demand they produce in single countries or on the estimation of the annual electricity that a specific facility may require according to its configuration and external weather conditions. Nevertheless, no studies were found about the impact that elevated data centers power demands may cause on the energy system of a specific region.

Koomey [7] is certainly the forerunner for estimating and predicting the electricity use of data centers in a single country and worldwide. The methodology he set was reused successively for updating his results [8] but it was also recently followed by Shehabi et al. [9] and Stobbe et al. [10] for reports respectively in the U.S. and Germany. It consists in exploiting the data provided by companies that track the installed base of servers and other IT units (network ports, disks, etc.) in a specific country through the shipments of these goods. An energy model is then applied to each category, usually taking into consideration an average utilization factor and technological advances. As a result, an average power draw for each type of component is obtained, which yields to an electricity consumption after the multiplication by the working hours. Although this is a proper bottom-up approach that would be perfectly coupled with bottom-up energy system models such as TIMES-DK, it contains two major drawbacks. On the one hand, data about goods shipments are not easily accessible due to their elevated cost. On the other hand, estimates about future shipments are limited to 5-10 years and they do not match well with the longer time horizon of energy system models (30-50 years).

Instead, a wide variety of models for abstracting the power consumption of a single facility at an hourly and annual level is available in literature. Most of them are centered on a dynamic representation of cooling systems, which may constitute the highest source of inefficiency in data centers. In fact, they all simulate different configurations of free cooling strategies, which aim at exploiting external weather conditions for cooling down data centers, thus reducing the energy consumption. As an example, Ham et al. [11], [12] created a model for a modular data center in Seoul with the purpose of capturing the interaction between the servers and the cooling system operating conditions under different configurations. Similarly, Depoorter et al. [3] created a model for coupling the servers power consumption and heat dissipation with a specific free cooling strategy in a sample facility but analyzing the total annual electricity consumption in various sites. They also attempted to look at the power production mix of these locations, but they did not develop a simulation for assessing the impact on the respective power systems. Gozcu et al. [4] instead simulated the annual power consumption of four diverse free cooling strategies in several places throughout the world, giving less importance to the representation of the IT equipment. A further approach was found in Pelley et al. [13] who attempted to describe the average behavior of a typical data center by means of a parametric model.

However, despite an abundance of examples, no studies explicitly centered on large data centers were found in literature. Nevertheless, they were still a relevant starting point thanks to the modular structure that large data centers usually possess, as further detailed in Chapter 4.

## 1.4 Methodology overview

In accordance with the purpose of this thesis, long-term analyses of the Danish energy system were developed through the creation of a new sector for large data centers (“LDC”) in the bottom-up model TIMES-DK starting from a dynamic representation of their electrical and thermal behavior (Figure

1). A thorough description of the adopted methodology is presented in the following chapters, while this section aims at underlying its major features.

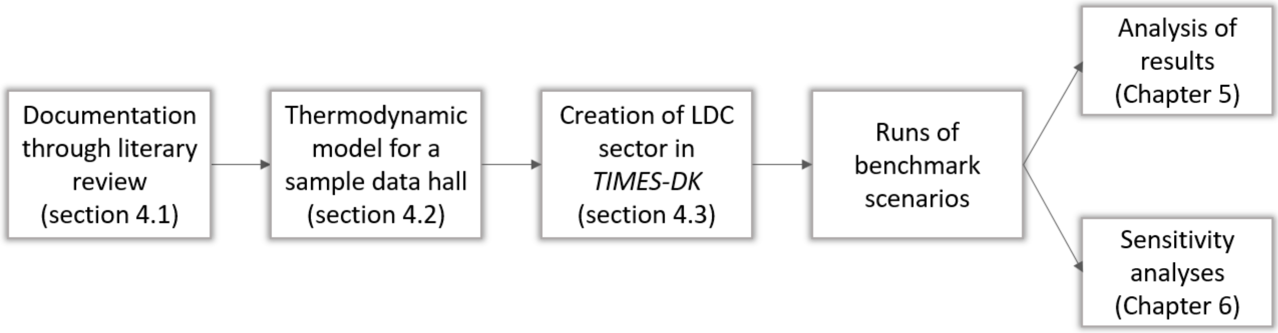


**Figure 1 –** Methodology overview

TIMES-DK deals with the annual levels of the commodities exchanged within the system, but they can be specified with a finer temporal resolution depending on the model's settings. However, TIMES-DK divides a year in only 32 shorter periods (known as "time-slices") while data centers' power consumption is dependent on a daily workload profile and on hourly external weather conditions (because of the cooling system). Furthermore, TIMES-DK requires the description of the entire large data centers' sector.

Therefore, considering that a typical large data center has a modular structure constituted by an ensemble of "data halls", a dynamic model for a reference data hall was first built and its hourly results were then rescaled in space (to fit with the expected installed size of large data centers in Denmark) and aggregated in time (to fit in TIMES-DK temporal resolution). The power required by the IT equipment of the data hall (servers and other working units) was assumed to be equivalent to the thermal load it produces. Therefore, the dynamic model is, actually, a thermodynamic model, whose purpose is to estimate the additional power required by the cooling system together with the available waste heat. Since these two quantities do not only depend on the thermal load produced by the IT equipment, but also on external weather conditions, an annual temperature profile for Western Denmark was taken from the Danish Meteorological Institute [14] as a reference for outdoor climate. The result of this first step is an hourly profile for the power of the cooling system and the waste heat within a reference year.

Secondly, the model for large data centers in TIMES-DK was built on a key element: the projection concerning their expected annual power consumption until 2040 [15]. In addition, the outcomes of the thermodynamic model allowed to characterize the variations in the power consumption of the cooling system and in the available waste heat in each time-slice. Therefore, large data centers are implemented as processes that must satisfy the exogenous power demand taken from [15] which consume the electricity produced by the available technologies in TIMES-DK. Moreover, they require supplementary electricity for their cooling systems but produce an amount of waste heat that can be recovered.

In the end, two scenarios ("Frozen Policy" and "Carbon Free in 2050") were selected for comparing the outcomes before and after the inclusion of this new sector. Furthermore, sensitivity analyses were developed due to uncertainties on the annual power consumption of large data centers and on the power they will require after 2025. The former aspect is addressed considering a daily variable workload profile (instead of a flat one which is the current guess of official sources), while the latter is dealt with the assumption that the number of facilities hosted in Denmark will keep increasing until 2040.

## 1.5 Structure of the thesis

This thesis is divided into eight chapters. The first is an introductory chapter with the aim of summarizing the context, the purpose and the methodology of this thesis. The contents of the succeeding chapters, instead, are listed below.

*Chapter 2* presents the tool used for assessing the impact of large data centers on the Danish energy system: TIMES-DK. First, it contains a brief overview on the panorama of energy system models and then it provides an insight on the family of TIMES models. Finally, it is concluded with the description of TIMES-DK.

*Chapter 3* contains a description of the object of this thesis: the large data centers. First, it delineates their key features (with particular attention to the cooling system and waste heat reuse) and then it outlines the current situation and the future perspectives in Denmark.

*Chapter 4* deepens the methodology that has been followed. First, it presents a simplified representation of large data centers; then, it describes the thermodynamic model and finally it outlines the implementation in TIMES-DK, presenting, as well, the key features of the two scenarios chosen as a benchmark ("Frozen policy" and "Carbon Free in 2050").

*Chapter 5* analyzes the results from the runs of the two scenarios with and without the presence of large data centers. Therefore, it mainly contains comparative analyses regarding the power sector, $CO_2$ emissions, waste heat recovery and marginal prices.

*Chapter 6* is devoted to sensitivity analyses. In particular, the objects of these analyses are: the effects of an assumed daily variable workload profile and the influence of possible additional facilities built after 2025.

*Chapter 7* lists possible further developments, while

*Chapter 8* contains the conclusions.

# Chapter 2

# Energy systems modelling

A wide portfolio of energy system models is available nowadays, even though a well-marked dichotomy was established in the past years based on their approach: top-down or bottom-up. The formers represent an entire economy through a restricted number of aggregate variables and, due to their market-oriented approach, they cannot represent with detail the existing and future technologies in the energy sector. Consequently, they are typically used for analyzing the evolution of energy prices and macroeconomic variables but are not suitable for comparing the effects of different energy policies. Bottom-up models, instead, are technology explicit and focus primarily on the energy sector of an economy: energy-using technologies are specified through their inputs, outputs, unit costs and technical-economic characteristics and they are linked together by the commodities they exchange. Therefore, this latter class of models is particularly suitable for scenario analyses.

Despite this richness, in the existing literature there aren't examples of studies analyzing the interaction between large data centers and the energy system of a country. The reason is twofold: very few countries possess a detailed model of their own energy system with a perspective onto its possible future scenarios and the diffusion of large data center facilities is rather recent and localized in specific areas. However, Denmark is an interesting case study since it possesses both these features: it has a comprehensive bottom-up long-term model for its energy system (TIMES-DK) and is going to host a bunch of new large data centers.

Therefore, this chapter presents the main features related to bottom-up long-term energy system models and characterizes the peculiarities of TIMES-DK.

## 2.1 Model process generation

The first step for creating a comprehensive bottom-up model of a large energy system is to identify the most suitable modelling features for the purposes of the user. The choice is usually the result of a trade-off between various aspects: simplicity, accuracy, manageability and computational costs just to mention the most important.

Then, each model has its own mathematical formulation, but the set of equations commonly requires three types of inputs:

- The technical-economic characterization of the energy system under study, which is dependent on information availability from public or private datasets;
- The model's drivers, which may be either macro-economic parameters influencing end-use demands or may be directly projections of end-use demands;
- The selected scenarios, usually expressed in the form of constraints.

Subsequently, two fundamental steps must be followed in order to guarantee the robustness of the model: these are the verification and validation procedures. The former attests the correctness of the mathematical syntax, the latter proves that the model truly represents the behavior of the system under study.

Finally, the model can be solved with appropriate algorithms and results can be analyzed.

## 2.2 Classification of energy systems bottom-up models

A wide variety of bottom-up energy models is currently employed around the world depending on users' interests. Nevertheless, two broad categories can be identified based on the rationale behind the representation of the system behavior: simulation models and optimization models.

The formers are descriptive models trying to reproduce a simplified operation of the system under study [16]. In fact, they illustrate the response of a system, in terms of possible impacts and costs/benefits, to specific input settings (technical variables and policies). In addition, simulation models can be either static, when they represent the operation of a system in a single period, or dynamic, when the outputs of a period are affected by the outputs of previous periods [16].

On the other hand, optimization models follow a particular objective function to compute the optimal values of all the variables within the system, while satisfying given constraints. In fact, they optimize energy investment decisions endogenously [16]. However, their perspective over a broad time horizon (up to 50 years) implies a coarser temporal resolution with respect to simulation models.

Therefore, both classes of bottom-up models are often used for scenario analyses requiring a thorough description of system technologies. For this reason, TIMES-DK appears to be extremely suitable for the purpose of this thesis.

## 2.3 The case of the TIMES model generator

The Integrated MARKAL-EFOM System (TIMES) is an economic model generator for energy systems at different spatial scales: local, national, multi-regional or global. Provided a technology-rich basis in input, it aims at representing the energy dynamics within the desired system over a multi-period time horizon. Therefore, it is usually applied to in-depth analyses of the entire energy sector, but it may also be used for single sector studies, such as electricity and district heating sector [17]. The Energy Technology Systems Analysis Program (ETSAP), an implementing agreement of IEA, is responsible for the development of this tool.

A TIMES model combines the bottom-up approach with an optimization process (based on linear programming): given estimates of end-use energy service demands and given estimates of existing energy equipment stocks, characteristics of available future technologies and present and future primary energy sources, the model aims at supplying energy services at minimum global cost. This is achieved by making decisions on equipment investment and operation, primary energy supply and energy trade between each region. Constraints imposed by the user, such as environmental constraints, influence the decisions made by the model [17].

All these features make TIMES models suitable for the exploration of possible energy futures based on divergent scenarios, since they lead to a coherent organization of the system under study depending on the assumptions implemented by the user.

### 2.3.1 Time horizon

The time horizon is the temporal extension covered by the model and is divided into a number of time-periods chosen by the user. Each period may have a variable length: it is up to the user to define how many years they contain. The year in the midst of a period is called “milestone year” and TIMES variables, such as capacities and flows, evolve linearly between successive milestone years [17].

The initial period is usually a single past year, over which the model has no freedom because the quantities of interest are fixed by the user to their historical values. In fact, the model needs to be calibrated to standard energy statistics in order to represent as accurately as possible the energy system under study.

In addition to time-periods, there are user-defined time divisions within a year: the “time-slices”. For instance, a typical classification distinguishes between seasons, weekdays/weekends and day/night

(Figure 2). These subdivisions are relevant when the mode and cost of production of an energy carrier changes during a year. Nevertheless, it should be noted that time-slices are still too aggregated in time to fully capture possible mismatch problems between demand and supply that may arise during a year (e.g. in the power grid).

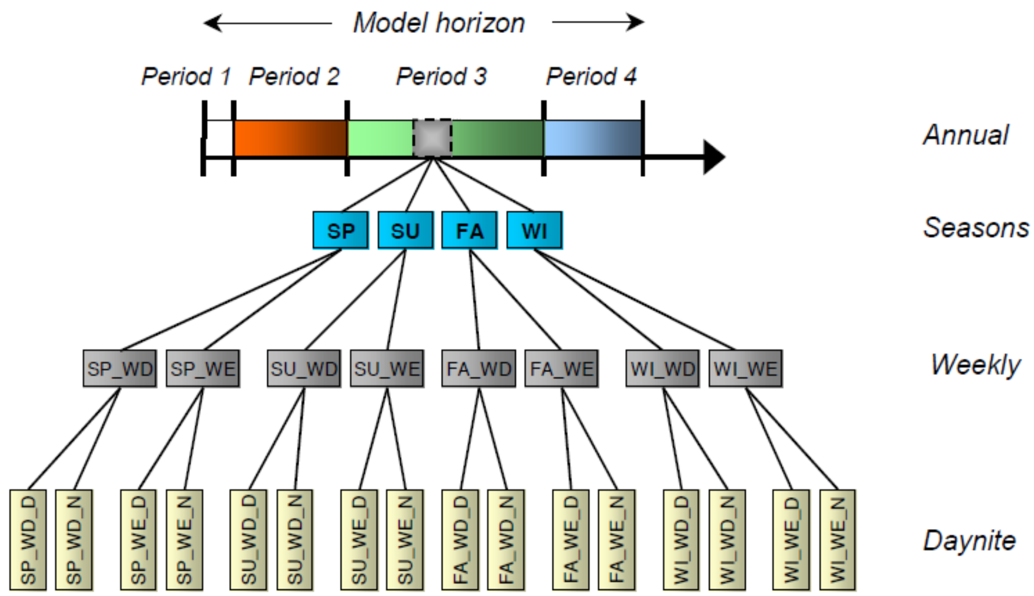


**Figure 2 –** Typical time-slice classification in TIMES models [17]

### 2.3.2 Elements of the Reference Energy System (RES)

A crucial concept in bottom-up energy modelling is represented by the Reference Energy System (RES), which is the graphical representation of the relationship between all the elements of a broad energy system. In TIMES models the RES is always constituted by a set of three entities:

- *Commodities*, which are energy carriers, energy services, materials, monetary flows and emissions;
- *Processes*, which are the technologies able to produce, consume or transform the commodities. For instance, they may be mining or import/export processes (primary sources of commodities), or transformation activities such as power plants and refineries, or end-use demand devices, such as cars and heating systems, which transform energy into a demand service;
- *Commodity flows*, which are the links between processes and commodities. Thus, they have the same nature of their respective commodities, but they are attached to processes.

Depending on the technological detail and on the extension of the energy system under study the RES may be extremely complex to visualize, but its concept is rather simple as shown in Figure 3, which depicts a portion of a hypothetical RES containing a single energy service demand, "House space heating" [17].

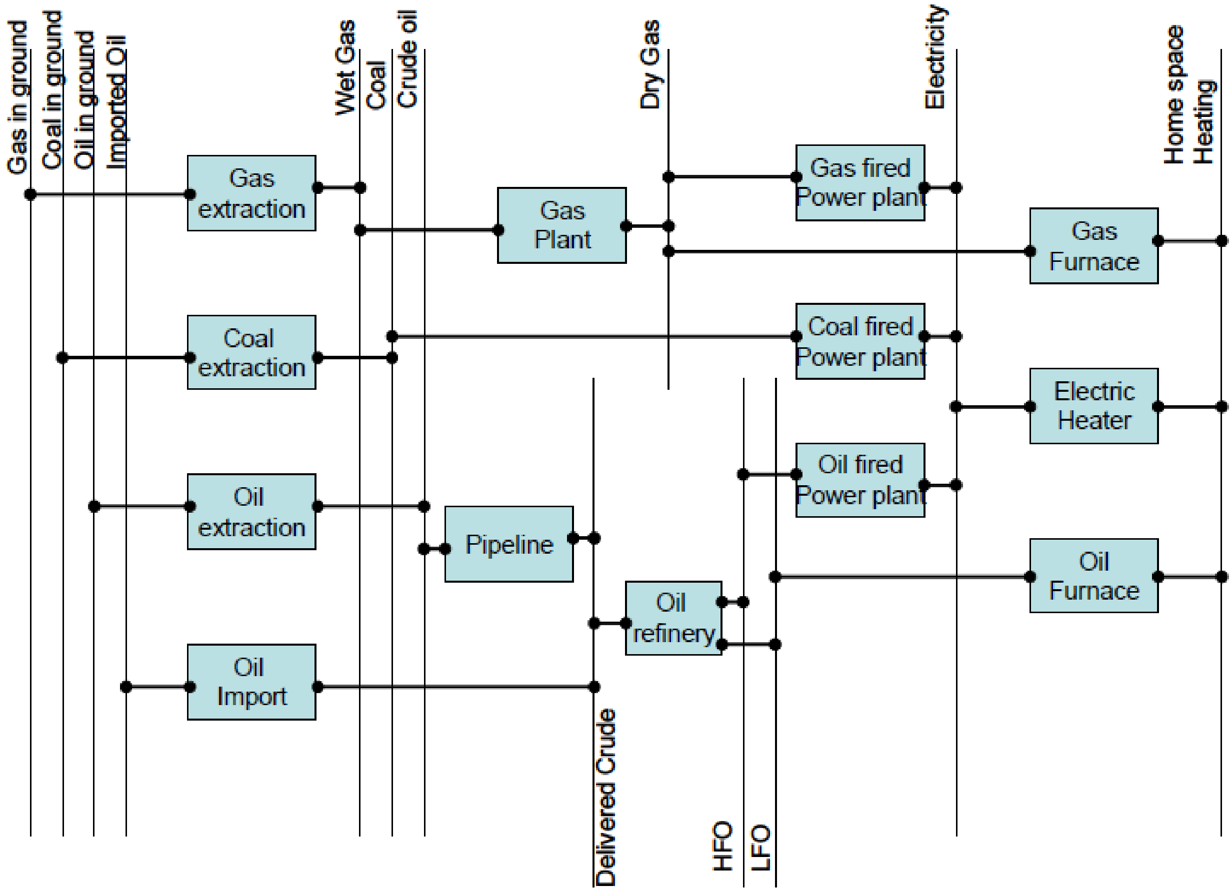


**Figure 3 –** Sample portion of a RES [17]

However, the level of detail at which the energy system is represented does not depend only on the quantity of processes, commodities and flows inserted, but also on the *attributes* chosen to characterize them. A wide range of parameters is available and each of them activates a precise equation when specified. These attributes may be of different nature: technical, economic, policy parameters, bounds or advanced parameters (such as the *vintage* characterization). Moreover, some factors can be also attached to the entire RES: currency conversion values, region-specific time-slice definitions, region-specific general discount rate and reference year for calculating the discounted total cost (the objective function) [17].

### 2.3.3 The economic rationale behind TIMES

As already mentioned, TIMES models return the configuration of the energy system under study over a time horizon as a result of an optimization problem whose objective is supplying energy services at the minimum global cost, or, equivalently, whose objective is maximizing the total economic surplus (the sum of all suppliers' and consumers' surpluses). This is obtained through the computation of an equilibrium since, according to the Equivalence Principle, when all markets are in equilibrium the total economic surplus is maximized [17].

However, TIMES models focus their attention only on the energy market and neglect the other economic markets. For this reason, they can be considered *partial equilibrium* models. In general terms, "a market is said to have reached an equilibrium at prices p* and quantities q* when no consumer wishes to purchase less than q* and no producer wishes to produce more than q* at price p*" [17].

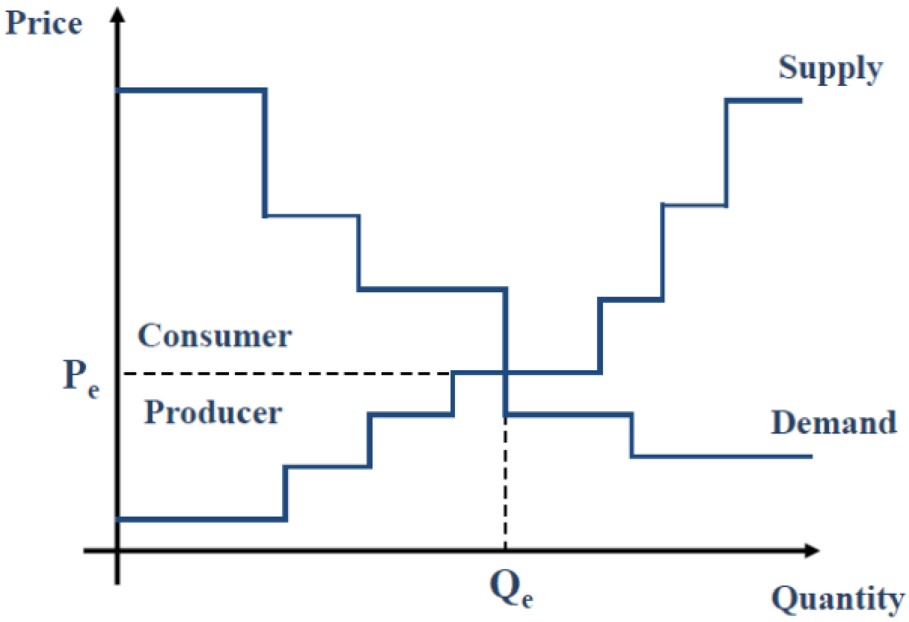


**Figure 4 –** Example of supply and demand curve in TIMES models [18]

In TIMES, the suppliers of a commodity are the processes able to produce that commodity, while the consumers are those processes utilizing it. In microeconomics, the set of suppliers can be represented through the *inverse production function*, which plots the marginal production cost of a commodity as a function of the quantity supplied. In TIMES, as in other linear optimization models, this curve is entirely determined endogenously: each technology has its own marginal cost to produce a specific quantity of a commodity. Therefore, the model establishes a merit order for satisfying the demand according to the cheapest technologies. That is why the supply curve has a step-wise shape: each step corresponds to a different technology. On the other hand, the demand curve is implicitly built if the commodity is an energy carrier (which is endogenous in the model). When the commodity is an energy service, instead, the modeler has the task to define it through the specification of elasticity factors (Figure 4).

Summarizing, a supply-demand equilibrium has its economic rationale in the maximization of the total surplus, which is the objective pursued by TIMES models. However, the mathematical method used to obtain it must be adapted to the mathematical properties of the model, which for TIMES are the following [17]:

- Linear relationship between inputs and outputs of a technology. This means that a technology may be implemented at any capacity, but allows adopting Linear Programming techniques, which reduce the computational time;
- Total economic surplus is maximized over the entire time horizon;
- Energy markets are competitive with perfect foresight. Competitive markets are characterized by perfect information, which is extended to the whole planning horizon in TIMES, thus each agent has complete knowledge of the market's parameters, present and future.

Eventually, as a result of these assumptions, the market price of each commodity is equal to its marginal value in the overall system [17].

## 2.4 A relevant example: TIMES-DK

TIMES-DK is one application of the TIMES model generator to represent the energy system of a country. TIMES-DK includes the complete Danish energy system and covers investments throughout the whole modelling horizon. It is the result of the collaboration between the Danish Energy Agency (DEA), the Technical University of Denmark (DTU) and E4SMA [19]. So far it has been used for analyses in the residential heating sector and in transports. Petrović et al. [20] assessed the importance of an improved modelling of residential heat pumps in the Danish energy system until 2050 through TIMES-DK. Then, Karlsson et al. [21] used it for proposing the best solution (from an energy system perspective) between heat pumps, biomass boilers and district heating for supplying heat to a housing community near Roskilde. Finally, Tattini et al. [22] exploited TIMES-DK to suggest efficient decarbonization measures in the transport sector, trying to replicate the human behavior in transport mode adoption. Moreover, ongoing projects aim at utilizing it for supporting and guiding the decisions of Danish policy makers.

From the point of view of the structure, TIMES-DK contains two communicating regions (Denmark East and Denmark West) which are also linked to the neighboring countries through interconnectors. Then, a further geographical classification for certain processes and commodities is specified considering the location of district heating areas (Figure 5).

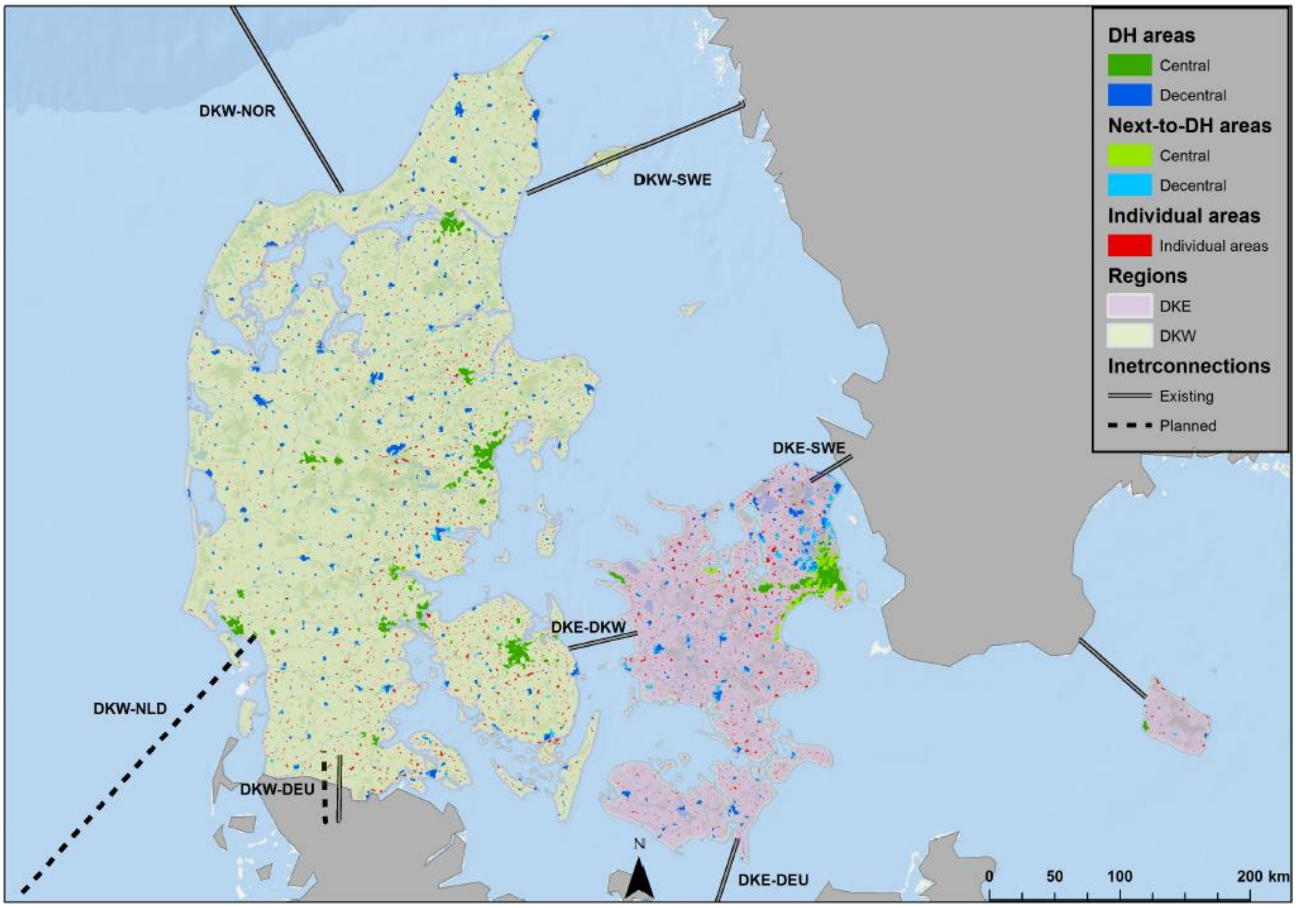


**Figure 5 –** Geographical classification in TIMES-DK [19]

Furthermore, the model groups the entire energy system into five main sectors: supply (SUP), power and heat (ELC), industry (IND), residential (RES) and transports (TRA). The former sector regards the primary energy commodities, which can be imported, exported or domestically extracted. Then, conversion technologies (usually belonging to ELC sector) transform the primary energy commodities in secondary energy commodities, which are needed by end-use technologies to satisfy the end-use sector service demands (i.e. for industry, residential and transport sectors) [19].

As far as time is concerned, instead, TIMES-DK is calibrated for 2010, which is the model base year, and has a modelling horizon until 2050. The time horizon is further subdivided into shorter periods (1-5 years) and each year is the sum of 32 non-sequential time slices, representing seasonal (4 seasons), weekly (working/non-working days) and daily variations. On the daily level, TIMES-DK adopts a peculiar classification: every hour of the year is organized into four categories, as shown in Figure 6 [19]:

A) "High wind production – low power demand";

B) "High power demand – low wind production";

C) "No photovoltaics (PV) production";

D) "Rest".

This subdivision aims at capturing the critical situations for the power system, i.e. when a mismatch is present between the power demand and the power production from variable renewable energy sources. However, it is important to notice that not all modelled parameters are defined at this level of detail.

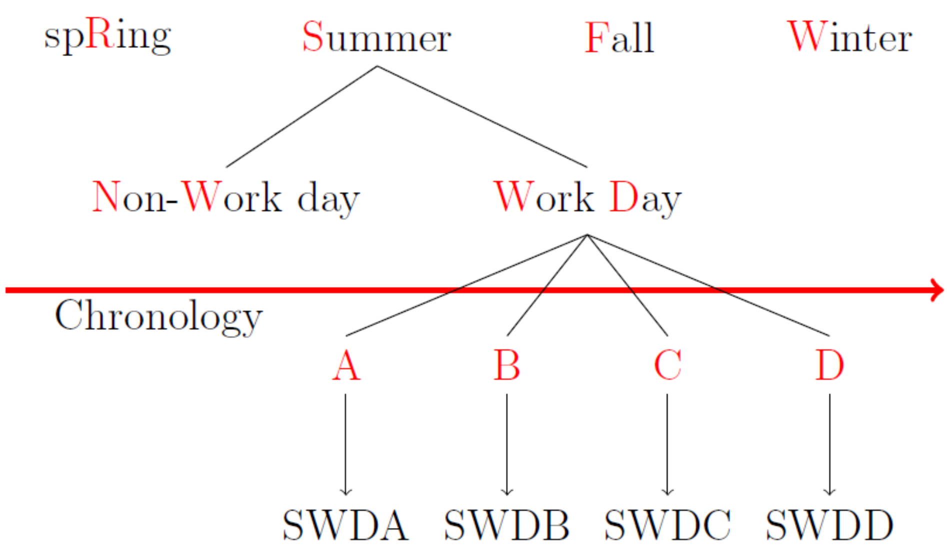


**Figure 6 –** Time-slice classification in TIMES-DK [19]

To summarize, TIMES-DK is a powerful tool for analyzing medium and long-term evolution of the Danish energy system under a set of constraints. In fact, thanks to the inclusion of the most important sectors of the economy, it allows to analyze the relationship between the supply-side and the end-use sectors from a system perspective. In addition, its modular structure allows implementing model expansions in a simple way, especially in terms of new technologies and sectors. Therefore, these features make TIMES-DK a suitable tool for assessing the impact of large data centers on the Danish energy system.

# Chapter 3

# Overview on large data centers

This chapter introduces the key features of large data centers through the comparison with small and medium-size ones. Particular attention is devoted to their equipment, cooling strategies and possibilities for reusing waste heat. In conclusion, current and future developments of this industry in Denmark are delineated.

## 3.1 Definition and purpose

Barroso [23] defines data centers as "buildings where multiple servers and communication gear are co-located because of their common environmental requirements and physical security needs and for ease of maintenance".

However, there is a fundamental difference depending on their size: small and medium data centers typically host a large number of small applications, each running on dedicated hardware infrastructure, which is separated from other systems in the same facility. Large data centers, instead, run a small number of large applications (Internet services) and rely on homogeneous hardware and software platforms [23].

This substantial difference depends on the owner of the facility: small/medium data centers usually belong to companies that rent their servers or applications to other organizations, while large data centers, also known as "warehouse-scale computers", belong to companies such as Google, Amazon, Facebook, Apple and Microsoft, which need huge infrastructures to power their online services offered worldwide [23].

Therefore, large data centers, which are the focus of this thesis, display the same components of the smaller ones, but require bigger facilities and are operated differently. Consequently, from an energy point of view, the latter usually do not exceed 10 MW of installed IT capacity [24], while large data centers may even require more than 100 MW overall.

## 3.2 How large data centers work

As previously mentioned, data centers have always the same task, independently on their size: offering IT services through their servers, which, in turn, need to be kept at an appropriate temperature.

Therefore, there are two functional units always present: the IT equipment and the cooling system. The former is responsible for the majority of the power consumption within the facility, while the latter has the task to maintain the correct conditions in the working environment, causing an overhead on the facility power draw. However, new facilities are becoming very efficient. In fact, recent designs have tried to integrate these particular buildings with the surrounding environment through two actions: the implementation of "free cooling" strategies and the reutilization of waste heat.

### 3.2.1 IT equipment

Servers are the core of data centers: they are in charge of conducting the computational operations. In addition, two other fundamental components join them: memory disks and network ports. The former allows storing information; the latter allow its transmission. These three elements are essential for data centers and are always present.

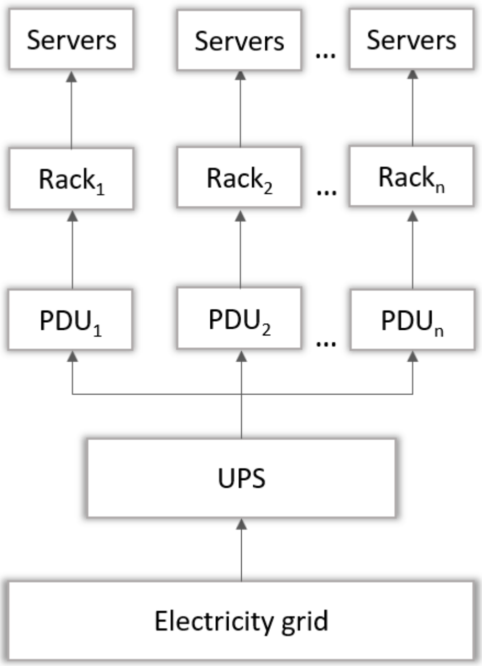


**Figure 7 –** Representation of a typical IT equipment in data centers [12]

However, they cannot be connected directly to the electricity grid for two main reasons: on the one hand, a buffer is needed in case of power shortage; on the other hand, the supply voltage should be lowered to meet the components specifications. Consequently, two additional devices are usually needed: Uninterruptible Power Supply systems (UPSs), which are batteries, and Power Distribution Units (PDUs), which are electronic tools able to modulate current and voltage. In general, a PDU supplies a group of servers, known as “rack”, while a UPS supplies several PDUs (Figure 7), even though the architectural design is flexible and other configurations may be adopted [23].

Summarizing, the set of these components can be referred as “IT equipment”, a term that thus encloses all the power-consuming units within a data center except from cooling system and lighting.

### 3.2.2 Cooling system and cooling strategies

Cooling the servers and the other units is one of the most important tasks the facility has to guarantee. Typically, closed loop systems are the preferred solution: a circuit of air removes the heat from the IT equipment and transports it to a heat exchanger. A consolidated solution for the data center’s air circuit is the adoption of raised floors in combination with hot aisle containment, which reduces the risk of mixing between hot and cold air. Units known as CRAHs (Computer Room Air Handlers) or CRACs (Computer Room Air Conditioners) then grant recirculation and fresh air supply (Figure 8) [23].

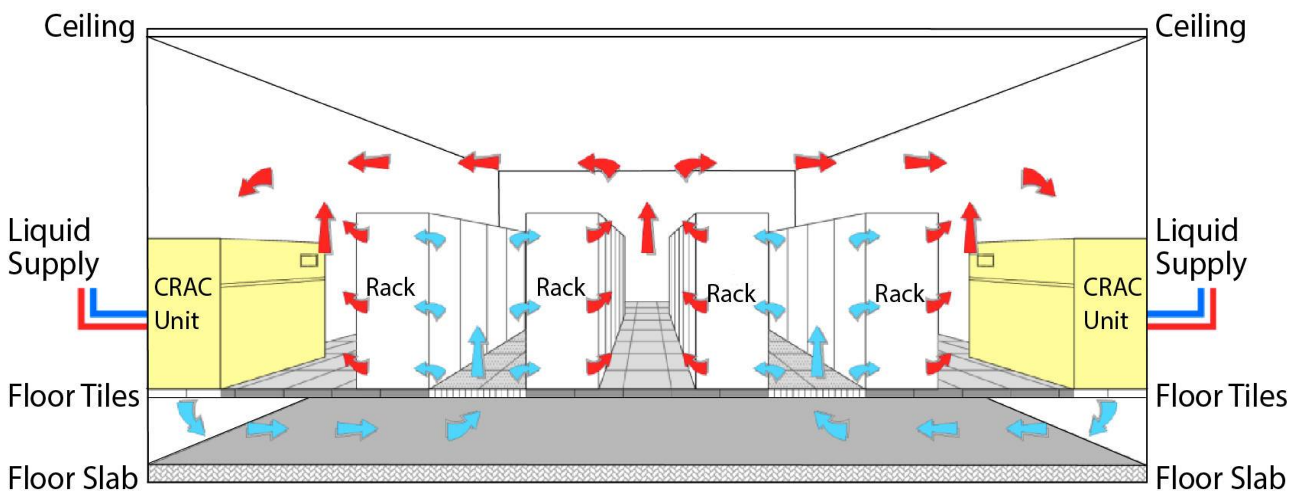


**Figure 8 –** Typical solution for cooling data centers [23]

The ways in which the heat is then rejected outside are various, but a common approach is to employ a vapor-compression chiller plant, which allows to supply chilled water to CRAH units (Figure 9a). Unfortunately, this solution may be synonym of great inefficiency, as these cooling systems alone could consume up to 50% of the total power draw [25]. For this reason, new data centers also implement effective "free cooling" strategies, which are based on the utilization of outdoor air whenever possible. Consequently, the cooling system electricity consumption drops to about 10% of the total facility consumption [26], [27] because only fans are needed for most of the time, but the optimal locations are restricted to colder regions.

Free cooling strategies are not univocal and different solutions are available. Nevertheless, the substantial difference is between a *direct* or an *indirect* use of outdoor air. The former option puts in direct contact the servers with the external air (Figure 9b), while the latter exploits heat exchangers to cool down the air inside the data center (Figure 9c, Figure 9d). The advantages in terms of energy consumption are evident, but there are some drawbacks to consider: direct air economization implies a more difficult control of contaminants, while indirect air economization requires higher investment costs, since the chiller plant is still needed for backup (Figure 9b, Figure 9c, Figure 9d).

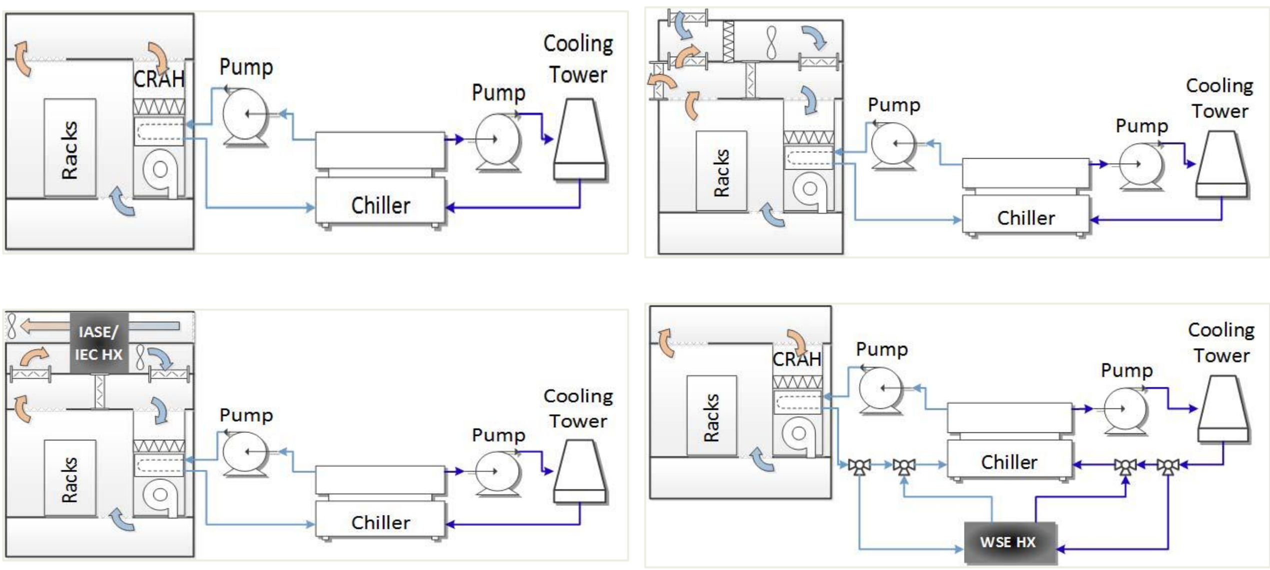


**Figure 9 –** Top left case (a) does not contain air economization. Top right case (b) represents direct airside economization. Bottom left picture (c) shows indirect airside economization (IASE). Bottom right picture shows water side economization [4].

### 3.2.3 Possibilities for waste heat recovery

Almost all the power absorbed by the IT equipment is finally converted into heat, mostly due to the dissipation in the CPU circuits. Therefore, a relevant stream of waste heat flows out of data centers through the cooling system. Unfortunately, this thermal energy has a very low quality, but the recent climate challenge has risen the interest in capturing and reusing it [28].

According to Ebrahimi et al. [28], there are several available techniques considering data centers' operational thermodynamic conditions: district/plant/water heating, absorption cooling, direct power generation (piezoelectric and thermoelectric), indirect power generation (steam and ORC), biomass co-location and desalination/clean water. However, the companies that are going to establish their large data centers in Denmark have already expressed their interest in reusing the excess heat for district heating purposes [29], [30], thus only this option is briefly discussed.

The most effective heat recovery for district heating comes from liquid cooled servers, where captured temperatures are in the range of 50-60 °C [28]. However, the design of these cooling systems integrated inside the processors is complex and expensive. Moreover, advances in the technology of heat pumps allow reusing low-grade heat with still an acceptable efficiency of the pump, which means a COP between 3 and 4 even when the heat source is around 20 °C and the heat demand is between 60 and 90 °C [31]. Consequently, also air-cooled servers producing a stream with a temperature of 30-45 °C become interesting for district heating purposes.

## 3.3 The industry of data centers in Denmark

### 3.3.1 Current situation

Historically, Denmark has not been a particularly attractive country for the data center industry for two main reasons: the size of data centers and the lower relevance of the service sector with respect to other European countries. In fact, as reported by Avgerinou et al. [32], the average power size of the current data center stock in Europe is around 2 MW, a dimension that makes their integration in an urban context extremely suitable. Consequently, most of the companies decided to build their data centers in those highly urbanized cities whose economy is mainly based on offering services, such as London, Paris, Milan and Amsterdam.

This situation is also confirmed by the website "datacenter map" [33], which indicates that Denmark possesses 29 colocation facilities, a number significantly smaller compared to other countries: 244 in the United Kingdom alone, 185 in Germany, 142 in France, 91 in the Netherlands and 62 in Italy. Moreover, half of them are located in Copenhagen, a further demonstration of the fact that these types of data centers require a relatively large urban context.

Furthermore, also their power consumption in Denmark is in line with what has been presented so far: according to the energy matrices of Denmark Statistics [34], the subsectors indicated as "*Computer programming, consultancy services relating to IT and similar activities*" and "*Information services*" consumed together 240 GWh of electricity in 2015, corresponding to 0.53% of the whole national power demand in the same year.

### 3.3.2 Potential and future investments

Recent tendencies in the IT sector are driving a shift in the size of data centers all over the world: from small/medium scale to hyperscale [9], [10]. Consequently, this new typology cannot be fitted in an urban framework but necessitates plenty of space just outside the cities, which Denmark is ready to offer thanks to the availability of large sites at a competitive price [35]. However, location availability is not the only reason that makes Denmark one of the most attractive countries for future investments in large data centers, there are other key factors [35]:

- Political, environmental and economic security and stability;
- High availability of the power grid (99,997% grid uptime);
- Possibility of being fed by 100% renewable energy (thanks to a high renewable mix in the power production and to green certificates);
- Competitive electricity prices (thanks to the recent abolition of PSO tax);
- Low latency and excellent fiber connectivity with both Central Europe and the United States;
- Energy efficiency through free air cooling strategies fostered by the cool climate;
- Possibility to reuse waste heat through its injection in district heating networks;
- Speed in obtaining building permits.

In order to exploit these features, the Danish Ministry of Foreign Affairs has instituted "Invest in Denmark", a national promotion agency with the task of attracting capital also in the data center industry. As a result of its mediation and thanks to the favorable conditions that the country offers, three well-known companies decided to invest in Western Denmark. Here is a list of the projects:

- Apple is building two hyper-scale data centers close to Viborg and Åbenraa;
- Facebook chose Odense for having a new data center online by 2020;
- Google bought two large-scale sites close to Fredricia and Åbenraa (but without announcing any plans for the near future).

However, despite their positive impact on the national economy, they may constitute a significant burden for the Danish power system. In fact, according to the estimates made by Energinet [15], their annual power consumption is predicted to be 3,6 TWh when fully operative in 2025, which would correspond to 8% of the national power consumption registered in 2015.

Therefore, this thesis is focused on the effect that large data centers will have on the Danish energy system precisely because their additional power demand is not negligible and rather near in the future.

# Chapter 4

## TIMES model for large data centers

This chapter introduces the adopted methodology for modeling large data centers in a TIMES environment. Three steps have been followed and they are reflected in the structure of this chapter. First, a simplified representation of data centers' thermal and electrical behavior was identified (4.1), then a thermodynamic model was created to capture variations in the power consumed by the cooling system and in the waste heat (4.2), finally the implementation in TIMES-DK was developed based on the outcomes of the thermodynamic model (4.3).

### 4.1 Simplified representation of large data centers

The existing scientific literature still lacks models focused on data centers of large dimensions (50-150 MW), while it is rich in models centered on small/medium data centers. Nevertheless, large data centers often display a modular structure and the whole facility is subdivided into "data halls", whose size is comparable to small/medium data centers. As an example, the large data center under construction in Viborg is constituted by a series of data halls of 4 MW [29].

Therefore, the first important assumption that has been adopted regards the scalability: the behavior of a single data hall, based on models present in literature, is considered representative for the whole large facility. This is also consistent with the implementation in TIMES-DK, where the most important inputs are two adimensional fractions: the power consumption of the cooling system and the amount of waste heat per unit power consumption of the IT equipment in each time-slice (see Figure 26 and Figure 27).

The working principle of the reference data hall is rather simple (see 3.2): the IT equipment withdraws power from the grid releasing heat, which is removed from the facility through the cooling system, whose power consumption is also dependent on external weather conditions. These affect both the amount of outdoor air flow through the fans (when free cooling strategies are adopted) and the COP of chiller plants. A schematic view of the adopted concept is presented in Figure 10, based on the sketch introduced by Depoorter et al [3].

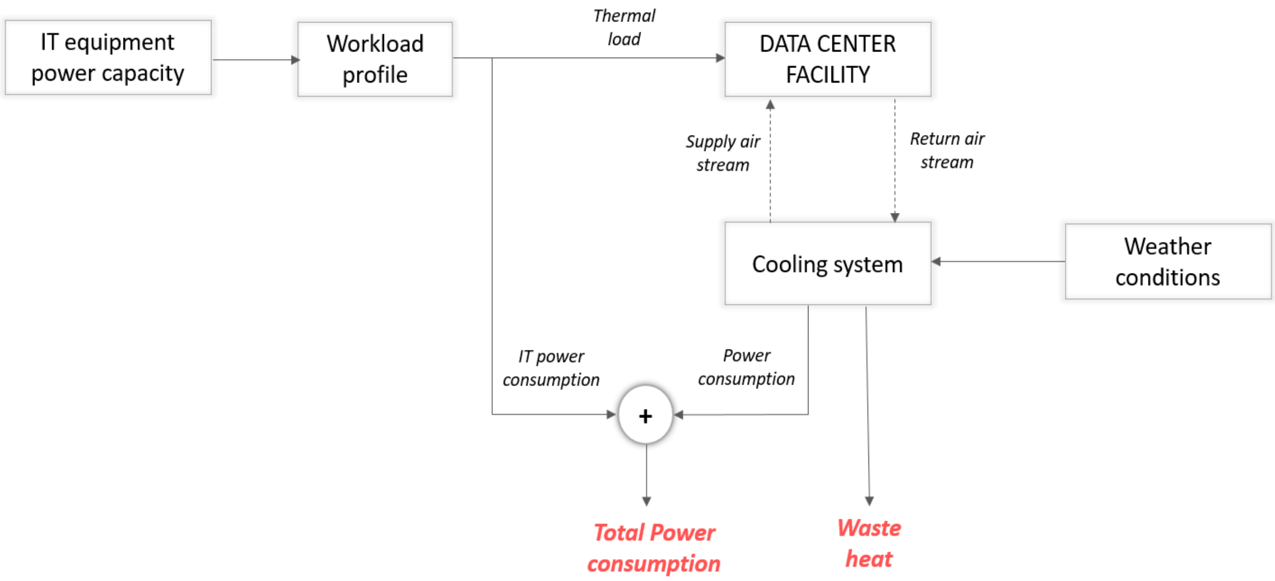


**Figure 10 -**Sketch for data center simplified representation

The IT equipment of the sample data hall is not subdivided into its components, but it is simply represented by the overall amount of power it requires. This choice is due to uncertainties in the configuration of these data halls. This implies that management strategies, such as consolidation, cannot be captured by this model or, equivalently, that the load is evenly distributed among the equipment. Nevertheless, modelling the IT equipment as a single entity requiring a specific amount of power is still relevant, particularly for the interaction with the cooling system.

Another important assumption regards the estimation of the thermal load, which is considered equal to the amount of power consumed by the IT equipment [3], [11], as presented in Figure 11. This hypothesis also confirms that a thorough characterization of the IT equipment is unnecessary for this study. In addition, the internal heat production within the facility is predominant: therefore, heat transfer through the building envelope is neglected [3].

Summarizing, all the simplifying assumptions done in the model are listed below:

- Perfect scalability: the behavior of a single data hall is representative of the whole data center;
- The IT equipment is considered a unique entity and it is not declined into its components, which also implies an even load distribution among the equipment;
- The power consumed by the IT equipment is entirely converted into heat;
- Heat exchanges through the building envelope are negligible compared to internal heat production due to the IT equipment.

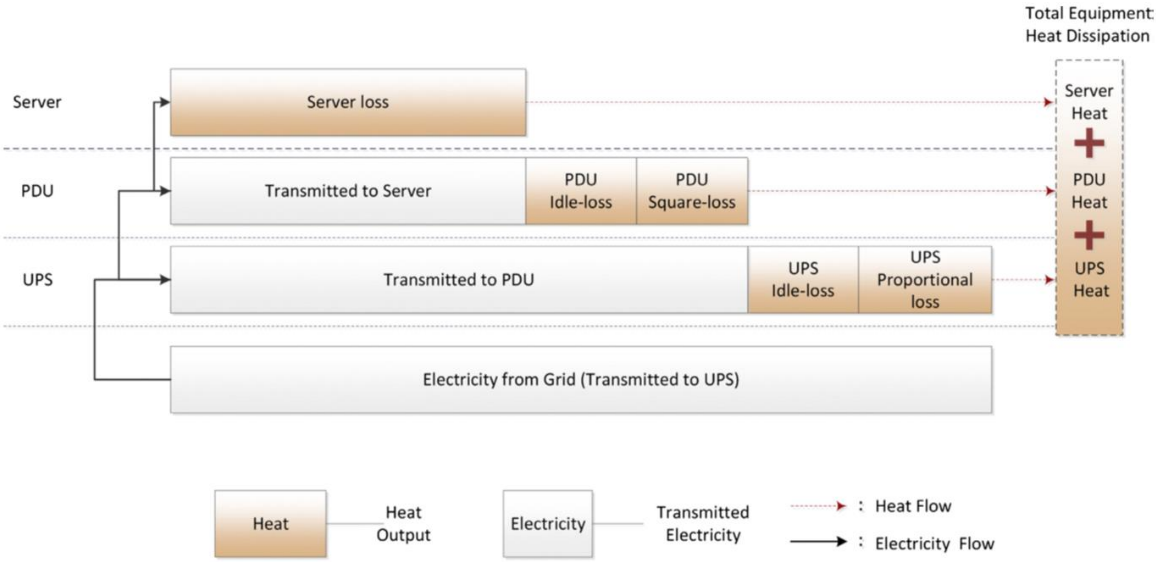


**Figure 11 –** Heat and power flow through IT equipment [11]

### 4.1.1 Workload profile

Since large data centers gained attention in recent times, little information about their utilization profile is available and the companies that already own them worldwide are reluctant to share data about daily power consumption. Therefore, a common hypothesis is to consider a flat daily load profile [15], [29]. As Figure 12 shows, this is clearly a prudent choice, but it is the only one that official sources have considered so far. Consequently, it is taken as a reference in this thesis.

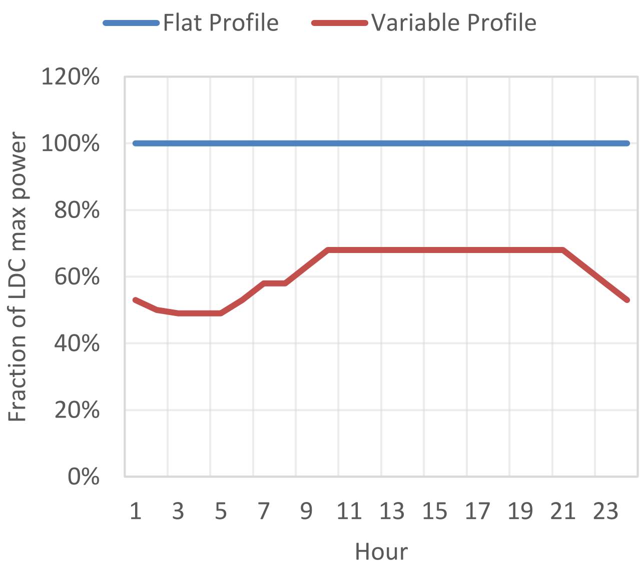


**Figure 12 –** Flat workload profile assumed for main results [15], [29] and variable profile assumed for sensitivity analyses [24]

However, the purpose of these large data centers is selling internet services to European customers (see 3.1). Consequently, it is likely they have a workload profile variable within the day, depending on the customers' requests. Therefore, a "web workload profile" [24] is also considered as an alternative to the flat profile.

## 4.2 Thermodynamic model for cooling system and waste heat

A thermodynamic model for the sample data hall is necessary for estimating hourly variations in the power consumed by the cooling system and in the available waste heat during a reference year, since these quantities strongly depend on external weather conditions. Although TIMES-DK works at a more aggregate time scale (a year is divided into 32 time-slices), a thermodynamic model with an hourly resolution is still meaningful when outcomes are aggregated at the time-slice level. In fact, without this procedure TIMES-DK would simply allocate the annual value of these quantities to the time-slices in proportion to their temporal length. It should be noticed that the aggregation to the time-slice level is possible thanks to the classification of each hour of a year according to the principles described in 2.4: a seasonal subdivision (winter, spring, summer, fall), a weekly subdivision (working days, non-working days) and a daily subdivision ("A", "B", "C" or "D").

Among all the assumptions, two main inputs characterize this thermodynamic model: the outdoor air temperatures and the IT load. Outdoor temperatures are obtained as an average of the hourly values in Western Denmark in 2010 [14], which are already used in TIMES-DK for varying the COPs of air heat pumps. The IT load, instead, is responsible for the thermal load on the data hall and may be flat or variable within the day as explained in 4.1.1. The remaining thermodynamic assumptions are summarized in Table 1.

The assumed configuration of the cooling system consists in a vapor compression air cooled chiller associated with an indirect airside economizer (IASE), as shown in Figure 9b. The amount of air recirculating through the IT equipment is calculated imposing a temperature rise of 15 °C in eq.(1):

$$V_{IT} = \frac{\dot{Q}_{IT}}{\rho c_P \Delta T_{IT}} \tag{1}$$

However, the airflow supplied to the data hall should not be the exact amount required to cool the equipment due to by-pass effects [4], [24]. Therefore, the total airflow recirculating inside the data hall is expressed by eq.(2):

$$V_{DC} = \frac{V_{IT}}{1 - BP} \tag{2}$$

**Table 1 –** Key assumptions for the thermodynamic model

| Nomenclature | Units | Value | Description |
|---|---|---|---|
| $c_p$ | MJ/(kg*K) | 0.001 | Specific heat of air |
| $\rho$ | kg/m$^3$ | 1.2 | Air density |
| $Q_{IT,max}$ | MW | 4 | Max power consumed by IT equipment [29] |
| $\Delta T_{IT}$ | °C | 15 | Temperature increase across IT equipment [3] |
| $T_{DC,sup}$ | °C | 20 | Data center supply temperature [24], [36] |
| $\Delta T_{PP}$ | °C | 3 | Minimum allowed temperature difference in heat exchanger [24] |
| $T_{OA,max}$ | °C | 17 | Maximum outdoor temperature for free cooling ($T_{OA,max} = T_{DC,sup} - \Delta T_{PP}$) [24] |
| BP | - | 10% | Airflow by-pass [24] |
| $\eta_{fan}$ | - | 75% | Fan efficiency (incl. Motor) [24] |
| $V_{DC,max}$ | m$^3$/s | 246 | Maximum recirculated airflow inside data hall (calculated) |
| $V_{OA,max}$ | m$^3$/s | 246 | Maximum external airflow (assumed equal to $V_{DC,max}$) |
| $\Delta p_{fan,OA,max}$ | Pa | 350 | Pressure drop over economizer fan [24] (confirmed also by [11]) |
| $\Delta p_{fan,IASE,max}$ | Pa | 400 | Additional pressure drop for moving DC air through HX [24] |
| $\Delta p_{fan,CRAH,max}$ | Pa | 700 | Pressure drop for recirculating air inside data hall [24] |
| $T_{air,rejected}$ | °C | 20 | Assumed air outlet temperature after heat recovery |

Figure 13 represents the cooling system operational mode according to outdoor air conditions. When they satisfy the data center supply requirements (zone 1), outdoor air is processed through the heat exchanger to cool the recirculating airflow inside the IT room. When outdoor air is outside zone 1, instead, the chiller is used. Notice that no significant sources of moisture are present in data centers, therefore neither humidification nor dehumidification processes are required for the recirculating air.

The outdoor air, instead, could benefit from an evaporative cooling for certain values of adiabatic saturation temperatures (zone 2) [3], [24]. However, this second zone was neglected due to the lack of data on hourly humidity conditions in Western Denmark.

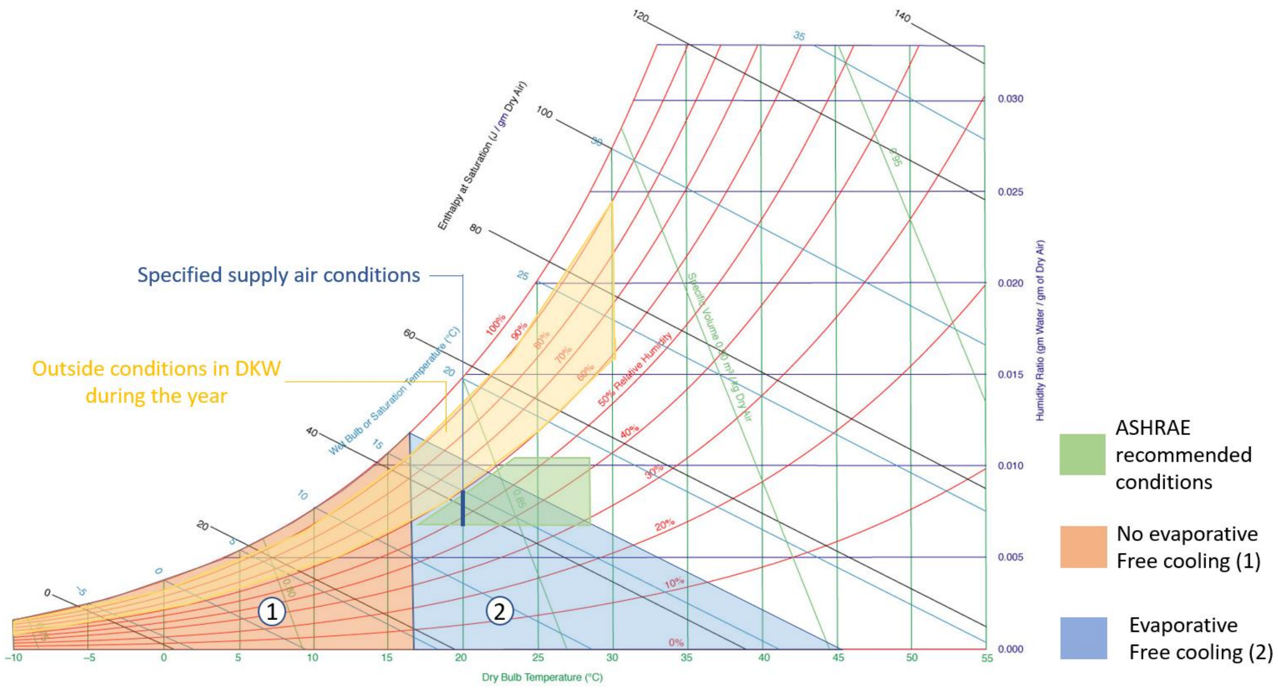


**Figure 13 –** Cooling system operational modes according to outdoor air conditions

The operational mode of the cooling system strongly affects its power consumption. However, in both cases the air is recirculated inside the data hall through the fans of CRAH units. Therefore, their power demand is always present and it is calculated according to eq.(3):

$$P_{fan,CRAH} = P_{fan,CRAH,max} * \left(\frac{V_{DC}}{V_{DC,max}}\right)^{2,5} \quad (3)$$

where:

$$P_{fan,CRAH,max} = \frac{V_{DC,max} \Delta p_{CRAH,max}}{\eta_{fan}}$$

$$V_{DC,max} = \frac{Q_{IT,max}}{\rho c_p \Delta T_{IT}} * \frac{1}{1 - BP} \cong 246\ m^3/s$$

Notice that the exponent is not set to 3 (as it would be expected from the theory) to account for variations in the fan efficiency with the airflow [24].

The flowchart of Figure 14 provides an insight on how the two different cooling strategies described before affect the power consumption of the cooling system and the waste heat.

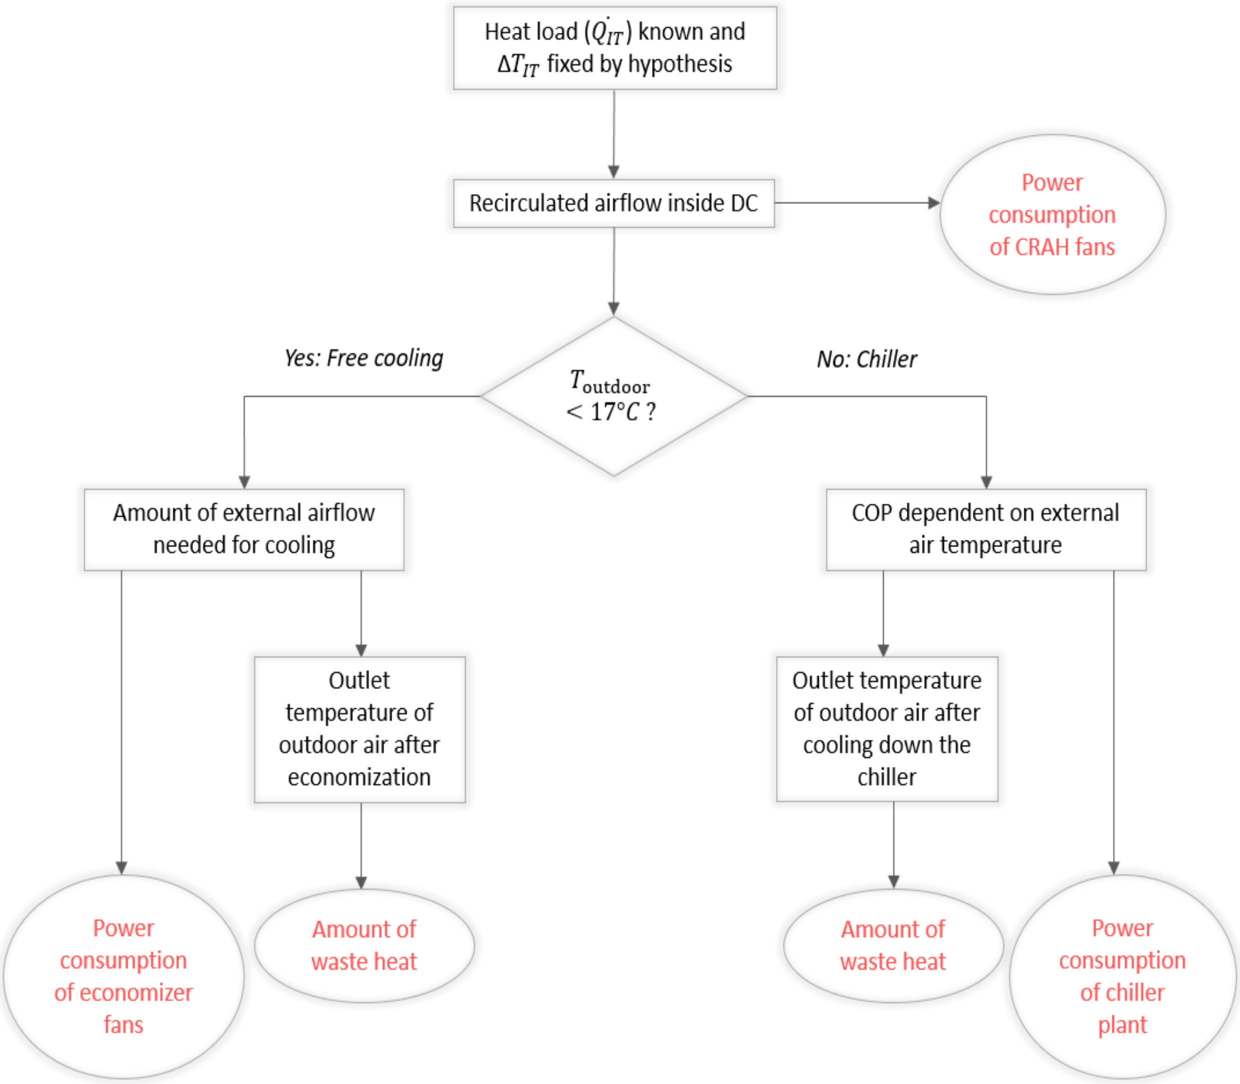


**Figure 14 –** Flowchart for calculating the power consumed by the cooling system and the amount of waste heat

It should be noticed that, even when outdoor air has a temperature above 17°C, a partial economization could be implemented, thus reducing the thermal load to be removed with the chiller. However, this situation was not considered.

### 4.2.1 Free cooling mode

Figure 15 presents the scheme adopted when the free cooling mode is available (zone 1 in Figure 13). In this case the power consumption of the cooling system is only due to the two additional fans required to overcome the pressure drops added by the heat exchanger of the economizer.

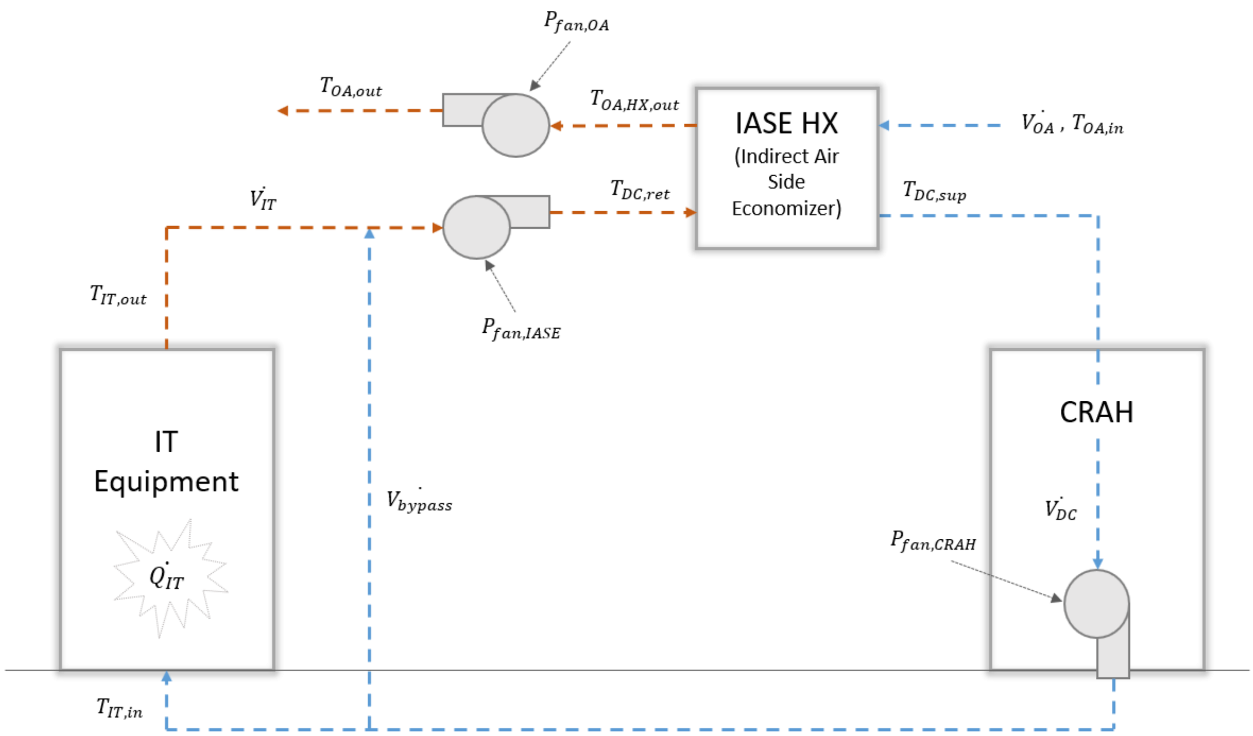


**Figure 15** -Scheme adopted for free cooling mode

Two temperatures are set across the economizer heat exchanger: the data hall supply temperature and the outdoor air temperature at the exit of the HX. The former ($T_{DC,sup}$) is equal to 20°C, while the latter ($T_{OA,HX,out}$) is 3°C lower than the return temperature ($T_{DC,ret}$), expressed by eq.(4).

The return air temperature is obtained considering an adiabatic mixing between the by-pass airflow and the IT airflow and adding the heat dissipation of the fan, which is slightly variable with the workload profile of the data hall. Therefore, $T_{DC,ret}$ is constant only if the workload profile is flat. However, variations registered with a flexible workload profile may also be neglected (less than 0.1 °C).

$$T_{DC,ret} = \left[(1 - BP) \ast T_{IT,out} + BP \ast T_{IT,in}\right] \ast \frac{(1 - \eta_{fan})P_{fan,IASE}}{V_{DC} \ast \rho \ast c_p} \qquad (4)$$

Figure 16 shows an example of temperature differences across the HX during spring.

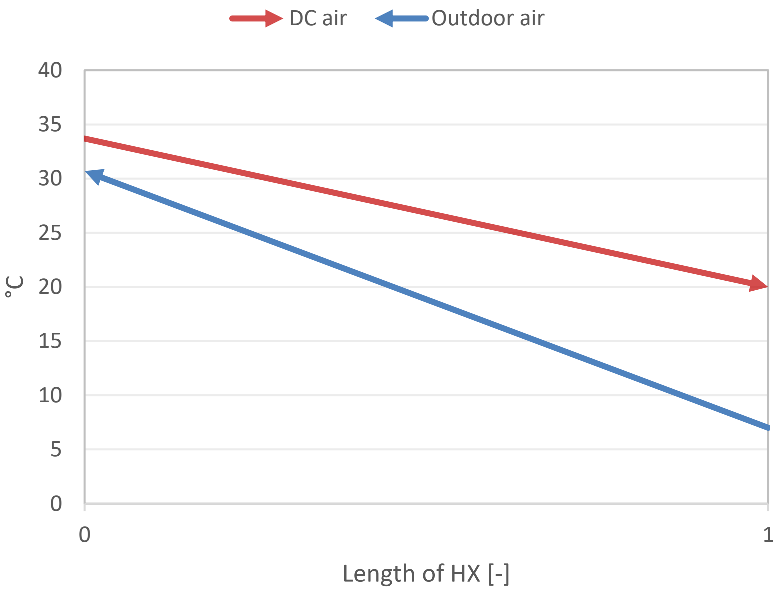


**Figure 16 –** Example of temperature differences across the economizer

Since $T_{OA,HX,out}$ is almost constant, the airflow taken from the external environment ($V_{OA}$) strongly depends on the outdoor temperature ($T_{OA,in}$) as shown in eq.(5):

$$V_{OA} = V_{DC} \frac{(T_{DC,ret} - T_{DC,sup})}{(T_{OA,HX,out} - T_{OA,in})} \tag{5}$$

The power consumption of the two additional fans ($P_{fan,IASE}$ and $P_{fan,OA}$) are then calculated according to eq.(3), but with different values for the maximum pressure drops (see Table 1) and assuming that the maximum outdoor airflow coincides with the maximum airflow recirculated inside the data hall.

In conclusion, when free cooling is available the power consumption of the cooling system is indicated by eq.(6). A graphical representation is given in Figure 17, which shows the power consumption of the data hall's cooling system in a typical spring day, when outdoor temperatures still allow using the free cooling technique. It is evident how external temperatures influence the total power consumption due to higher air flow rates processed by the fans.

$$P_{cooling\ system} = P_{fan,CRAH} + P_{fan,IASE} + P_{fan,OA} \tag{6}$$

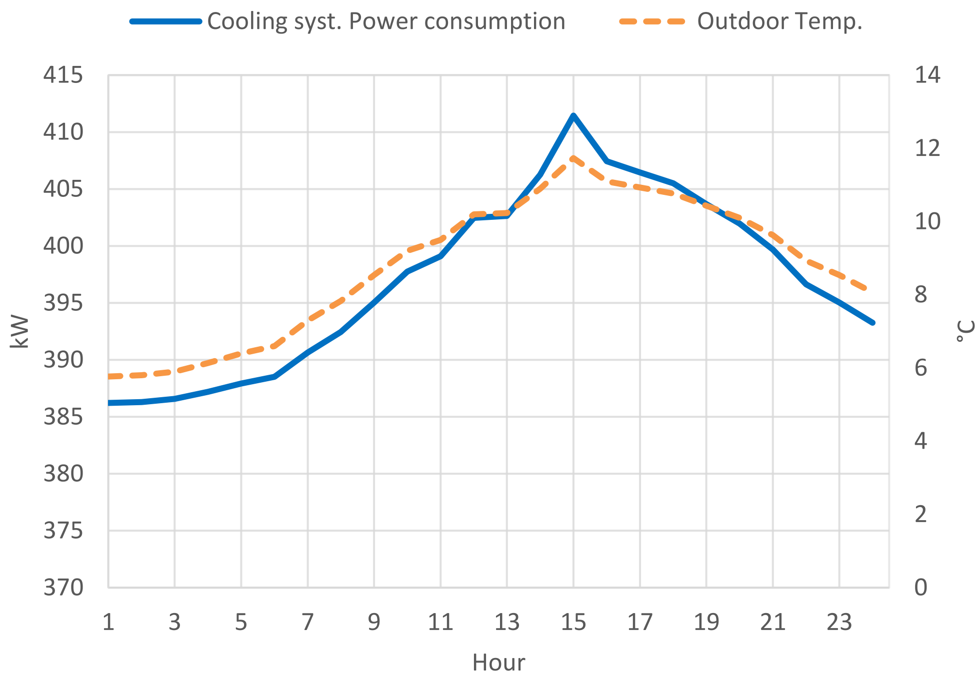


**Figure 17 –** Example of cooling system power consumption of the sample data hall vs. outdoor temperature (free cooling operation)

### 4.2.2 Cooling operated with chiller

Figure 18 introduces the cooling strategy adopted when external air conditions do not comply with the requirements for indirect airside economization.

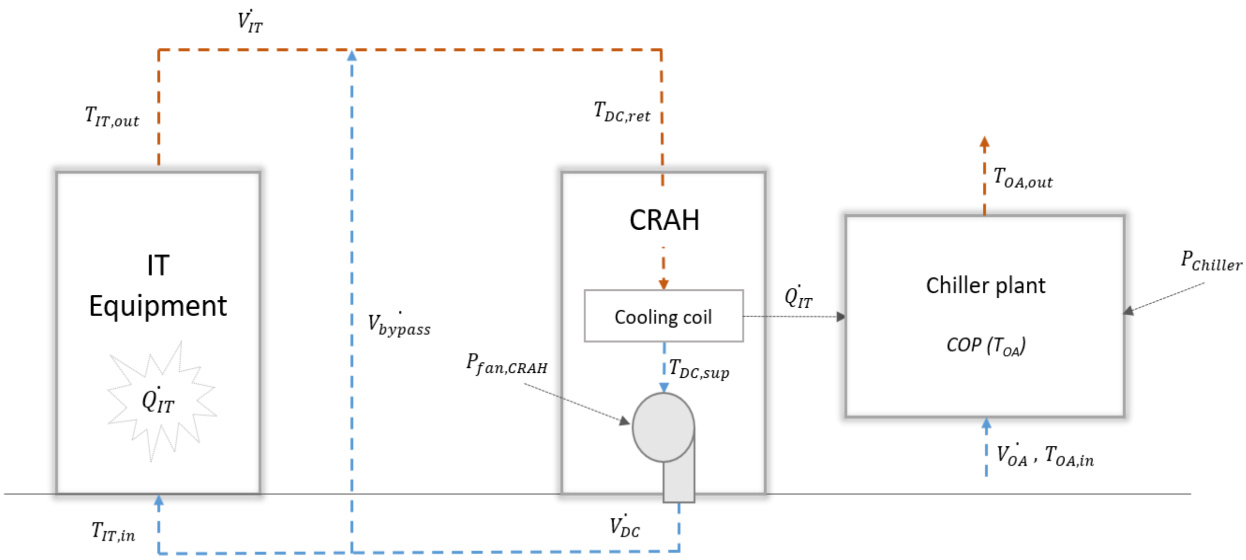


**Figure 18 –** Scheme adopted when the cooling is operated by the chiller plant

The chiller plant is comprehensive of water pumps, the chiller itself and the cooling tower. The COP of the plant is dependent on external temperatures (Table 2) [3], while its variation at different loads is not considered. A linear interpolation of the data in Table 2 was then used.

**Table 2 –** Dependence of chiller efficiency on external temperatures [3]

| **Outdoor T [°C]** | 0 | 5 | 10 | 15 | 20 | 25 | 30 | 35 | 41 |
|---|---|---|---|---|---|---|---|---|---|
| **COP** | 5.82 | 5.49 | 5.13 | 4.74 | 4.34 | 3.93 | 3.52 | 3.12 | 2.66 |

Therefore, the power consumption of the chiller plant is expressed by eq.(7):

$$P_{chiller} = \frac{Q_{IT}}{COP} \tag{7}$$

Inside the data hall, instead, the economizer is by-passed, thus the additional pressure drops of the free cooling mode should not be accounted.

In conclusion, when the chiller plant is activated the power consumption of the cooling system is expressed by eq.(8). A graphical representation is visible in Figure 19: when outdoor temperatures do not comply with free cooling requirements the chiller plant triplicates the total power consumption.

$$P_{cooling\ system} = P_{fan,CRAH} + P_{chiller} \tag{8}$$

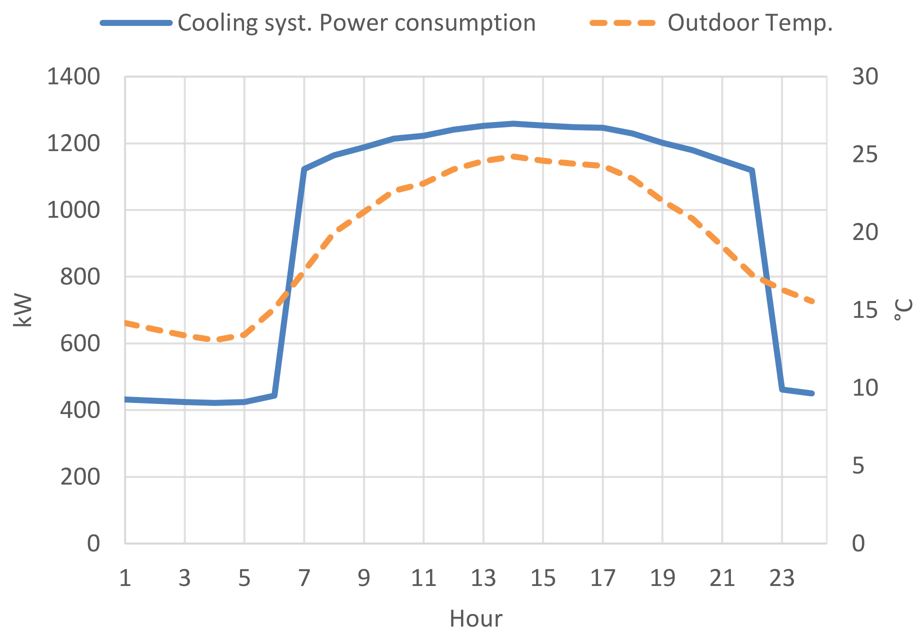


**Figure 19 –** Example of cooling system power consumption of the sample data hall vs. outdoor temperature (chiller enters in operation)

### 4.2.3 Available waste heat

TIMES-DK already allows some heat recovery from industrial processes generating excess heat according to three levels of temperature: 40°C, 80°C and 120°C. The first two levels need heat pumps to increase the heat temperature above 90°C, which is the supply temperature of district heating, while the third type of excess heat can be directly reused through heat exchangers. Although large data centers' waste heat has a very low quality (20°C), it can still be reused as suggested by the Danish Energy Agency technology catalogue [31]. Figure 20 contains a graphical representation of electric compression heat pumps, while Table 5 presents their main technological and financial characteristics. In addition, in order to decouple the DH grid and data centers, an intermediate heat exchanger should be foreseen [30] (Figure 21).

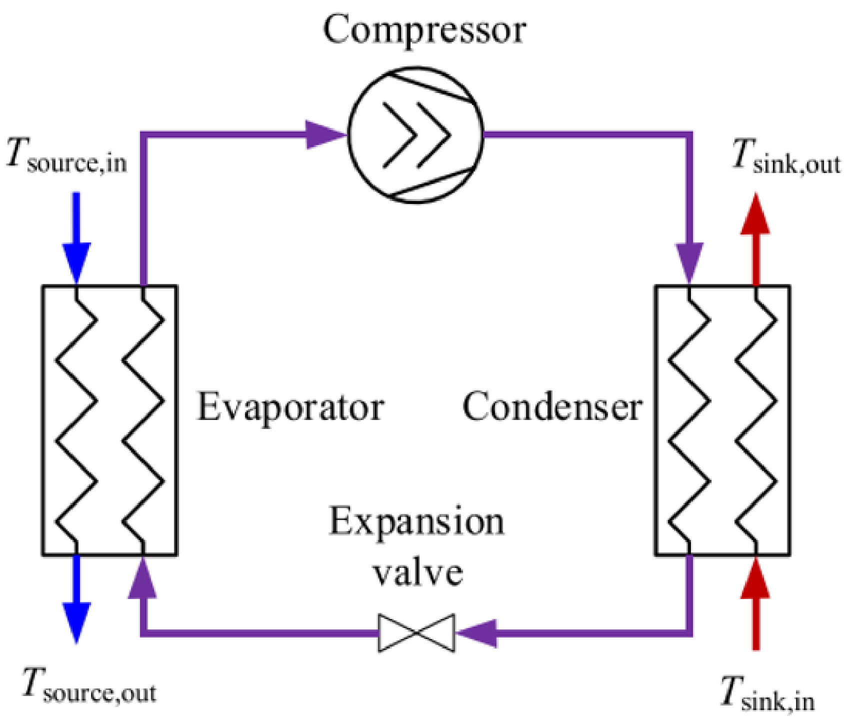


**Figure 20 –** Scheme of a HP for upgrading low temperature heat [31]

The temperature of the source is actually the logarithmic temperature, which is expressed by eq.(9):

$$T_{lm,source} = \frac{T_{source,in} - T_{source,out}}{\ln\left(\frac{T_{source,in}}{T_{source,out}}\right)} \tag{9}$$

The average heated outdoor air temperature from the sample data hall is 31°C. Assuming 5°C as temperature difference in the intermediate heat exchanger, the temperature rage of the heat pump source should be 15-26°C in order to comply with the logarithmic temperature of 20°C specified before. Therefore, the air stream coming from the data hall can be cooled down to 20 °C as shown in Figure 21.

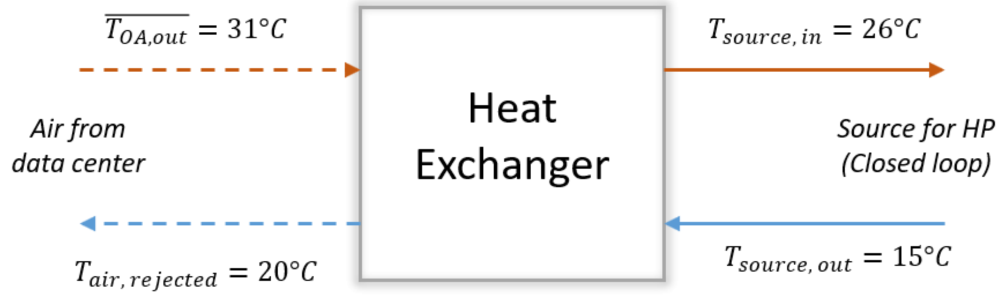


**Figure 21 –** Assumed working temperatures of intermediate heat exchanger between the data hall and the HP for waste heat recovery

Finally, the heat available at the heat exchanger is calculated according to eq.(10):

$$Q_{HX} = V_{air} * \rho * c_p * (T_{OA,out} - T_{air,rejected}) \quad (10)$$

Where:

| | *Free cooling mode* | *Cooling with chiller* |
|---|---|---|
| $V_{air}$ | $V_{OA}$ | $V_{OA,max}$ (fixed by hypothesis) |
| $T_{OA,out}$ | $T_{OA,out} \cong T_{DC,ret} - \Delta T_{pp}$ | $T_{OA,out} = T_{OA,in} + \frac{Q_{IT} + P_{chiller}}{V_{OA,max}\rho c_p}$ |
| $T_{air,rejected}$ | 20 °C | 20 °C |

## 4.3 Implementation in TIMES-DK

As indicated in 2.4, in TIMES-DK the entire energy system is divided into 5 sectors: supply (SUP), power sector (ELC), residential sector (RES), industry (IND) and transports (TRA). Although the industrial sector already represents information technologies, it was chosen to create a new sector for the large data centers, called “LDC”. The advantage of this implementation is the simplicity to add or remove the presence of large data centers in the model of the Danish energy system, thus making immediately evident their impact.

Moreover, large data centers have features that should be "activated" or "deactivated" according to the purpose of the different simulations. Therefore, separate files have been created to consider diverse demand projections and the characterization of commodities at the time-slice level (including or neglecting the hypothetical daily usage profile).

Summarizing, the following worksheets have been created:

- 1 *VT file* for the new sector;
- 3 *Scenario files* for the desired features of large data centers;
- 1 *SubRes file* for allowing investments in the reuse of waste heat.

### 4.3.1 Characterization of processes and commodities

A visual representation of the adopted model is given in Figure 22: boxes stand for technologies, while arrows denote commodities (the red one is the exogenous demand, i.e. the driver). Connections with the rest of TIMES-DK are performed through the electricity withdrawn from the grid ("ELCC") and the heat possibly injected in DH networks ("HET"). Moreover, to be consistent with the geographical classification of TIMES-DK (see 2.4), processes and commodities were duplicated into "central" and "decentral": the asterisk is intended to substitute the letters "C" and "D" respectively for central and decentral technologies/commodities.

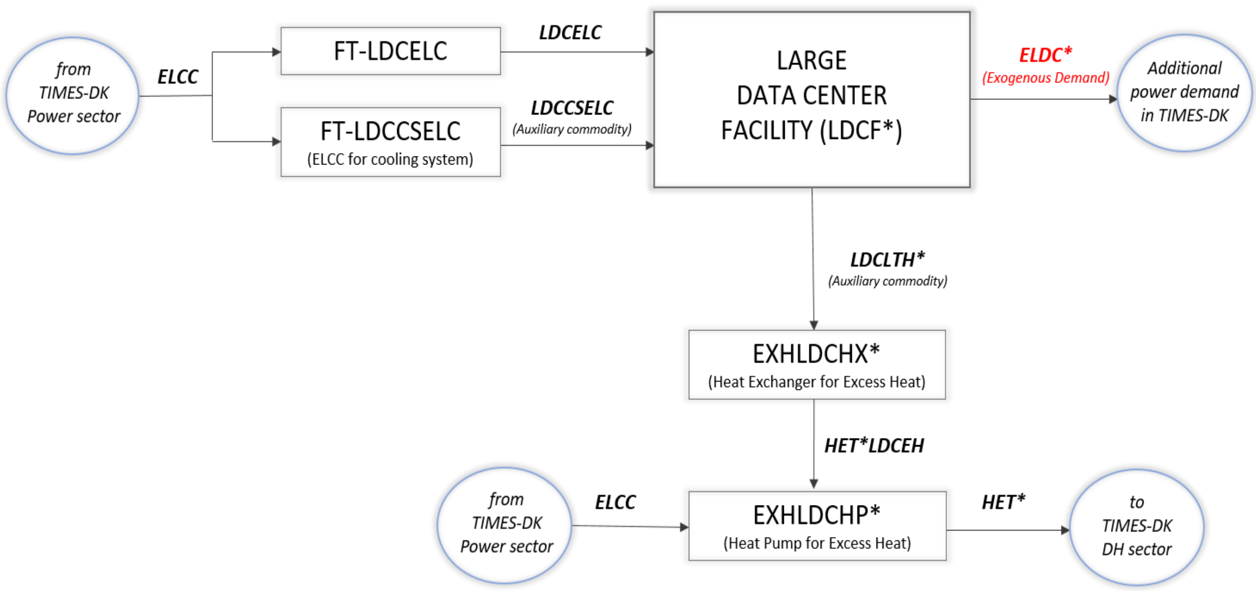


**Figure 22 –** Visual representation of the adopted model for the implementation in TIMES-DK

In the framework of TIMES models, large data centers facilities can be seen as demand technologies (here called "LDCF"), i.e. as processes that have as output commodity the estimated power demand ("ELDC"). On the other hand, the input is the electricity from the grid supplying the IT equipment ("LDCELC"). In addition, there are two other commodities to consider: the power for the cooling system ("LDCCSELC") and the low-temperature excess heat ("LDCLTH"). However, they are tied to the activity of the process (which is ELDC), therefore they are modelled as auxiliary commodities. In fact, this choice allows to represent them as a fraction of ELDC through the attribute called "VDA_FLOP" instead of expressing their annual or time-sliced values. Their characterization is presented in Table 3.

**Table 3 –** Characterization of "Large Data Centers Facilities" processes in TIMES-DK

| TechName | Comm-IN | Comm-IN-A | Comm-OUT | Comm-OUT-A | EFF | VDA_FLOP | AF | NCAP_TLIFE |
|---|---|---|---|---|---|---|---|---|
| **Technology Name* | *Input Commodity* | *Auxiliary Input Commodity* | *Output Commodity* | *Auxiliary Output Commodity* | *Efficiency* | *Process Input/Output tied to activity* | *Availability Factor* | *Lifetime of Process* |
| **Unit* | | | | | | *PJ per PJ activity* | | *Years* |
| LDCFC | LDCELC | | ELDCC | | 1 | | 1 | 50 |
| | | LDCCSELC | | | | 0.12 | | |
| | | | | LDCLTHC | | 0.58 | | |
| LDCFD | LDCELC | | ELDCD | | 1 | | 1 | 50 |
| | | LDCCSELC | | | | 0.12 | | |
| | | | | LDCLTHD | | 0.58 | | |

Following the customary in TIMES models, fuel technologies were also introduced. They are usually indicated with the prefix "FT-" and are good tools to better track the final consumptions by sector because they simply change names to commodities (see Table 4).

**Table 4 –** Characterization of "Fuel Technologies" in TIMES-DK

| TechName | Comm-IN | Comm-OUT | EFF | NCAP_TLIFE |
|---|---|---|---|---|
| **Technology Name* | *Input Commodity* | *Output Commodity* | *Efficiency* | *Lifetime of Process* |
| **Unit* | | | | *Years* |
| FT-LDCELC | ELCC | LDCELC | 1 | 100 |
| FT-LDCCSELC | ELCC | LDCCSELC | 1 | 100 |

To complete the model, new technologies in which the algorithm can invest are added. In this case they are heat exchangers ("EXHLDCHX") and heat pumps ("EXHLDCHP") to recover the excess heat, whose quality is upgraded for the injection in the district heating network (i.e. from "LDCLTH" to "HET"). Data regarding these processes refer to the Danish Energy Agency's technology catalogue [31] and are shown in Table 5 and Table 6.

**Table 5 –** Technical and financial characterization of electrical compression heat pumps for DH [31]

| | 2015 | 2020 | 2030 | 2050 |
|---|---|---|---|---|
| ***Energy/technical data*** | | | | |
| Total eff., net (%), annual average, waste heat 20° C | 440 | 500 | 600 | 740 |
| Planned outage (weeks per year) | 0.5 | 0.5 | 0.5 | 0.5 |
| Technical lifetime (years) | 25 | 25 | 25 | 25 |
| Construction time (years) | 0.5 | 0.5 | 0.5 | 0.5 |
| ***Financial data*** | | | | |
| Nominal investment (M€ per $MW_{heat}$) | 0.70 | 0.66 | 0.59 | 0.53 |
| Fixed O&M (€/$MW_{heat}$/year) | 2000 | 2000 | 2000 | 2000 |
| Variable O&M (€/$MWh_{heat}$) | 8.4 | 8.2 | 8.8 | 8.2 |

**Table 6 –** Technical and financial characterization of heat exchangers for waste heat recovery [31]

| | 2015 | 2020 | 2030 | 2050 |
|---|---|---|---|---|
| ***Energy/technical data*** | | | | |
| Total eff., net (%), annual average | 100 | 100 | 100 | 100 |
| Technical lifetime (years) | 20 | 20 | 20 | 20 |
| ***Financial data*** | | | | |
| Nominal investment (M€ per $MW_{heat}$) | 0.07 | 0.07 | 0.07 | 0.07 |
| Fixed O&M (€/unit/year) | 1250 | 1250 | 1250 | 1250 |
| Variable O&M (€/$MWh_{heat}$) | - | - | - | - |

### 4.3.2 The driver: power demand projections

The estimation of the power that will be required by these large data centers in Denmark is taken from [15]. This projection constitutes a driver for the model, because it induces the model to choose coherent investments in order to always satisfy their annual demand level.

The *Energinet* company is in charge to supply the demand of large data centers with new power stations. *Energinet* knows the capacity of the power stations that are under construction but ignores what will be their demand pattern over the time horizon. Therefore, they assumed a precautionary flat utilization profile with a 100% daily demand for the entire year (see 4.1.1). Consequently, in the absence of further clarifications from the owners of these data centers, the projection made by *Energinet* was taken as a reference for the main results of this thesis (presented in Chapter 5). However, it was split between "central" and "decentral" demand (Figure 23) to fit the geographical subdivision mentioned before: the only decentral data center is the one in Viborg, whose rated power corresponds to 40% of the total [29].

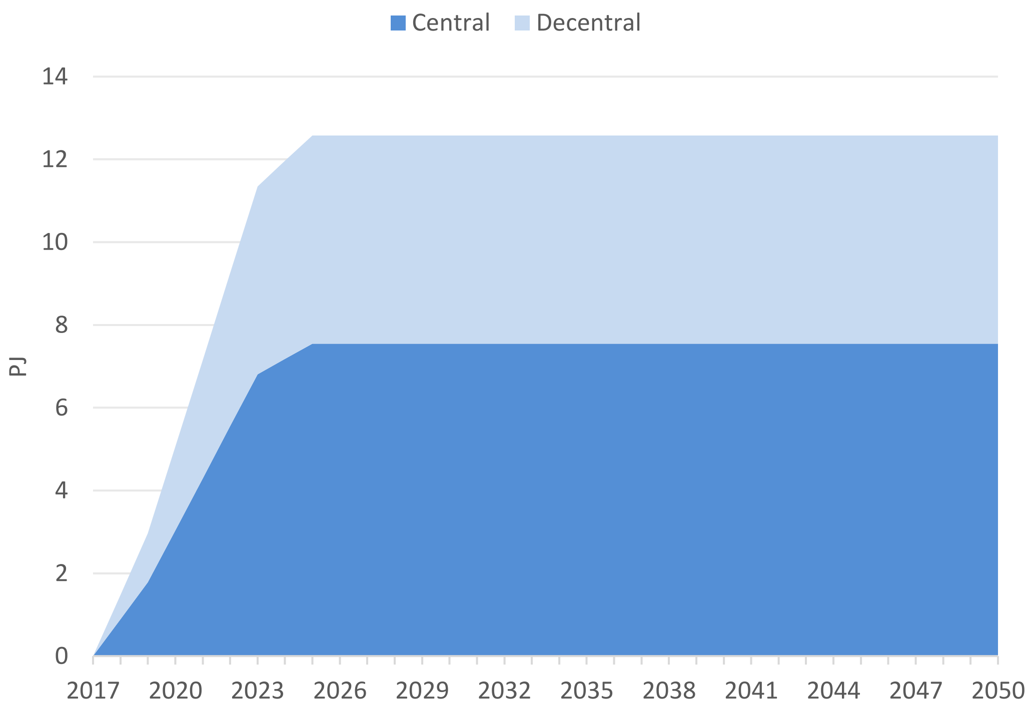


**Figure 23 –** Reference power demand projection for large data centers

Nevertheless, the future demand of large data centers in Denmark is still uncertain. Therefore, two additional projections were designed to address the following features:

- A daily variable workload profile, which causes a lower annual electricity demand with respect to the estimates made by *Energinet*. Figure 24 shows how the projection is affected when the flexible profile, introduced in 4.1.1, is simulated in each day of the year for the rated power of data centers in that year, which is simply obtained by dividing *Energinet*'s estimated MWh by the yearly 8760 hours;
- A continuous increased demand after 2025, as expected by DEA's analysts [37]. Also in this case the variable utilization profile was applied and differences are shown in Figure 24.

The effect of these two additional different projections is presented and discussed in Chapter 6, which is dedicated to sensitivity analyses.

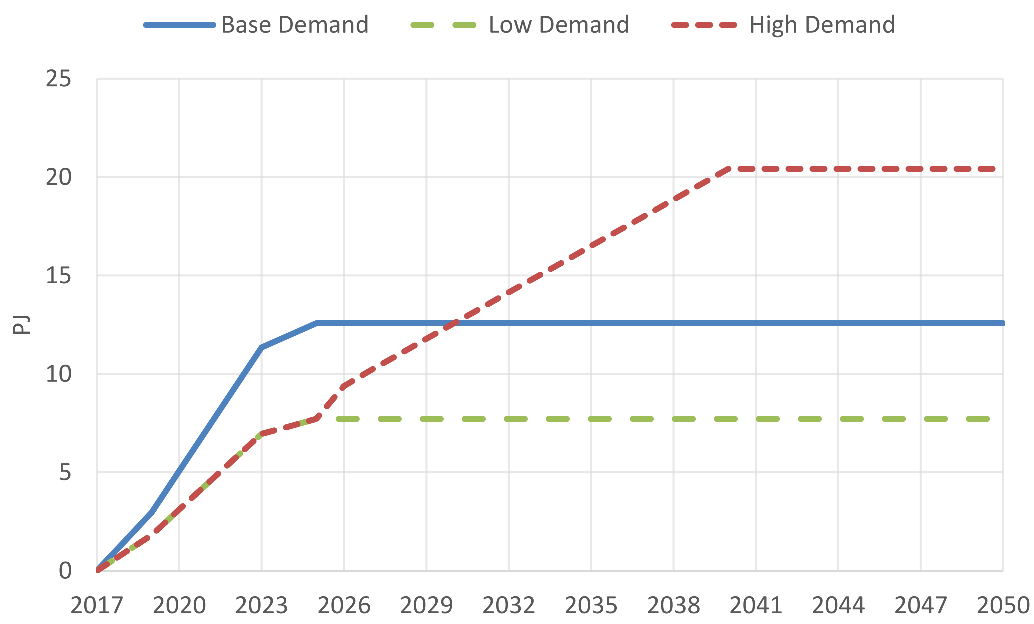


**Figure 24 –** Different LDC power demand projections: "base" (*Energinet*), "low" (*Energinet* + variable workload profile), "high" (DEA's projection + variable workload profile)

### 4.3.3 Representation at time-slice level

When a commodity is tracked at the time-slices level, TIMES models distribute the annual value of that commodity among the time-slices in proportion to their length, unless specified otherwise. The default subdivision follows the *YRFR* attribute, which indicates the relative duration of a time-slice with respect to a year. Therefore, when a daily flat utilization profile is assumed for data centers

(*Energinet*'s estimate), the default subdivision for the annual value of *ELDC* is adopted since the power demand is the same in each hour belonging to a time-slice. On the other hand, when a daily variable utilization profile is simulated, the annual value of *ELDC* should be subdivided differently because the hourly power demand is variable. In this latter case, the fraction of the annual power demand (*ELDC*) consumed in each time slice is manually specified through the aggregation of its hourly values. Figure 25 shows how the annual value of *ELDC* is spread over the time-slices in the framework of TIMES-DK, comparing the case with a flat and a variable workload profile. Unfortunately, the time scale of TIMES-DK does not allow to capture relevant differences between the two cases, but it is important to keep this distinction in mind when switching to a finer temporal resolution.

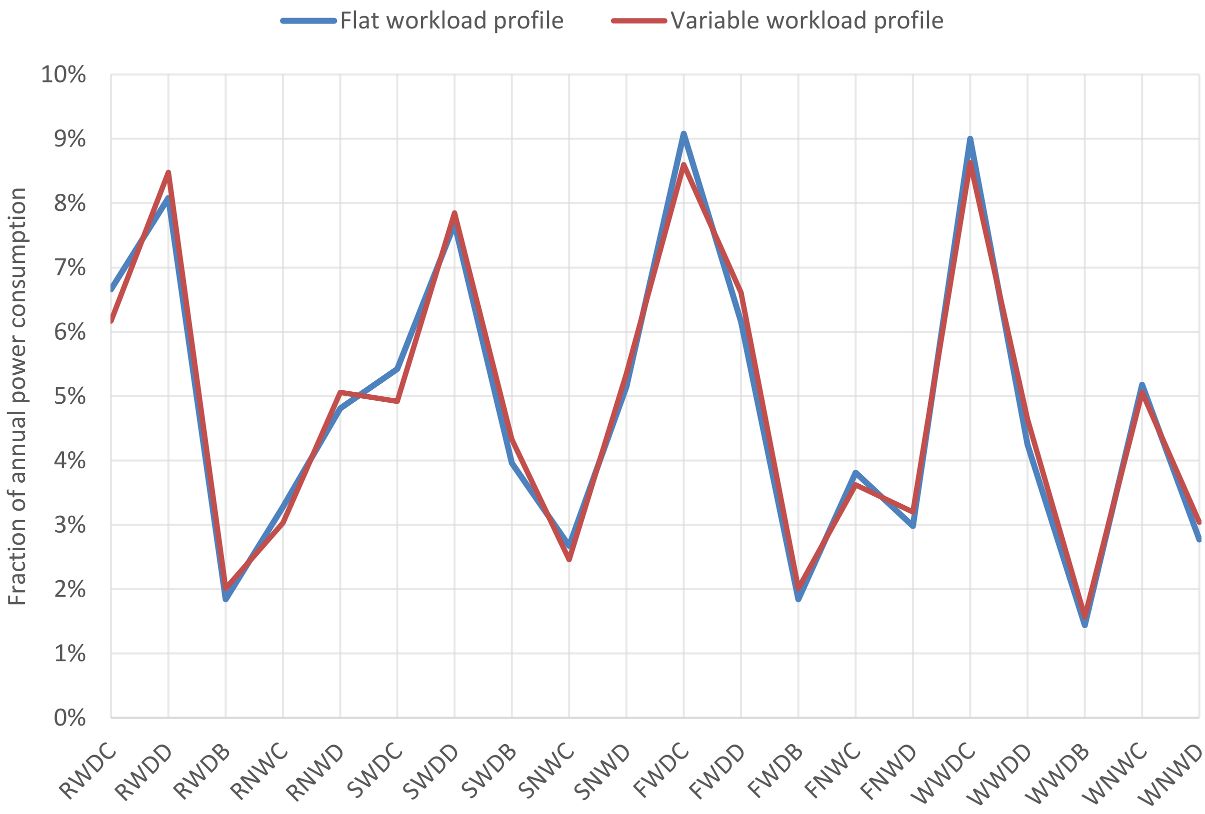


**Figure 25 –** Large data centers power consumption at the time slice level (in the most significative time-slices)

As already mentioned before, the power of the cooling system and the amount of waste heat are defined as auxiliary commodities dependent on *ELDC*. Therefore, their values at the time-slices level are specified by the attribute *VDA_FLOP* (PJ of commodity per PJ of data centers' activity). This is obtained through the aggregation of the hourly results retrieved from the thermodynamic model.

Figure 26 and Figure 27 show the inputs to TIMES-DK in the hypothesis of flat and variable workload profiles.

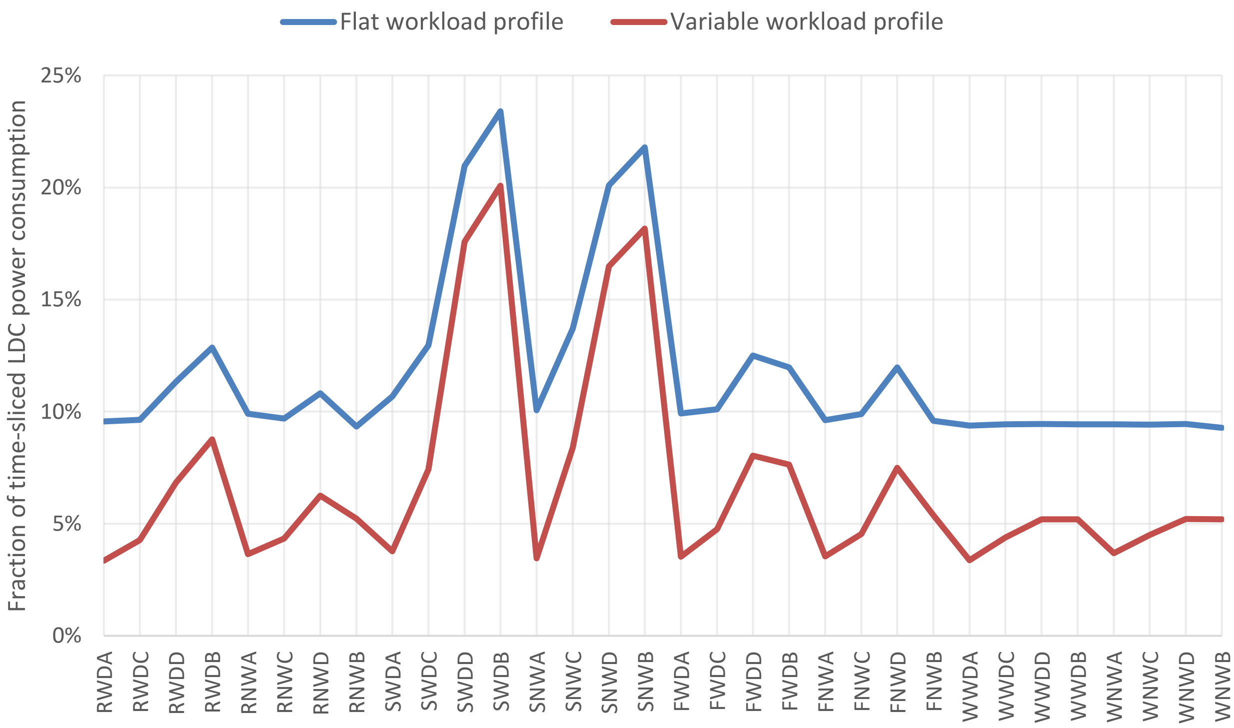


**Figure 26 –** Power consumption of cooling system at the time-slice level expressed as a fraction of the value of ELDC in each time-slice (input for TIMES-DK)

When a variable workload profile is used, it is evident that also the relative power consumption of the cooling system is reduced because it depends on the recirculated airflow (which is lower because the heat dissipated by the IT equipment is lower, while the temperature increase across the servers is kept constant by hypothesis). On the other hand, the relative waste heat produced by large data centers is practically independent on the workload profile because the recirculated airflow is directly proportional to the heat dissipated by the IT equipment, which in turn coincides with the power consumption of the IT equipment by hypothesis.

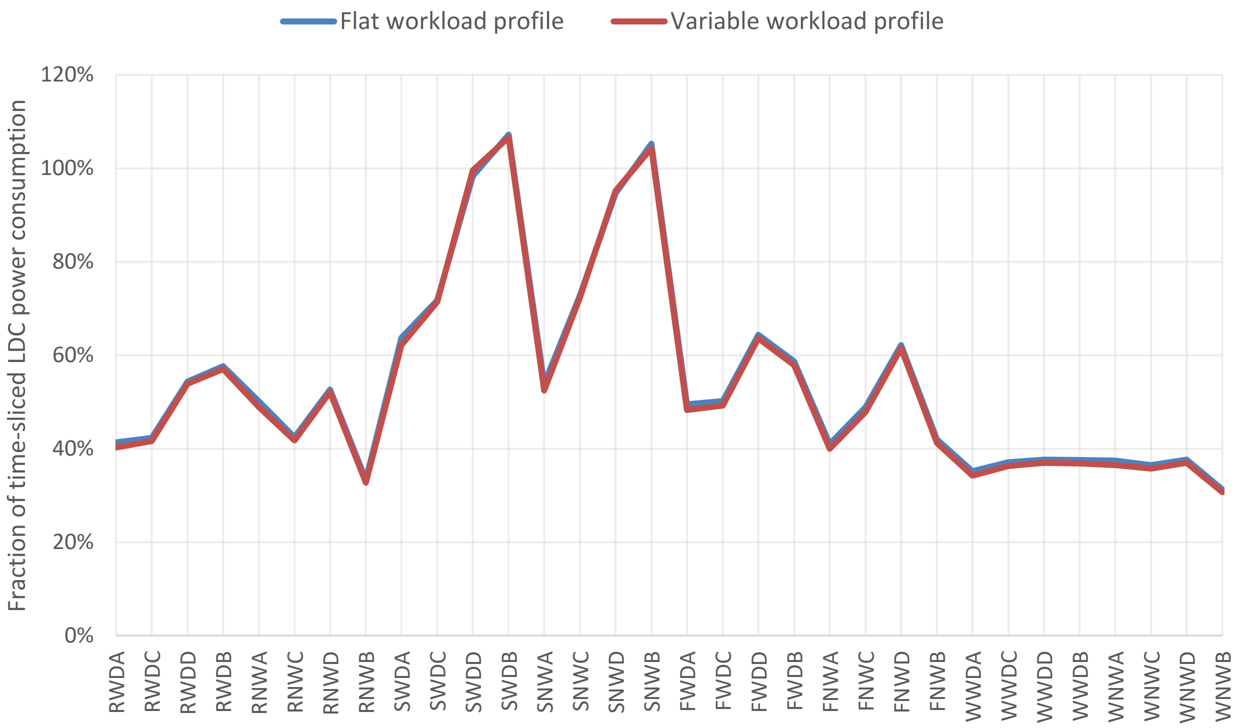


**Figure 27 –** Large data centers' waste heat at the time-slice level expressed as a fraction of the value of ELDC in each time-slice (input for TIMES-DK)

## 4.4 Features of benchmark scenarios

The two scenarios chosen for testing the introduction of large data centers in the Danish energy system have distinct purposes. The Frozen Policy scenario aims at understating what will be the shape of the energy system if current policies are not updated when they expire. Usually, most of the taxes are permanent throughout the time horizon, while subsidies have a limited duration. On the other hand, the Target scenario ("Carbon Free in 2050", also referred as *CF2050*) has the goal of proposing the socio-economic optimal configuration for the system. In fact, it reaches the target (no $CO_2$ emissions in 2050) desired by the society (social optimum) within the framework of a TIMES model (economic optimum). Therefore, in this scenario the energy market is not influenced by duties and subventions because, based on the outcomes, policy makers can subsequently shape taxes and subsidies in accordance with the desired target.

However, both scenarios share the same settings apart from the emission goal (only present in *CF2050*) and the taxes/subsidies (only present in *FP*). Common assumptions regard:

- Demand projections of different sectors;

- Base constraints, such as minimum heat production from heating-only plants (in ELC sector), maximum PV and solar heating roof installation or maximum use of natural gas (for residential sector), laws for blending fuels (in transports), etc.;
- Delivery costs of commodities and their characterization at the time-slice level;
- New technologies available in each sector;
- Emission factors and efficiencies;
- Maximum potentials for renewable technologies in the power sector (see Table 7);
- Discount rate fixed to 4% in both scenarios, which gives more importance to public investments rather than private ones.

In addition, the setting regarding imports and exports of electricity was slightly modified: after a preliminary run it was decided to limit the ratio between annual imports and exports to 1 (i.e. to a condition of net zero exports). In fact, results showed an evident tendency in exporting excessive quantities of electricity because of convenient prices, even though great uncertainties are related to these values especially in the future. This would have clearly blurred the influence of data centers on the Danish energy system, as a relevant part of new power capacity investments were only driven by favorable exports.

**Table 7 -** Capacity bounds for renewable power technologies

| Technology | Capacity bound in 2011 [GW] | Capacity bound in 2025 [GW] | Capacity bound in 2050 [GW] |
|---|---|---|---|
| *Wind onshore (WON)* | 6 | - | 8 |
| *Wind offshore (WOF)* | 20 | - | 50 |
| *Solar PV* | 0.8 | 4.4 | 9 |

Figure 28 and Figure 29 show the power production mix in the selected scenarios (without data centers). As expected, their outcomes are not very different. In fact, the policies adopted by the Danish government already point in the direction of maximizing the renewable share in the power sector (where around 60% of the generated electricity already comes from renewable sources). Therefore, if in *CF2050* scenario fossil fuels are disadvantaged by the long-term emission target, in the *FP* scenario they are similarly disadvantaged by taxes. However, it is interesting to notice that the $CO_2$

emissions cumulated in the time horizon are higher in the *CF2050* scenario because the constraint is set only in 2050.

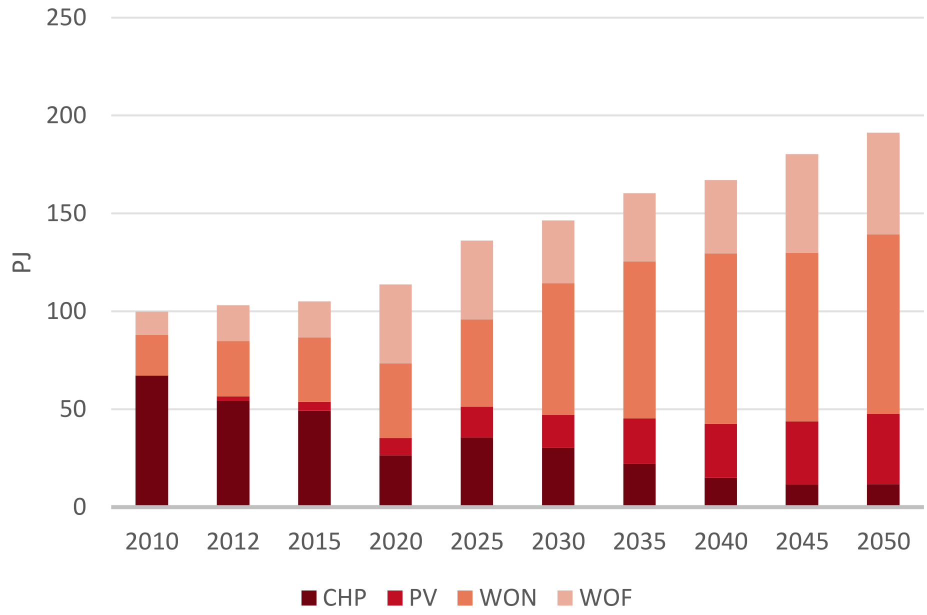


**Figure 28 –** Power production mix in Frozen Policy scenario

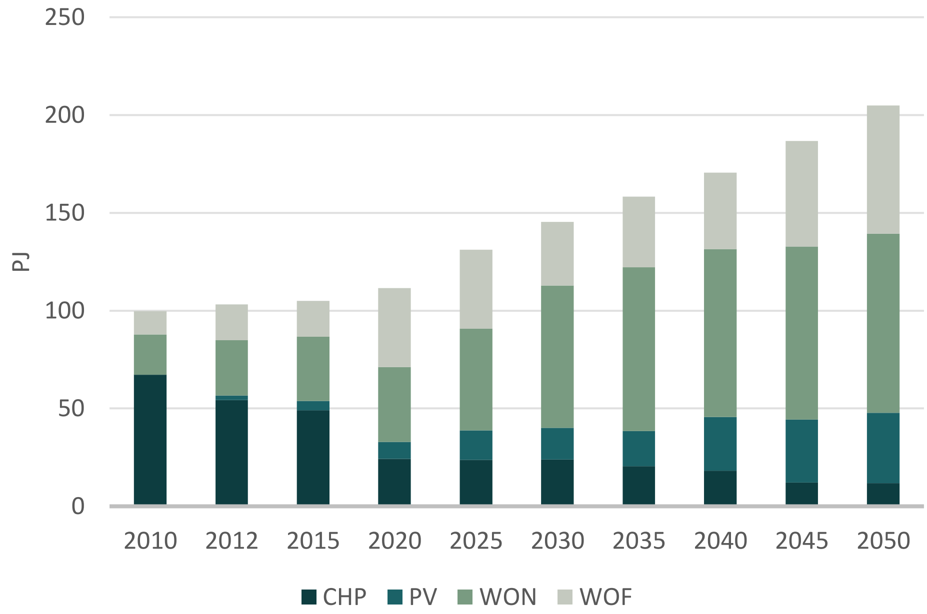


**Figure 29 –** Power production mix in Target scenario (CF2050)

# Chapter 5

## Results

This chapter presents the main outcomes of the thesis, obtained through the simulation of two scenarios: the Frozen Policy (*FP*) and the Target scenario ("Carbon Free in 2050", briefly indicated as *CF2050*). Each scenario was run twice, with and without large data centers (briefly referred as LDC).

The research question posed in the introduction ("What may the effects of large data centers on the Danish energy system be?") is here addressed. Therefore, results are presented comparing two configurations of the same scenario: *FP* without data centers vs. *FP* with data centers and *CF2050* without data centers vs. *CF2050* with data centers. In addition, relations between the two scenarios are also analyzed in order to evaluate how specific scenario settings influence the outcomes.

As far as the structure of this chapter is concerned, each section is devoted to different themes:

- The influence on the Danish power sector (section 5.1);
- The influence on $CO_2$ emissions (section 5.2);
- The possibility for waste heat recovery (section 5.3);
- The influence on economic parameters such as marginal prices of electricity and district heat (section 5.4).

Finally, the reader should keep in mind that the years indicated in the graphs are actually "milestone years", thus they represent a period of five years in accordance with the time definition in TIMES models (see 2.3.1).

## 5.1 Effects on the power sector

### 5.1.1 Investments in new power capacity

From the runs, it results that the additional power demand required by the large data centers changes the shape of the power sector. In fact, they do not only affect the amount of new power capacity to be installed, but they also influence the timing of the investments both in the Frozen Policy and in the Target scenarios (Figure 30, Figure 31).

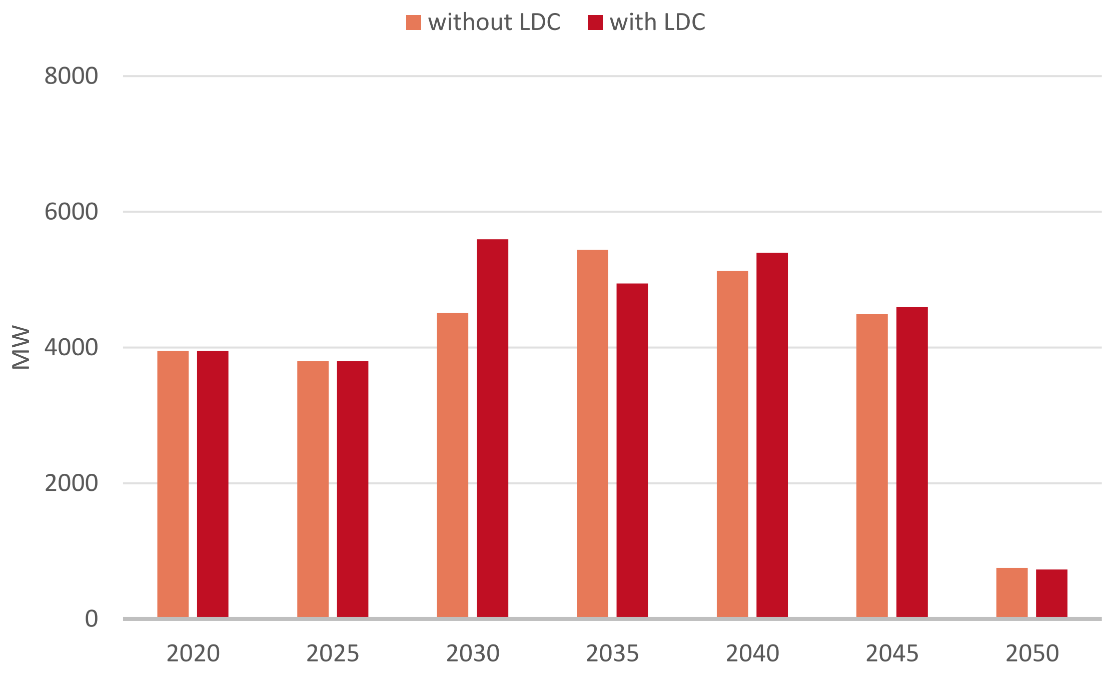


**Figure 30 –** New power capacity in Frozen Policy scenario with and without large data centers

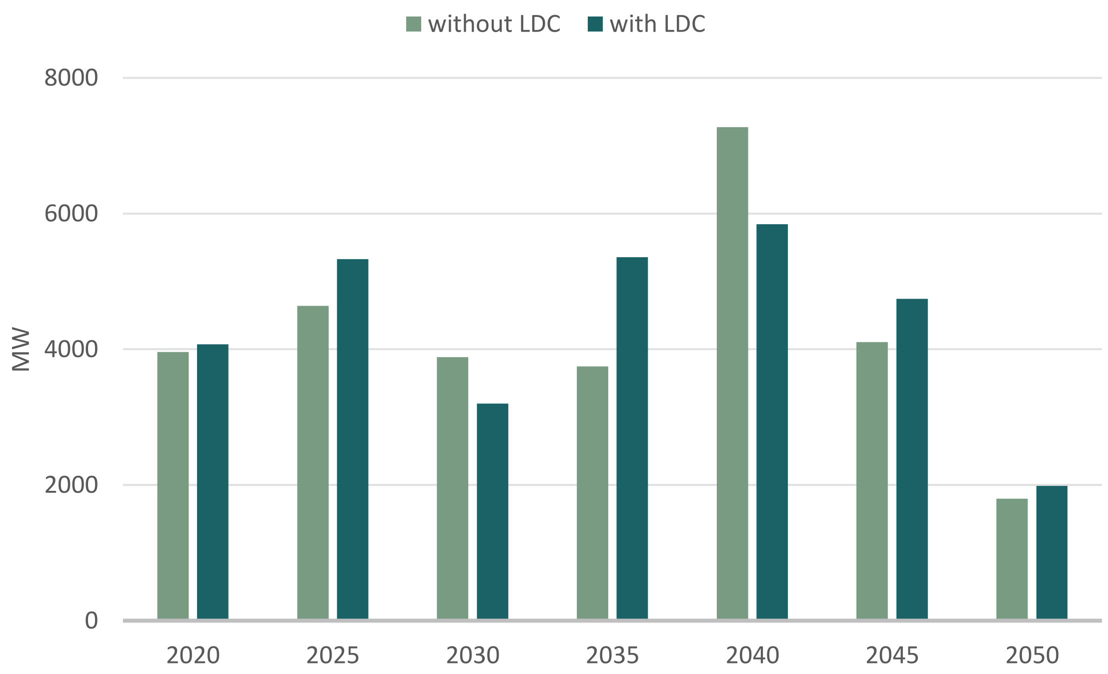


**Figure 31 -** New power capacity in Target scenario with and without large data centers

In both scenarios a clear pattern is visible: the presence of large data centers forces the model to anticipate investments in new power capacity. For instance, in the Target scenario (see Figure 31), the sum of the new capacity installed between 2025 and 2030 is the same regardless of the presence of large data centers in the model. However, when they are present, 685 MW of new capacity are moved from 2030 to 2025. This behavior is repeated in the following years, even though at the end of the time horizon the cumulated new capacity is obviously higher in the configurations with the data centers (respectively 940 MW more in the FP scenario and 1123 MW more in the Target scenario).

When detailing the new installations, it is visible that in both scenarios the additional new capacity required for the supplementary demand of data centers belongs to only one type of power plants: wind offshore. However, the optimal solution proposed by the model is not to invest immediately in it, but to postpone those investments after 2030 (Figure 32), when their cost is expected to be lower. Consequently, the model invests in other two technologies for the immediate future: wind onshore (in both scenarios) and photovoltaic (only in Target scenario), as shown in Figure 33 and Figure 34. The shift towards wind offshore technology is not only convenient, but also necessary in the long run because the potential for wind onshore and PV is constrained due to land occupation (respectively to 8 and 9 GW in 2050), while wind offshore is practically unlimited (bounded to 50 GW in 2050). Investments in CHP and thermal power plants, instead, are not affected.

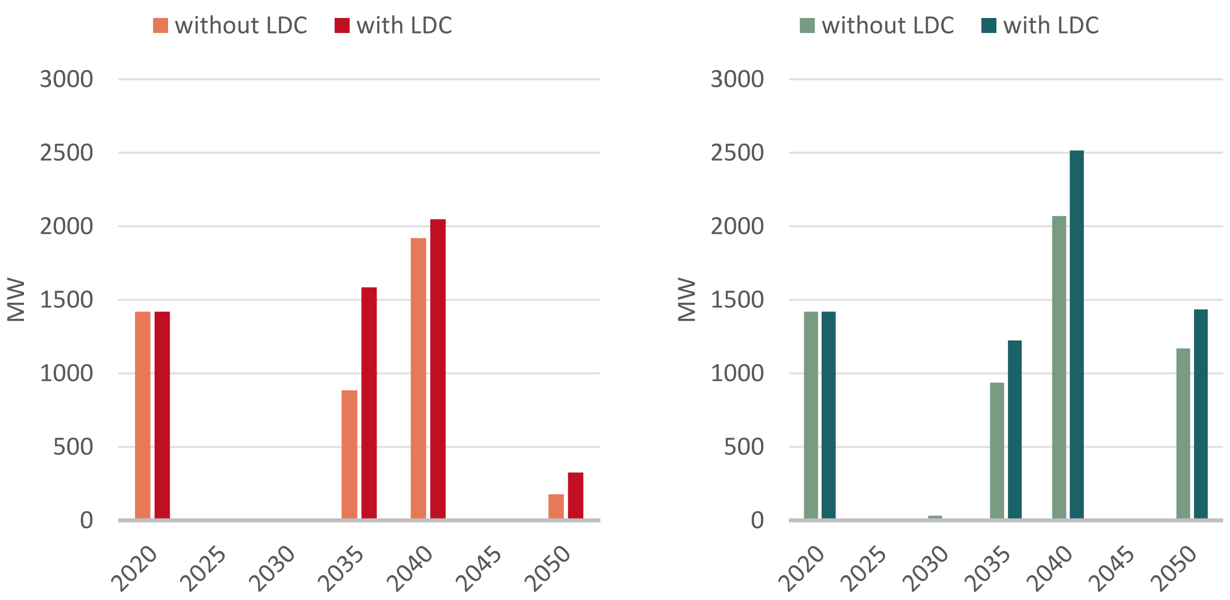


**Figure 32 –** New wind offshore power capacity: a) FP with and without LDC; b) CF2050 with and without LDC

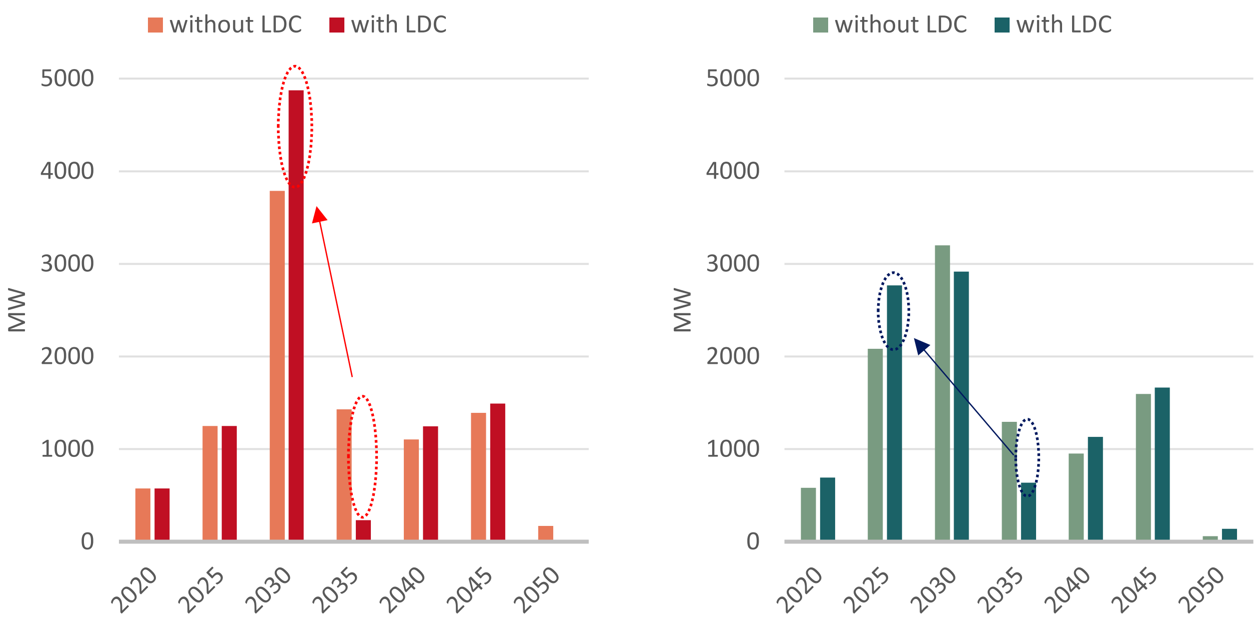


**Figure 33 -** New wind onshore power capacity: a) FP scenario with and without LDC; b) CF2050 with and without LDC

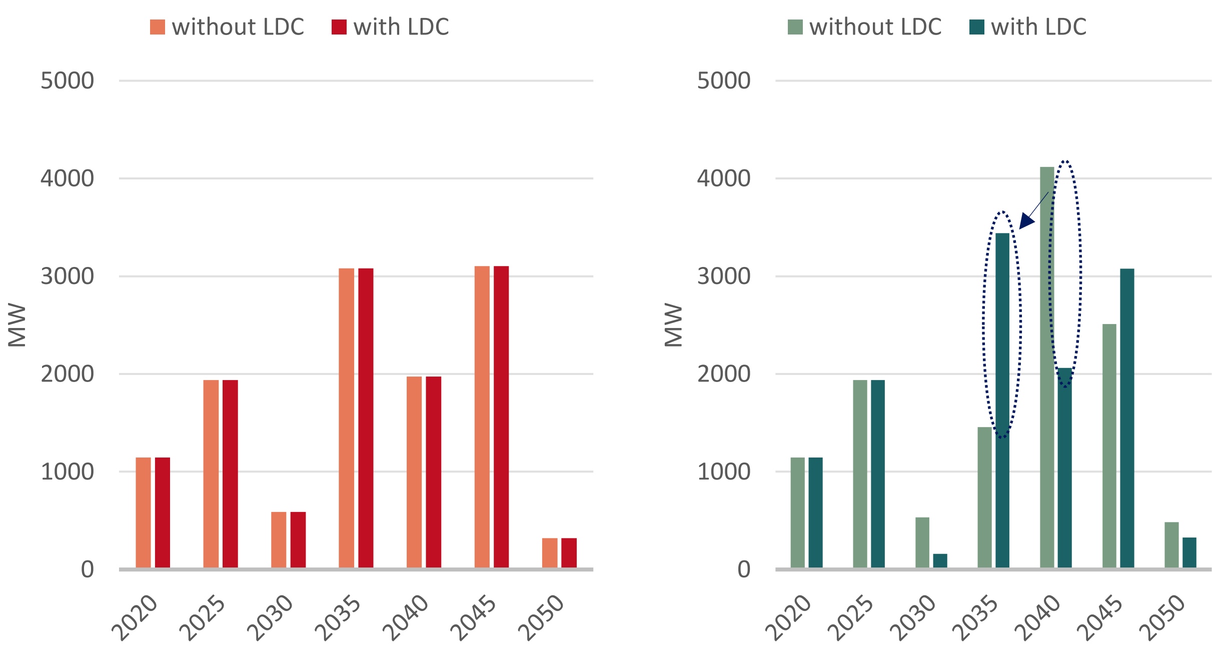


**Figure 34 -** New solar PV power capacity: a) FP scenario with and without LDC; b) CF2050 with and without LDC

### 5.1.2 Additional electricity production

The additional electricity production (Figure 35, Figure 36) reflects the new demand and the related additional investments made in new power capacity described before. Both the scenarios tend to satisfy the supplementary demand of large data centers with wind offshore technologies towards the end of the time horizon, while in the transition period (2020-2035) they rely on wind onshore and PV. Moreover, CHP plants have a key role in the first periods (especially in the FP scenario) despite the absence of additional new capacity. In fact, they are only partially exploited (their average full load hours are lower than 2900 within a year), while at the same time new renewable power plants are still quite expensive and there are no emission constraints in this period. Consequently, it is more convenient to extend the load factor of CHP plants than installing new renewable power plants in the first years. However, it should be noticed that the additional quota produced by CHP plants may be substituted by extra imports of electricity, even though they are not represented because of the constraint forcing the annual imports to be equal to the annual exports (see 4.4). It possible to notice also that, due to this constraint, monitoring the additional power production means monitoring the additional power demand required to the system in a year. Therefore, Figure 35 and Figure 36 show de facto how the supplementary power demand is allocated to the different energy sources at the annual level.

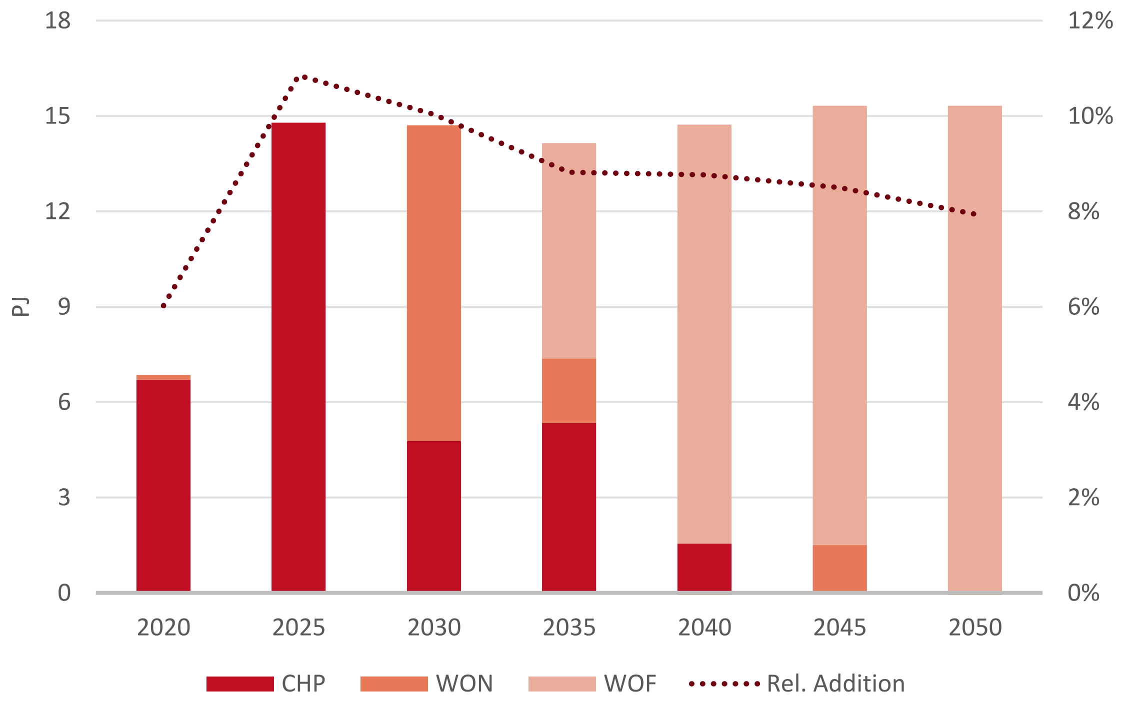


**Figure 35 –** Additional power production in FP scenario induced by large data centers

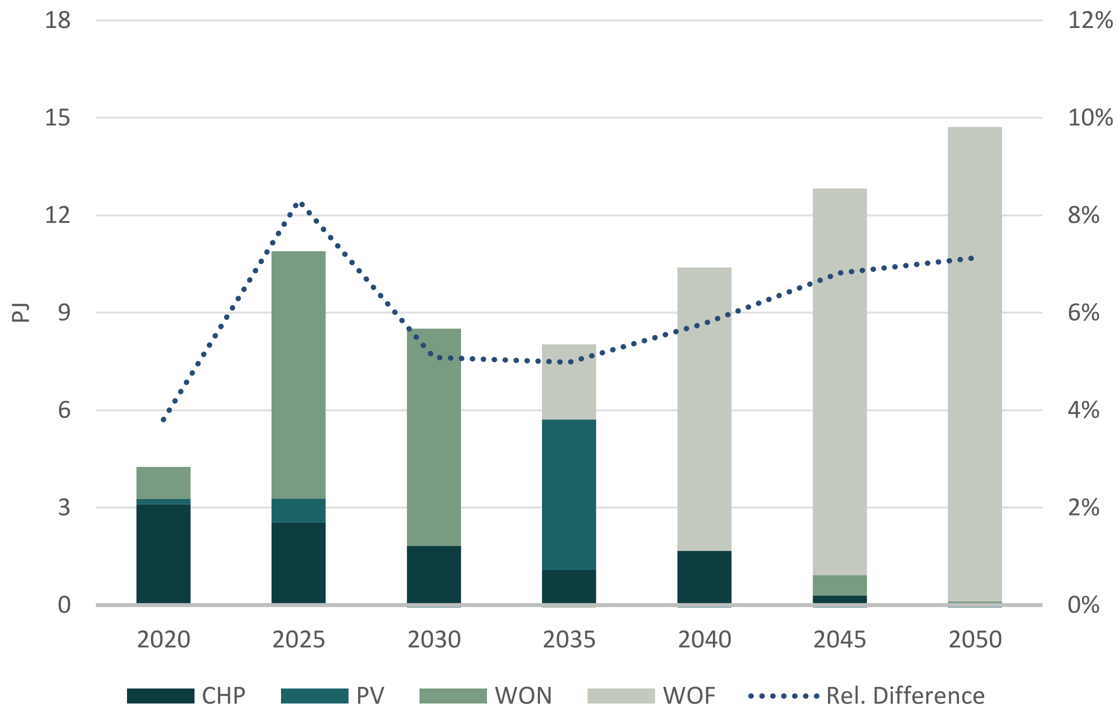


**Figure 36 -** Additional power production in Target scenario induced by large data centers

Although the electricity consumed by the large data centers can be entirely supplied with wind offshore plants in 2050 at the annual level, their instant power demand still requires programmable sources (such as biomass CHP), storage devices or imports. In fact, there is a mismatch between the additional power production from wind offshore plants and the data centers' power demand at the time-slice level in both scenarios (for clarity Figure 37 shows only the case of the FP scenario). As already discussed, TIMES-DK is not suitable for assessing demand/production matching problems precisely, but a discrepancy at the time-slice level certainly implies a mismatch at a finer temporal scale.

Another key point to be addressed is the shape of the two curves representing the additional power production in the two scenarios (Figure 35, Figure 36), which may be expected to follow the projected power demand of large data centers (Figure 23). However, this is true only in the Frozen Policy scenario, while in the Target scenario the additional electricity produced between 2025 and 2045 is lower than the annual power demand of large data centers. In fact, in the latter case their presence affects also the electricity demand of other two sectors: transports and industry (Figure 38, Figure 39).

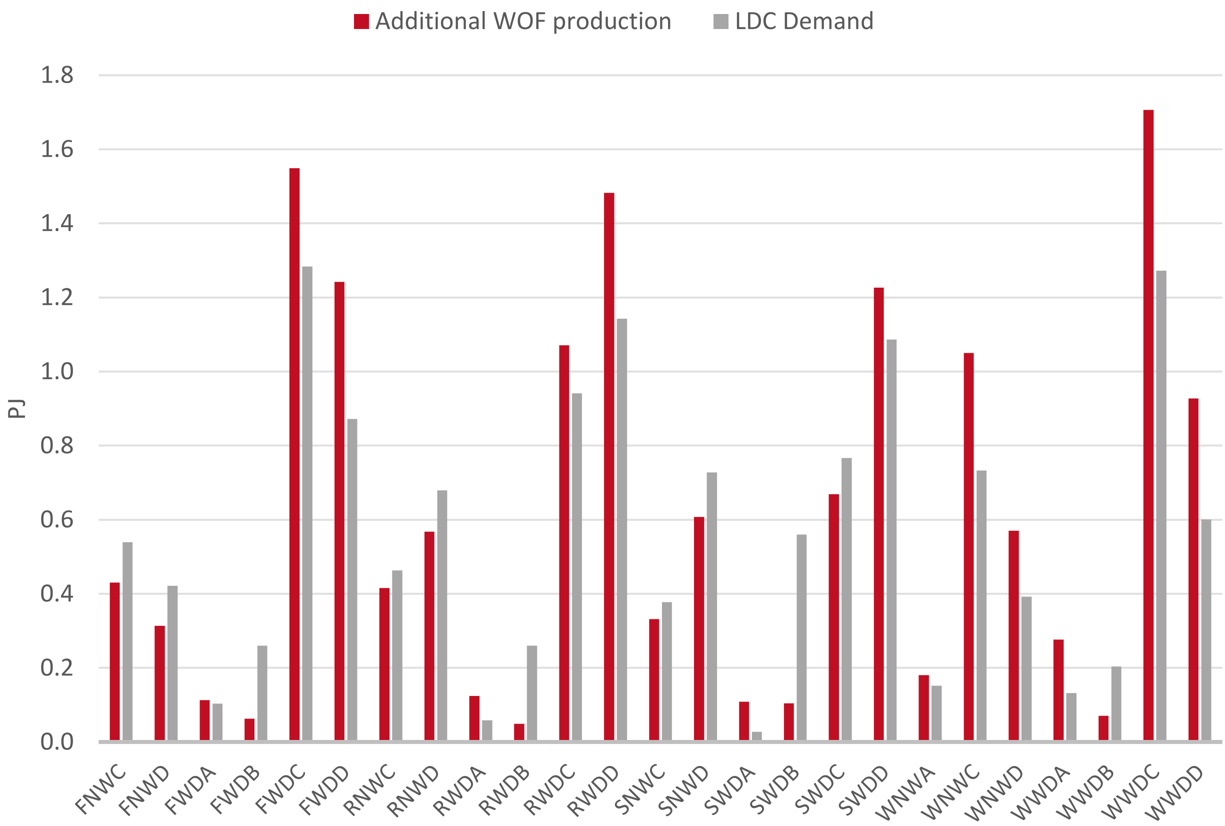


**Figure 37 –** Additional power production from wind offshore vs. LDC power demand in 2050 (FP scenario)

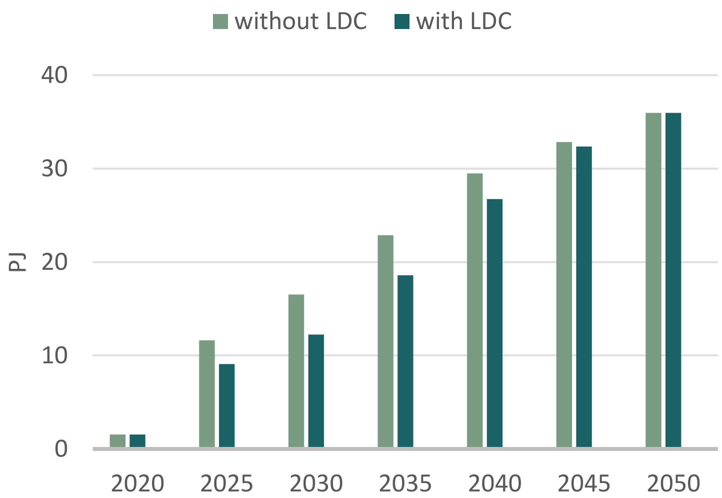


**Figure 38 -** Electricity consumption in transport sector (CF2050)

The higher influence is exerted on the transport sector, whose power demand faces a reduction of 24% in 2030. In fact, the electrification of the transport sector is delayed: it is not convenient to anticipate investments in renewable power capacity (as described in 5.1.1) and simultaneously increase the fleet of electric vehicles if there are no subsidies for renewables and the emission bound is set only in 2050 (features of the Target scenario). The power consumption in industry, instead, is slightly affected: the maximum relative reduction is only 3% in 2030. Causes are linked to the cost of electricity, which is higher in the period 2020-2040 due to higher investments in power capacity.

Consequently, part of the industrial heating that relies on electricity (heat pumps) is replaced by gas-fired technologies.

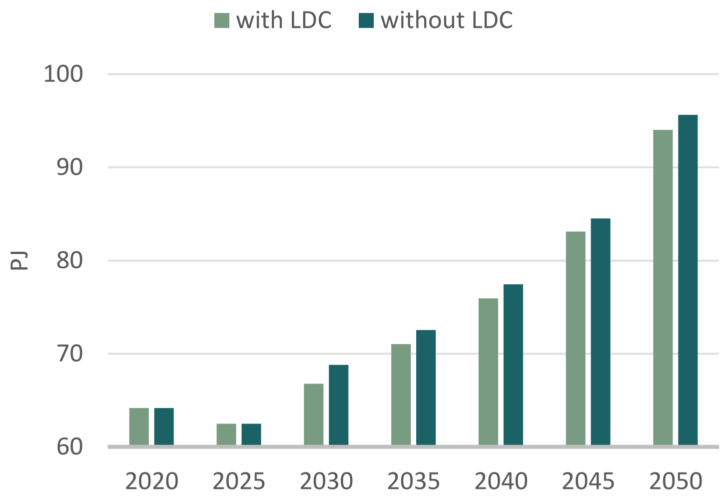


**Figure 39 -** Electricity consumption in industrial sector (CF2050)

## 5.2 Effects on $CO_2$ emissions

The additional $CO_2$ emissions induced by large data centers in both scenarios are a direct consequence of the power mix used to satisfy the surplus demand and the changes in the industrial and transport sectors (Figure 40, Figure 41).

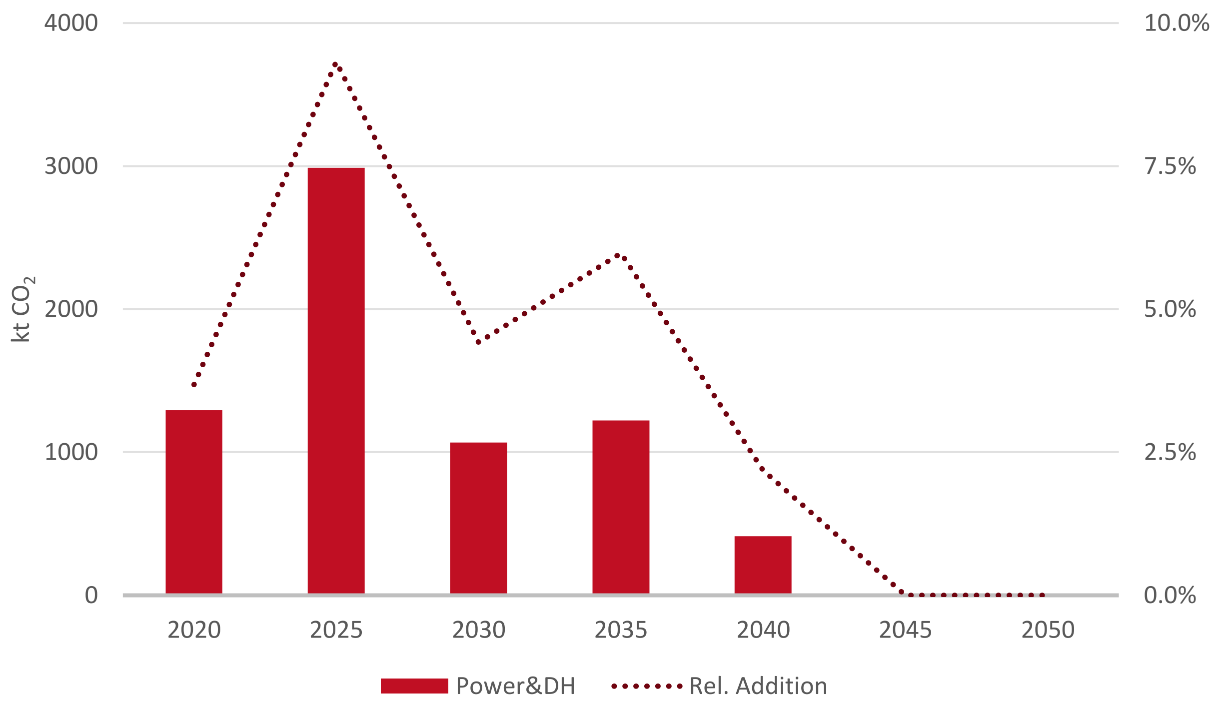


**Figure 40 –** Additional $CO_2$ emissions in FP scenario induced by large data centers

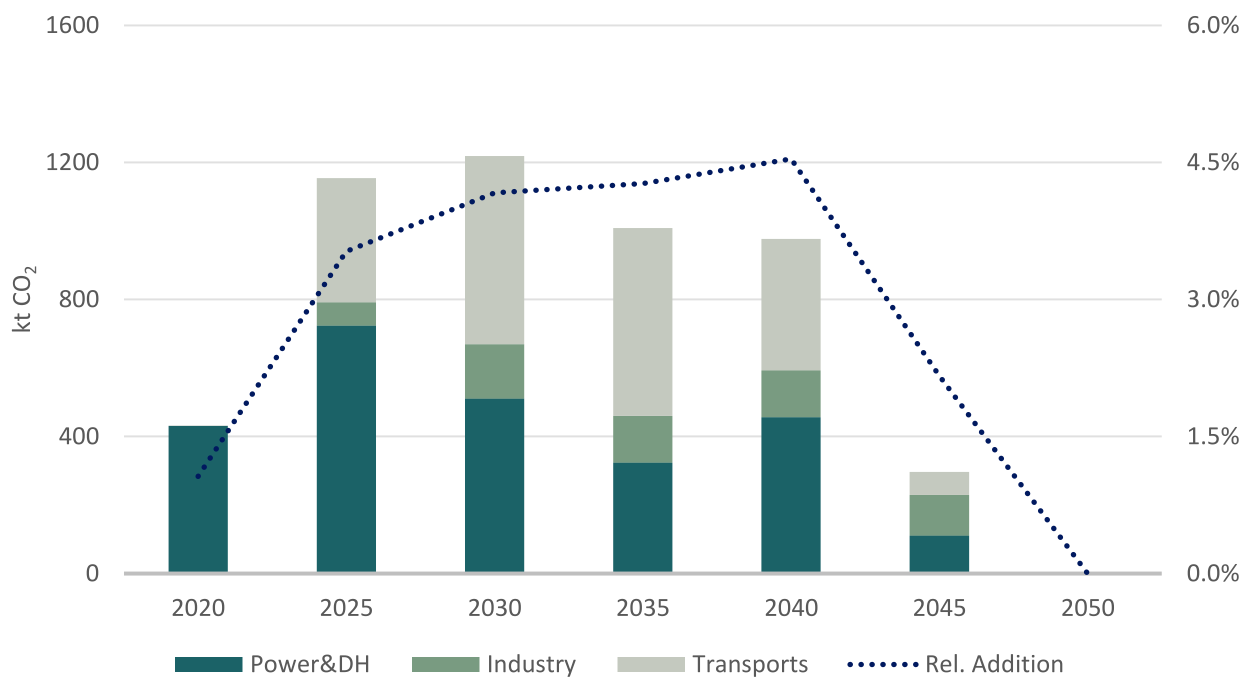


**Figure 41 –** Additional $CO_2$ emissions in Target scenario induced by large data centers

In the Frozen Policy scenario, the additional emissions are only caused by the extended use of CHP plants using coal as fuel (Figure 42). In fact, until 2040 it is still convenient to exploit the coal plants installed in the previous years, even though they are steadily dismissed towards 2050. Consequently, in the period 2020-2040 the average load factor of coal plants increases from 23% (FP without data centers) to 33% (FP with data centers). The peak of additional emissions is reached in 2025, when the presence of large data centers induces a 9% increase of $CO_2$ emissions in the FP scenario. However, it should be noticed that Denmark may avoid them by importing more electricity as mentioned before, although this choice may further increase the system costs.

In the Target scenario, instead, the extra emissions are not only related to the power sector, but also to industry and transports, even though their peak is halved with respect to the FP scenario. Coal-fired CHP plants and the lower electrification in the transport sector are the main causes of the additional emissions in the Target scenario with the data centers. The discussion regarding coal-fired CHP plants is similar to the one made for the FP scenario, while it still must be assessed what type of vehicles are responsible for higher emissions. These are vans (for freight transport) and cars (for passenger transport). In particular, the electrification of the vans with batteries is contrasted by vans fueled with a diesel blend (diesel + biodiesel), while full electric cars compete with hybrid ones, fueled with both electricity and a gasoline blend (gasoline + bioethanol). Figure 43 presents how the penetration of these vehicles is influenced in the Target scenario by the presence of large data centers.

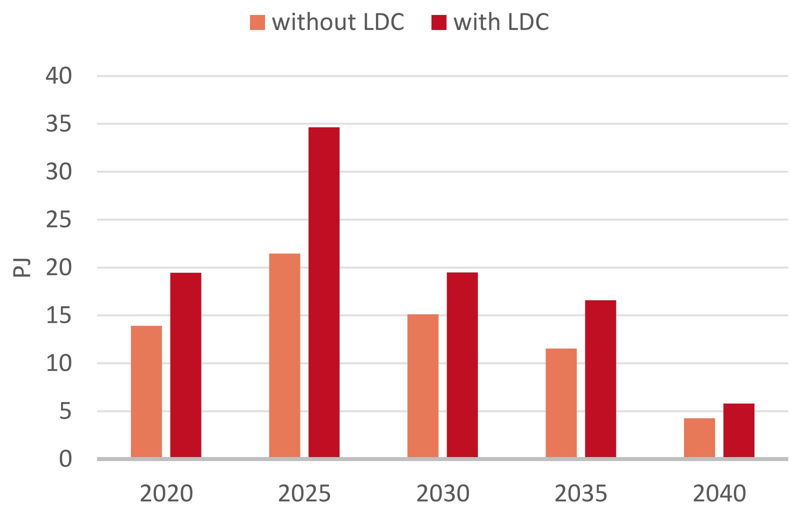


**Figure 42 –** Power production from coal plants in FP scenario

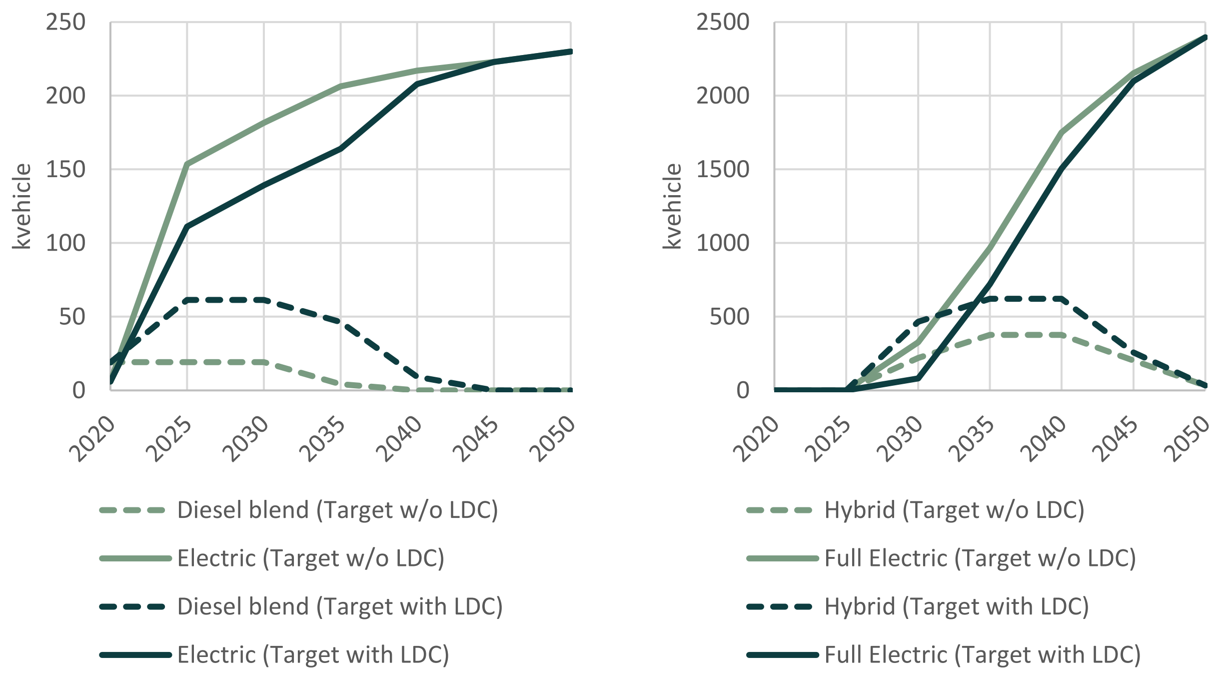


**Figure 43 –** Competition between electric and fossil fuel vehicles in CF2050 scenario with and without LDC: a) vehicles for freight transport (vans); b) vehicles for passengers transport (cars)

Finally, the additional $CO_2$ emissions caused by the industrial sector are due to the competition between electric heat pumps and absorption heat pumps supplied with natural gas. As Figure 44 shows, the presence of large data centers in the Target scenario advantages the latter type of heat pumps in industry.

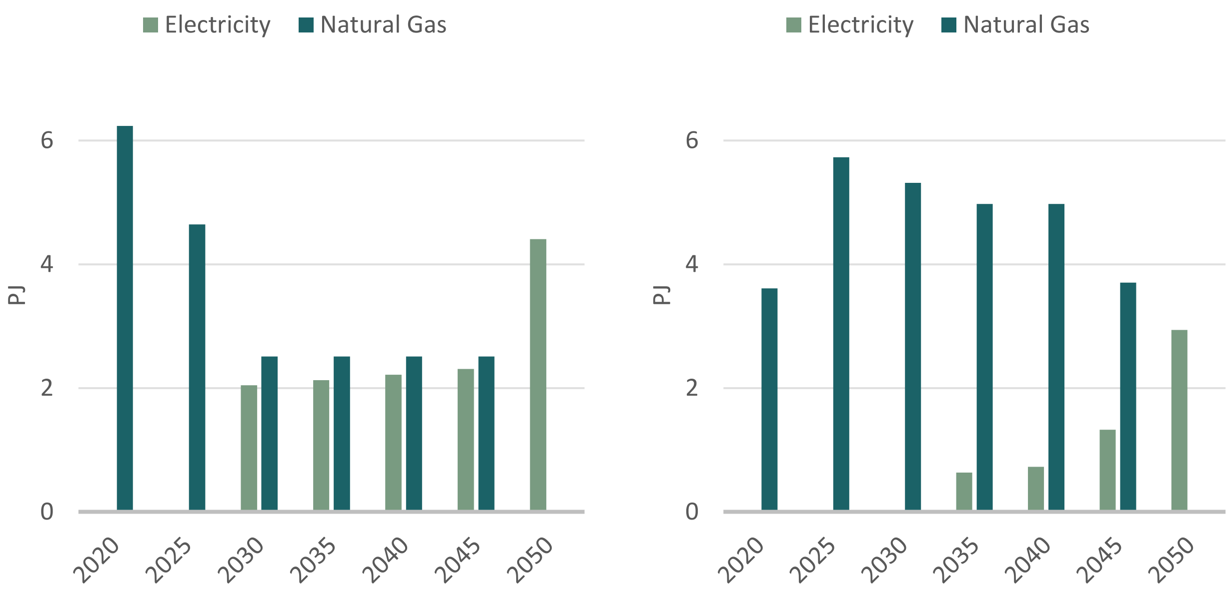


**Figure 44 –** Fuel consumption of industrial heat pumps: a) CF2050 without LDC; b) CF2050 with LDC

## 5.3 Waste heat recovery

In TIMES-DK, all the processes able to supply heat to central ("HETC") and decentral ("HETD") district heating grids are grouped into three sets: CHP plants, heating-only plants (HPL) and technologies recovering excess heat from industrial processes (EXH). Heating-only plants are largely preferred to produce *HETD*, while CHP plants are the main contributors to *HETC* (Figure 45). In fact, decentral DH areas host few industrial processes close to DH grids and few CHP plants, which are instead convenient in central DH areas characterized by higher specific heating demands. Notice that heating-only plants comprehend ambient air heat pumps, traditional boilers and renewable plants (solar and geothermal).

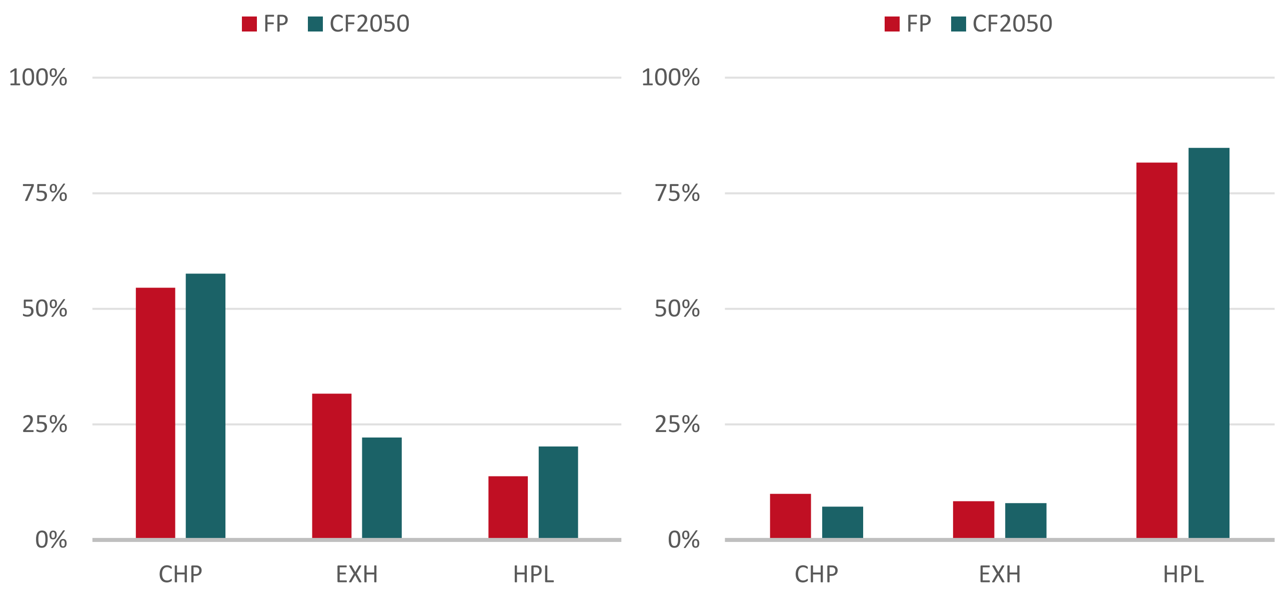


**Figure 45 –** Technologies used for producing distict heat: a) central; b) decentral

In both scenarios the total demand for district heating is the same through the years. Therefore, the waste heat from large data centers enters in competition with CHP, EXH and HPL technologies, while does not affect residential heating technologies. As Figure 46 and Figure 47 show, both scenarios confirm that a full waste heat recovery is convenient in 2050, but relevant differences are present in the other years.

Results indicate that in both scenarios large data centers' waste heat competes with a fraction of the heat produced by electric ambient air heat pumps. Certainly, large data center's waste heat has a higher temperature (31°C in average) compared to outdoor air (8°C annual average). Therefore, the efficiency of electric heat pumps connected to data centers is definitely higher (+25% according to [31]). However, also their investment and fixed O&M costs are higher compared to ambient air heat pumps (+11-14% according to [31]) because they require an additional heat exchanger to extract the waste heat.

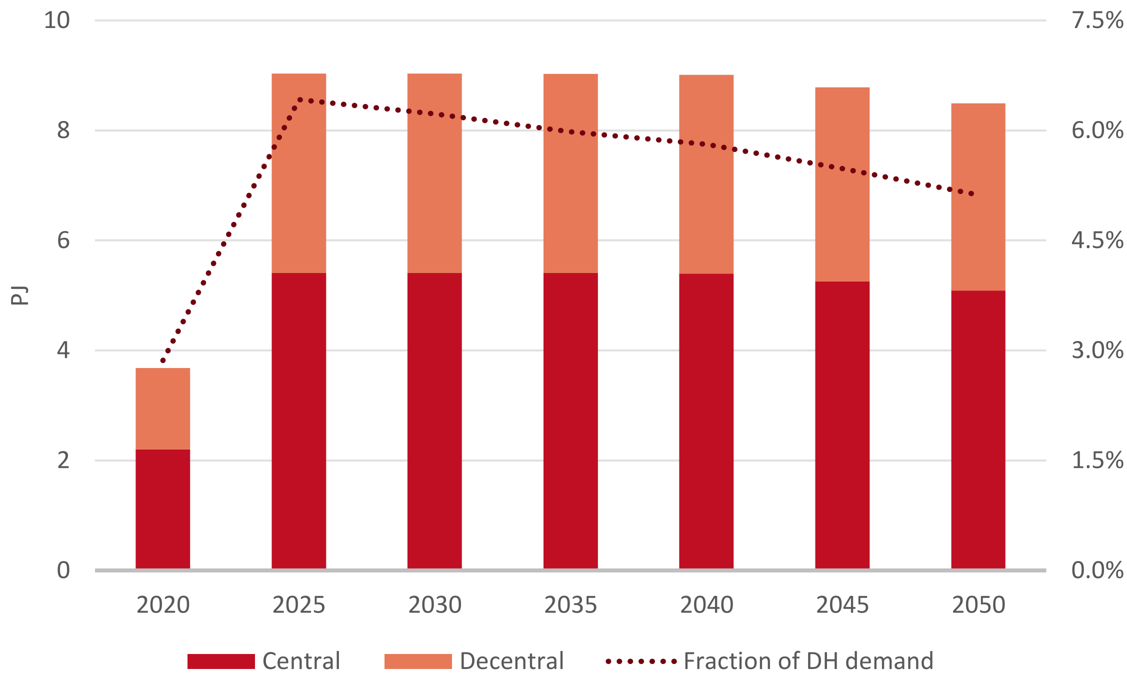


**Figure 46 –** Reuse of large data centers' waste heat in DH (FP scenario)

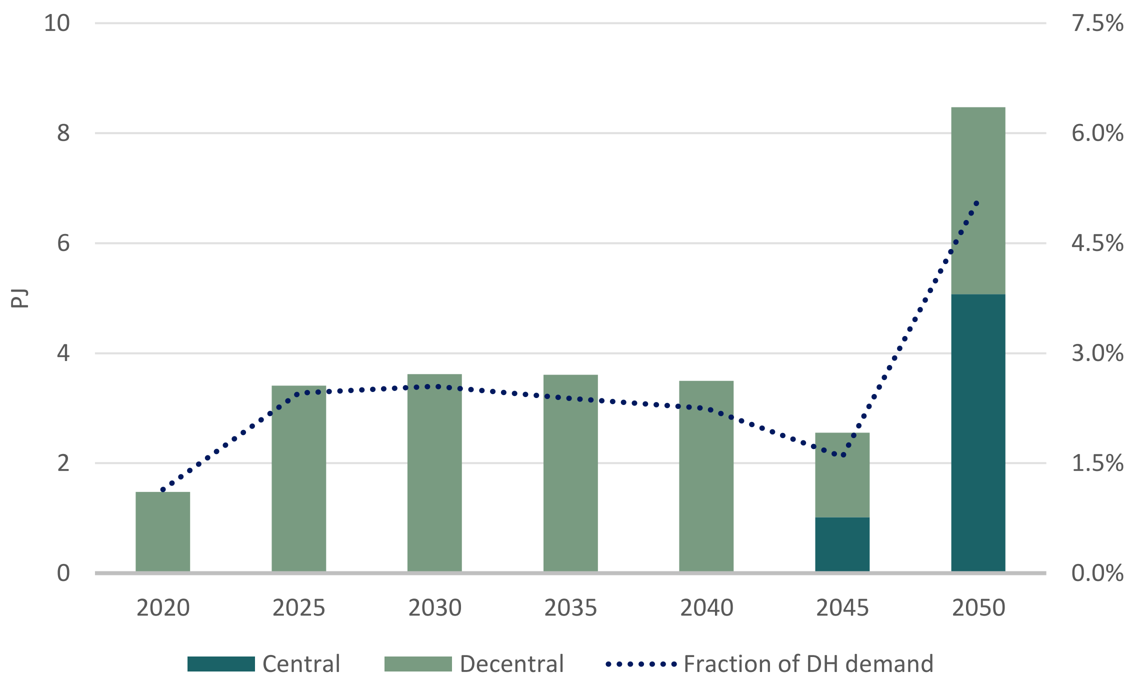


**Figure 47 –** Reuse of large data centers' waste heat in DH (CF2050 scenario)

Consequently, both *FP* and *CF2050* indicate that large data centers' waste heat reuse is always part of the optimal solution in decentral DH areas. Nevertheless, this is not true for central DH areas. The reason lays behind two factors: the marginal price of electricity and the thermal losses of the DH grid. In TIMES-DK, the annual thermal losses in central and decentral networks are respectively 15% and 22% of the annual heat produced. Therefore, in a framework of lower thermal losses (i.e. central grid) and lower cost of electricity (i.e. in *CF2050* scenario) the use of ambient air heat pumps is preferred to the recovery of large data centers' waste heat. The constraint on $CO_2$ emissions in 2050, however, gives a chance also to central large data centers in the Target scenario, as demonstrated by Figure 47.

Moreover, two important aspects should be remarked. First, the need for CHP plants to satisfy the power demand of large data centers (see 5.1.2) further penalize heating-only plants, in particular in the FP scenario where the heat coming from HPL technologies is reduced by 16% in 2025 and 2030 (if compared to the FP without data centers).

Second, TIMES-DK includes excess heat from biorefineries. However, this quantity is not affected by the presence of large data centers, not even in the Target scenario, where fuel blends are used to delay the electrification of the transport sector (see 5.2). In fact, the model has the freedom to set the optimal share of bio and fossil fuels for diesel and gasoline blends. As a result, the additional demand of fuel blends in the period 2020-2040 is satisfied with only fossil fuels, which are more convenient in the Target scenario thanks to the absence of taxation.

Summarizing, the highest potential for waste heat recovery corresponds to roughly 5% of the total heat produced for district heating. However, this number is only useful to understand the magnitude of their potential from a system perspective. In fact, the location of these facilities is well defined and the feasibility of their integration in the local district heating network should be evaluated according to the needs of the hosting municipalities.

## 5.4 Marginal costs

The marginal price of a commodity coincides with the cost for producing that commodity under perfect market equilibrium conditions (i.e. the situation reached by TIMES models). It represents the increase of the system cost for producing an additional unit of that commodity. Therefore, monitoring this quantity allows to assess how easily the system can face an increase in the demand of that commodity. For example, when the electricity demand is increased by one unit there are two possible outcomes: “either the system produces one more unit of electricity, or else the system consumes one unit less of electricity (perhaps by choosing more efficient end-use devices or by reducing an electricity-intensive energy service, etc.)” [17].

### 5.4.1 Electricity

Figure 48 presents the trend of electricity marginal prices (annual average) in the Frozen Policy scenario with and without the large data centers. The general increase towards 2040 is due to two factors: on the one hand, the increase of the power demand through the years may cause higher marginal prices depending on the technologies used to satisfy the demand (as it is witnessed by the inverse supply curve in Figure 4); on the other hand, investments in new capacity generate higher fixed annual costs. The general drop in the marginal price after 2040, instead, is linked to an oversized power capacity able to sustain higher power demands.

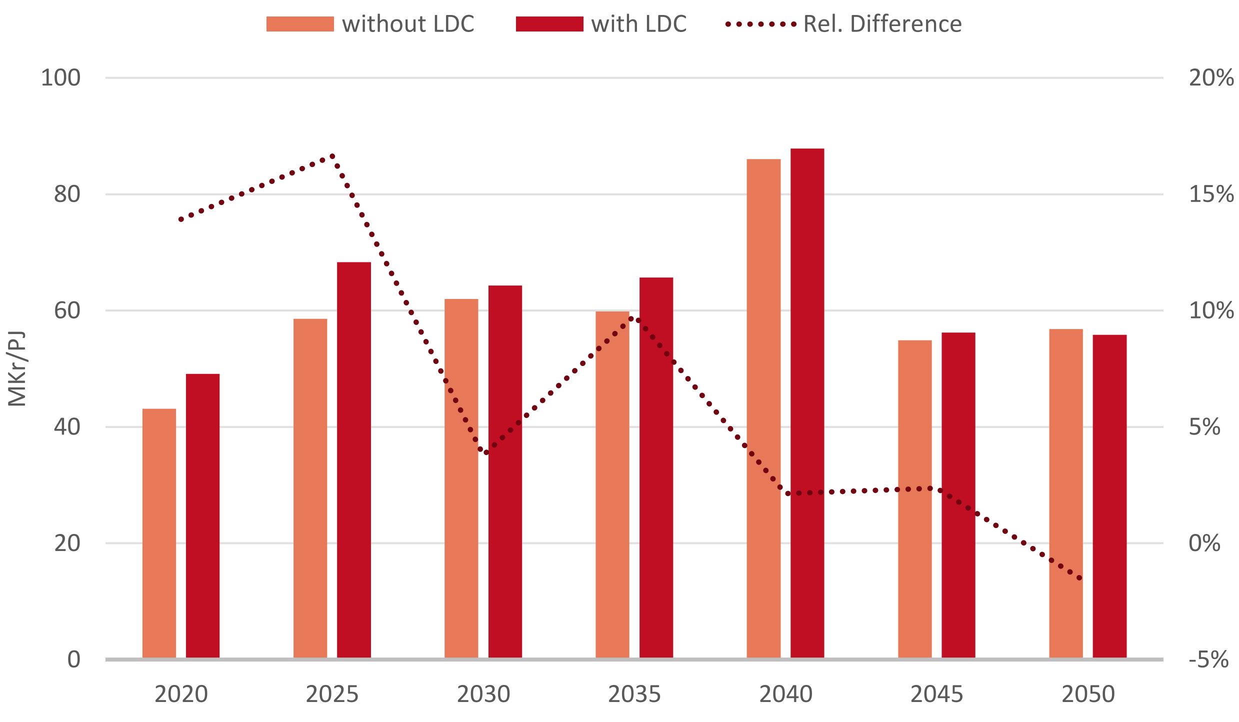


**Figure 48 –** Marginal price of electricity in FP scenario

When evaluating the effects of the large data centers on marginal electricity prices, it is visible that they induce an increase in the period 2020-2035, when CHP plants are required to produce more electricity. In fact, the highest relative difference (+17%) is registered in 2025, when the additional annual demand of large data centers is fully covered with CHP plants (Figure 35). When the system shifts more resolutely towards renewable sources, instead, the presence of large data centers has a minimal effect on the marginal price of electricity.

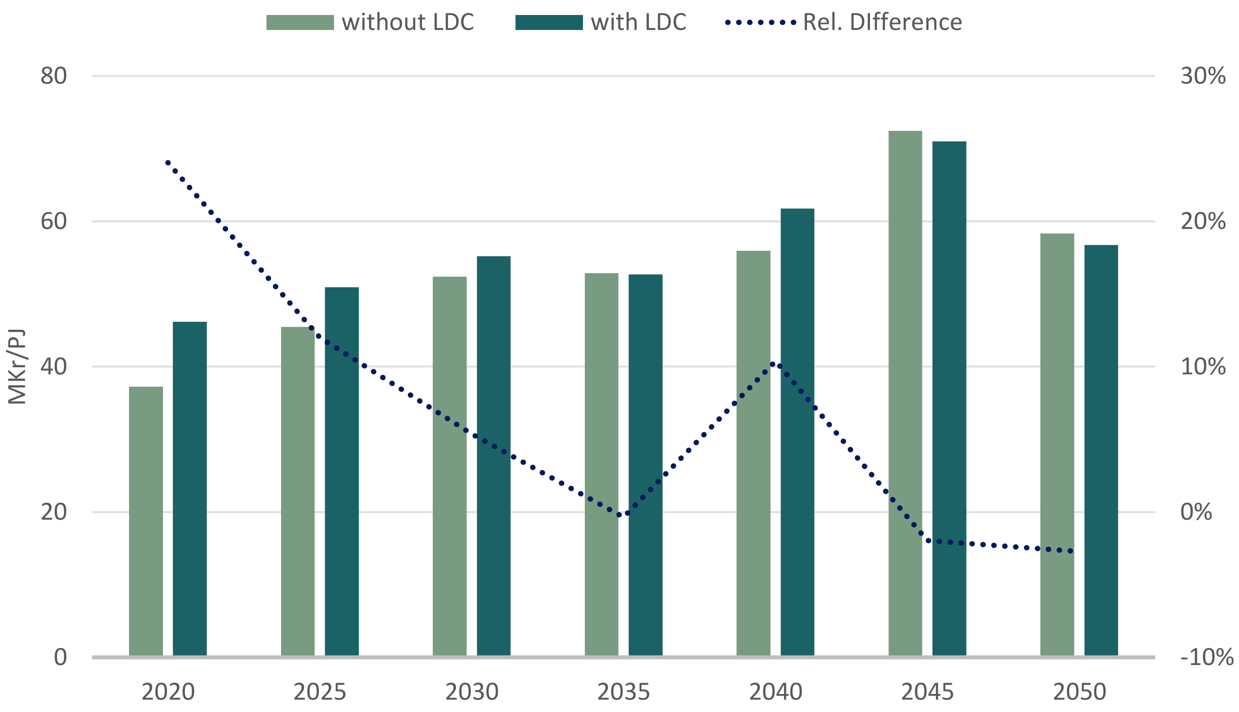


**Figure 49 –** Marginal price of electricity in CF2050 scenario

Figure 49 instead shows what happens in the Target scenario. The general increase of price culminates in 2045 (instead of 2040) because the emission bound in 2050 forces the model to choose more expensive investments. Reasons behind the general trend of the marginal price are the same presented for the FP scenario. However, it should be noticed that the absence of taxes in the Target scenario causes lower prices with respect to the FP scenario.

When evaluating the effect of large data centers in the Target scenario, it is immediately evident that the need to resort to CHP plants in 2020 is extremely expensive (+24% price increase). Additional marginal costs then steadily decrease as the role of CHP plants diminishes, with an exception in 2040 when a higher amount of electricity coming from this type of plants is again required (see Figure 36). However, it should be noticed that electricity imports could be a viable alternative in reality, even though they may still cause a cost increase, as already discussed in 5.1.2.

In conclusion, also the behavior of marginal electricity prices indicates that the power system is definitely ready to sustain the presence of large data centers from 2040 upwards, but the transition towards that period poses some issues.

### 5.4.2 District heating

The marginal price of district heat has the same meaning of the one for electricity, but it is referred to the sum of the heat produced by central and decentral technologies (i.e. *HETC* and *HETD*).

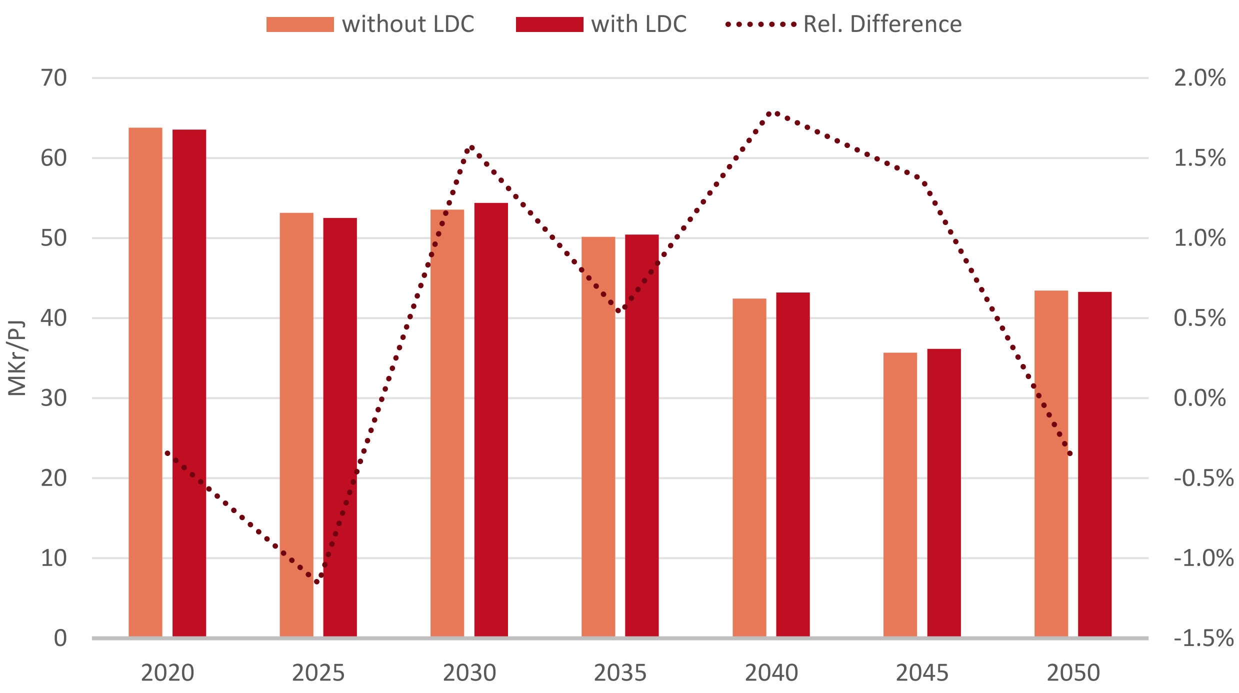


**Figure 50 –** Marginal price of district heat in FP scenario

Figure 50 presents the trend in the Frozen Policy scenario: the general decrease in price implies that the plants' thermal capacity is able to face the increasing heating demand from DH without issues. In this scenario, the presence of large data centers does not affect evidently the marginal prices of district heat. In fact, the waste heat from data centers substitutes a small fraction of the heat produced by ambient air heat pumps (see 5.3), thus it does not affect significantly the costs for deploying that technology.

In the Target scenario (Figure 51) a general decreasing trend is also present, but with a notable increase in 2050 because the cheaper $CO_2$-emitting technologies are completely banned. In this framework, the possibility to produce heat by simply investing in heat pumps and heat exchangers (for harvesting large data centers' waste heat) is more convenient than other solutions. In fact, in 2050 the district heating marginal price can be reduced by 6% if data centers are present. However, in the previous periods advantages from waste heat reuse are not so evident because of the heat coming from CHP plants: they are forced to produce an additional quota of heat because they partly cover large data centers' additional power demand (Figure 36), but this supplementary heat is more expensive. Consequently, the marginal price of district heating may even increase by 5%.

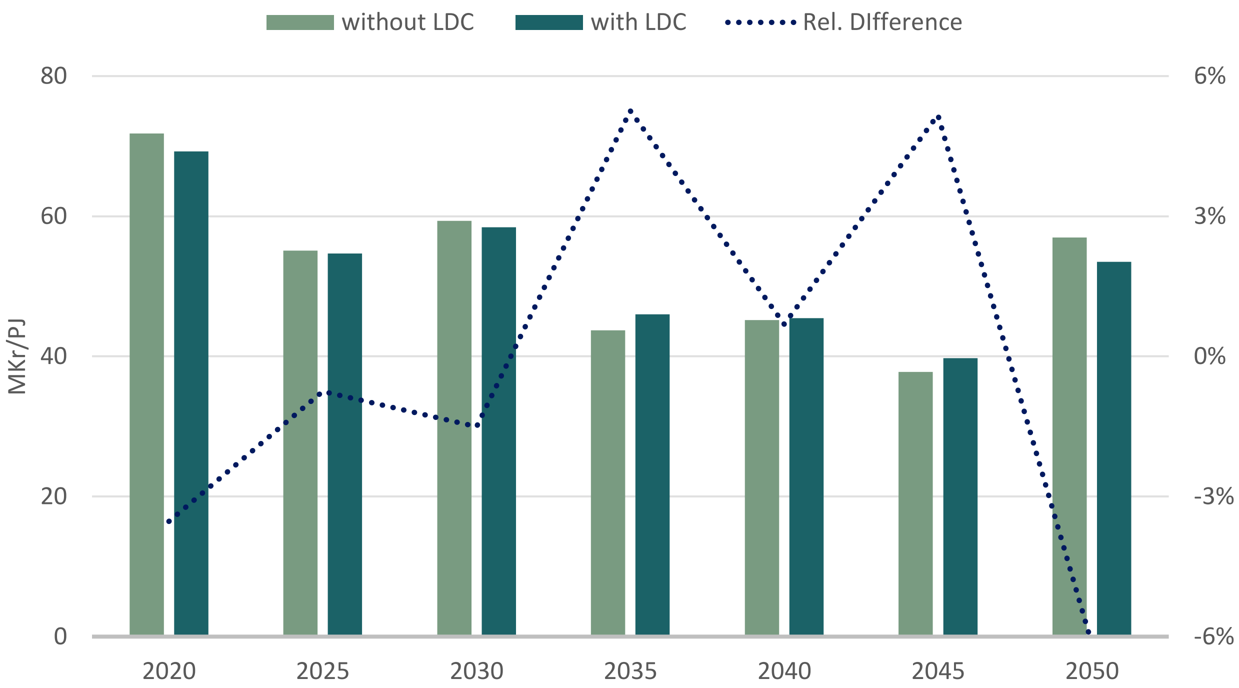


**Figure 51 –** Marginal price of district heat in CF2050 scenario

# Chapter 6

# Sensitivity analyses

Models supporting long-term scenarios analysis are usually characterized by elevated levels of uncertainties. In this application, the greatest uncertainty in modeling large data centers in TIMES-DK is linked to their expected power consumption. Therefore, the purpose of this chapter is to discuss how the optimal solution proposed by the model changes when the driver (i.e. the power demand) is modified. As already presented in 4.3.2, two additional demand projections are considered. In the former a daily variable workload profile is applied to the forecast made by Energinet, thus reducing the annual power demand. The latter, instead, is based on a different forecast (from the Danish Energy Agency) and the same variable workload profile, which cause a higher annual power demand compared to Energinet's projection starting from 2035.

Therefore, this chapter develops some sensitivity analyses on the four themes identified in the previous chapter, comparing the new outcomes with those obtained using Energinet's projection (discussed in Chapter 5). Comparisons are usually presented in relative terms, however, when absolute values are shown, the three projections are referred as:

- "Low LDC demand" (Energinet's projection + variable profile)
- "Base LDC demand" (Energinet's projection)
- "High LDC demand" (Danish Energy Agency's projection + variable profile)

In the end, some significant quantities belonging to the entire time horizon are introduced together with their uncertainties.

## 6.1 Sensitivity on the power sector

The two most important results of section 5.1 are here confirmed: variations in the annual power demand of large data centers are reflected on the power production from CHP plants in the period 2020-2035 and from wind offshore turbines in the period 2035-2050. In addition, the other two technologies (wind onshore and PV) display the same pattern regarding new power capacity installations compared to those obtained in 5.1.1. In fact, in both scenarios similar installations (same size) are made in the same period. Therefore, conclusions achieved in 5.1.1 are still valid, even when the additional power demand due to large data centers is varied.

Figure 52 and Figure 53 clearly show how an additional power production in the system relies on CHP plants between 2020 and 2035 and on wind offshore between 2035 and 2050 (for clarity only the *FP* scenario is presented, but *CF2050* displays the same pattern).

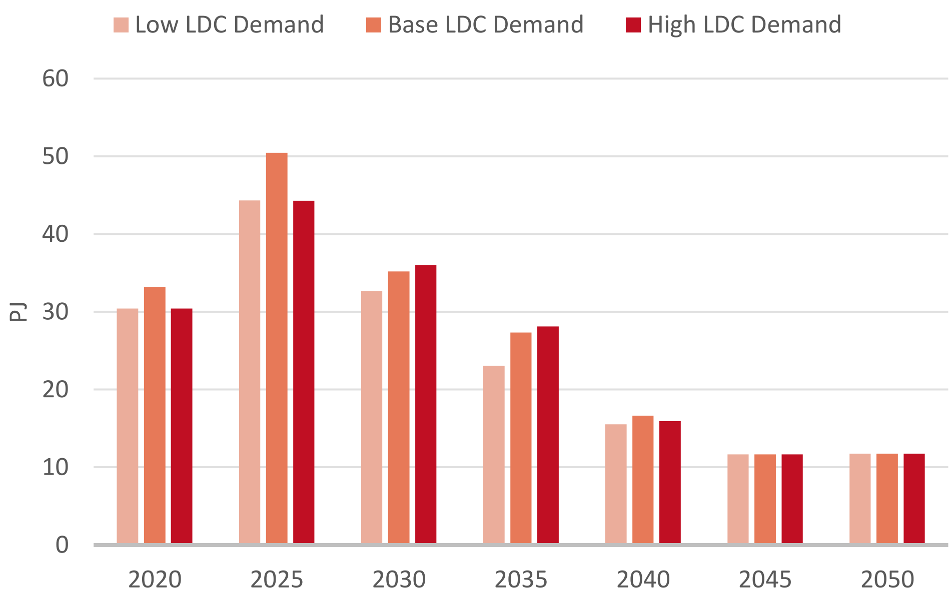


**Figure 52 –** Power production with CHP in FP scenario with different LDC power demands

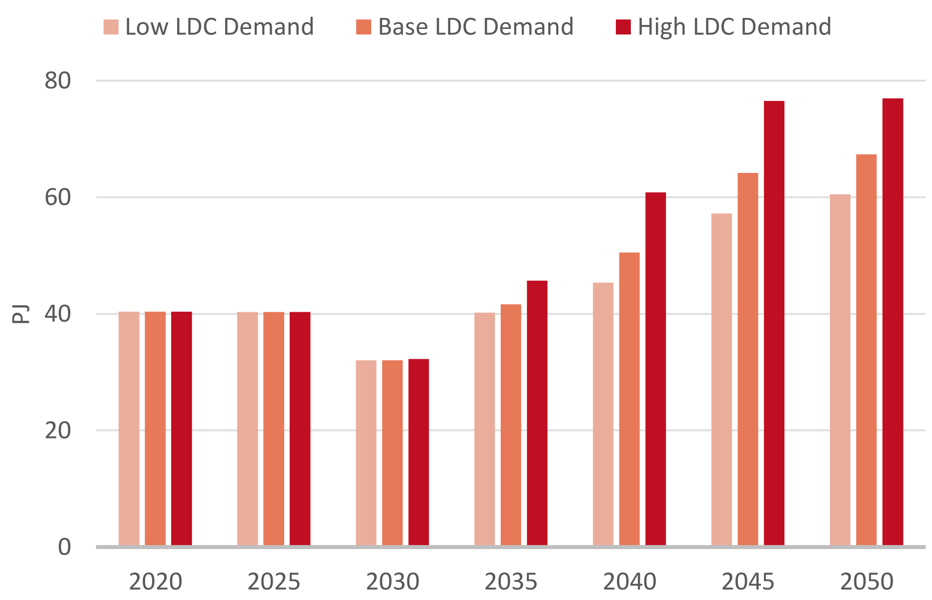


**Figure 53 –** Power production with WOF in FP scenario with different LDC power demands

Therefore, sensitivity analyses were developed on two different periods by summing the power demand of large data centers in 2020-2035 and in 2035-2050. Their relative values are presented in Table 8.

**Table 8 –** Large data centers' power demand variation in two periods (2020-2035 and 2035-2050) in different cases

| Case | LDC Power demand (2020-2035) | LDC Power demand (2035-2050) |
|---|---|---|
| *Base demand* | 100% | 100% |
| *Low demand* | 59% | 59% |
| *High demand* | 89% | 148% |

Consequently, Figure 54 and Figure 55 show how the power production of CHP plants and wind offshore is affected by changes in large data centers' power demand through the interpolation of the outcomes of the three simulations for each scenario.

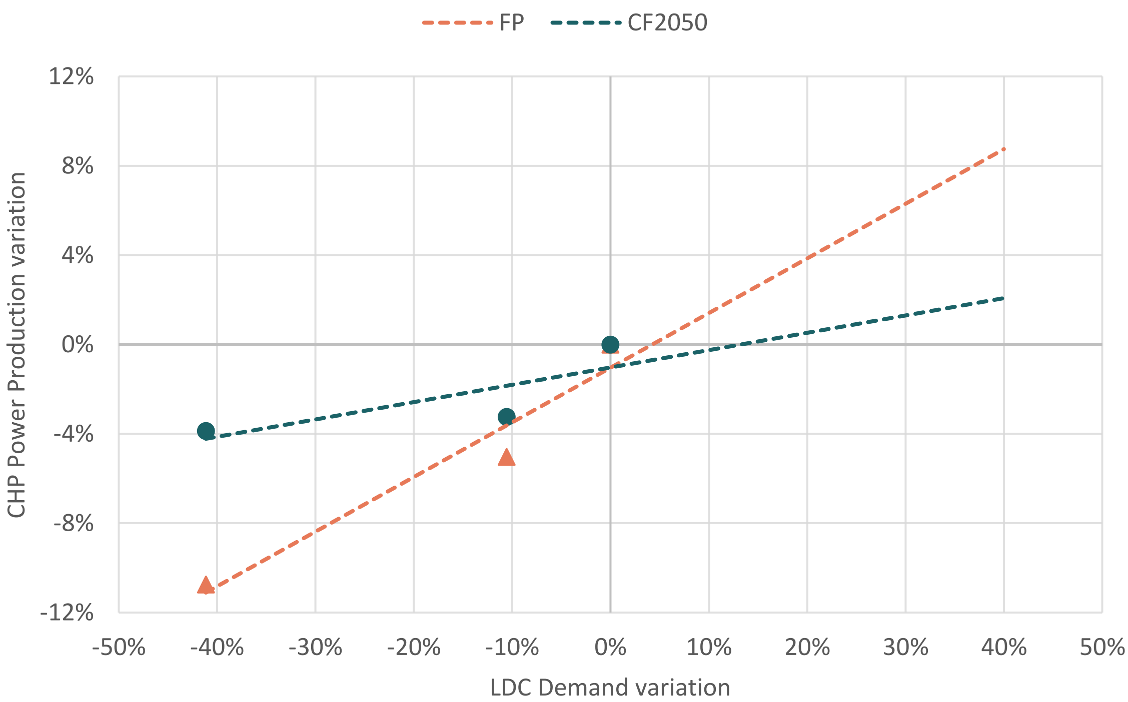


**Figure 54 –** Sensitivity analysis on total CHP power production in the period 2020-2035

The relevance of CHP plants is more evident in the Frozen Policy scenario, where a power demand variation of 10% in the period 2020-2035 causes a power production variation from CHP of around 2.4% in the same period. Moreover, in the FP scenario the linear trend seems a good interpolation ($R^2$ = 0.95). A sensitivity analysis in the Target scenario, instead, confirms the interrelation between large data centers and other sectoral power demands. In fact, the power production from CHP is rather stable (and other technologies do not change their output). This implies that the power demand of other sectors is also affected. Consequently, a linear interpolation cannot grasp entirely the relation between large data centers' demand and CHP power production ($R^2$ = 0.64).

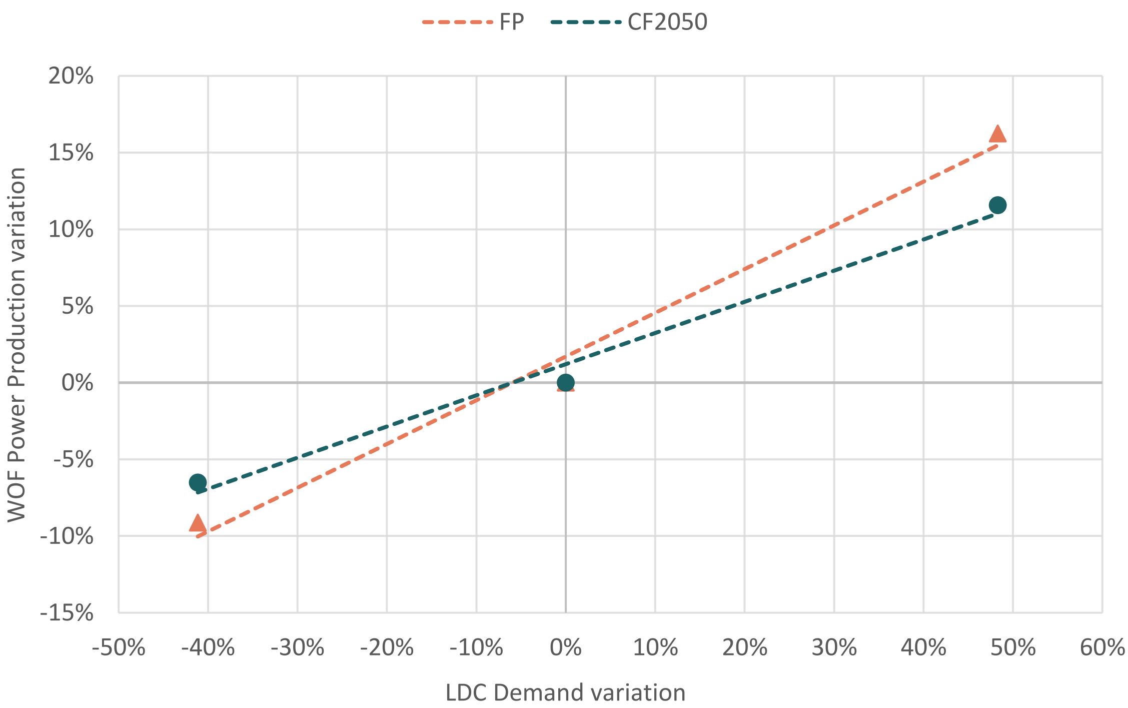


**Figure 55 –** Sensitivity analysis on total WOF power production in the period 2020-2035

Wind offshore technologies, instead, confirm to be the only solution on which the model relies for satisfying additional power demands in future. In fact, a power demand variation of 10% determines a power production variation from wind offshore turbines of 2.8% (in *FP*) and of 2% (in *CF2050*) in the period 2035-2050. Moreover, the linear interpolation is a good approximation for both scenarios ($R^2$ = 0.98). However, the power production from wind offshore in the Target scenario is again more stable compared to the FP scenario. In fact, power demands of other sectors are still affected until 2045, but in minor quantity (consequently $R^2$ of the linear interpolation is still close to 1).

## 6.2 Sensitivity on $CO_2$ emissions

$CO_2$ emissions induced by the presence of large data centers in the Danish energy system are also consistent with the pattern found in 5.2 and the variation of their power demand (remember that the "high demand" case displays a higher power demand only from 2035), as shown by Figure 56 and Figure 57. Moreover, these additional simulations also confirm that the main responsible for emissions in the Frozen Policy scenario is the power sector, while in the Target scenario industry and transports should also be accounted.

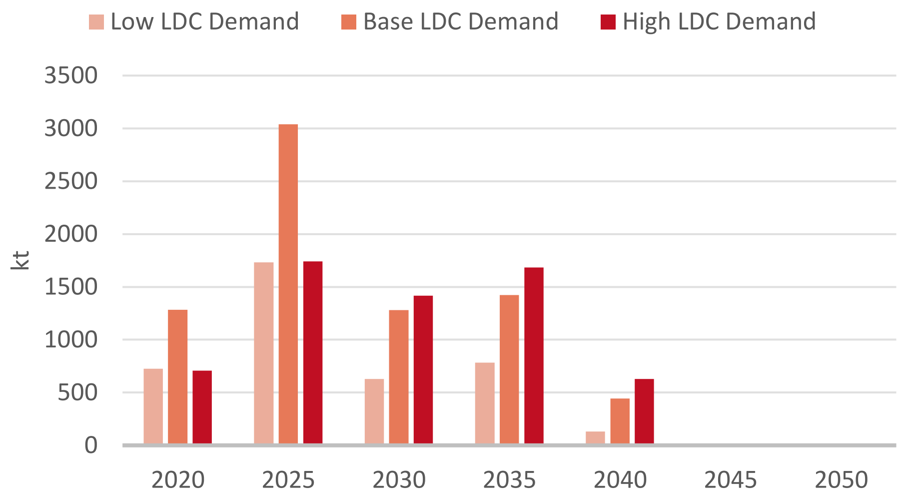


**Figure 56 –** Additional $CO_2$ emissions in FP scenario induced by different LDC demands

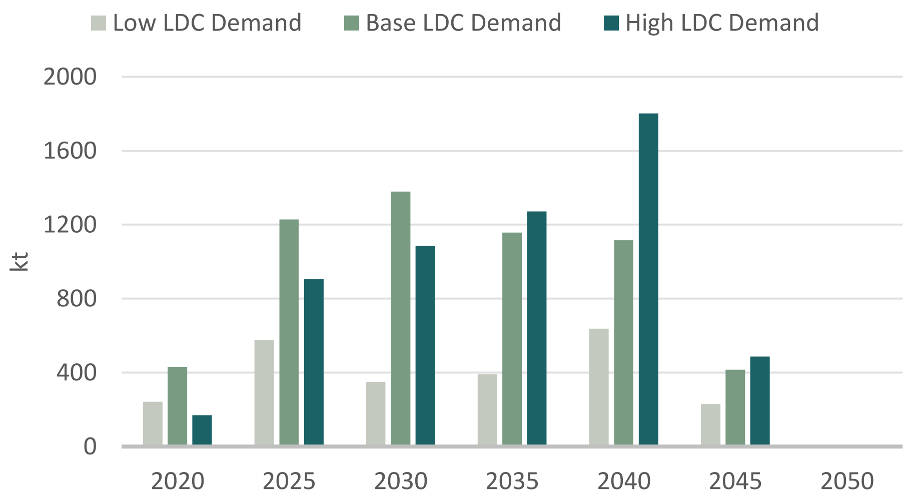


**Figure 57 –** Additional $CO_2$ emissions in CF2050 scenario induced by different LDC demands

The sensitivity analysis is here split into two periods (before and after 2035), as already indicated by Table 8. The reason is not only due to the major use of CHP until 2035, but also to the fact that the "high demand" case actually contains a higher power demand (with respect to the "base" case) only after 2035. Figure 58 and Figure 59 show how the cumulated $CO_2$ emissions of the Danish energy system are affected by different power demands of large data centers in these two periods.

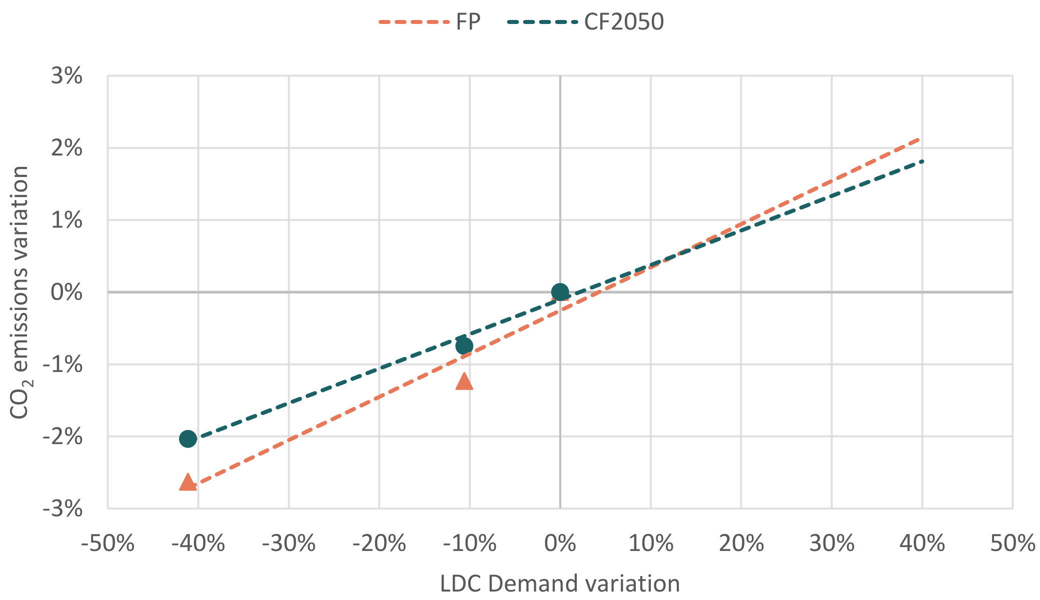


**Figure 58 –** Sensitivity analysis on $CO_2$ emissions in the period 2020-2035

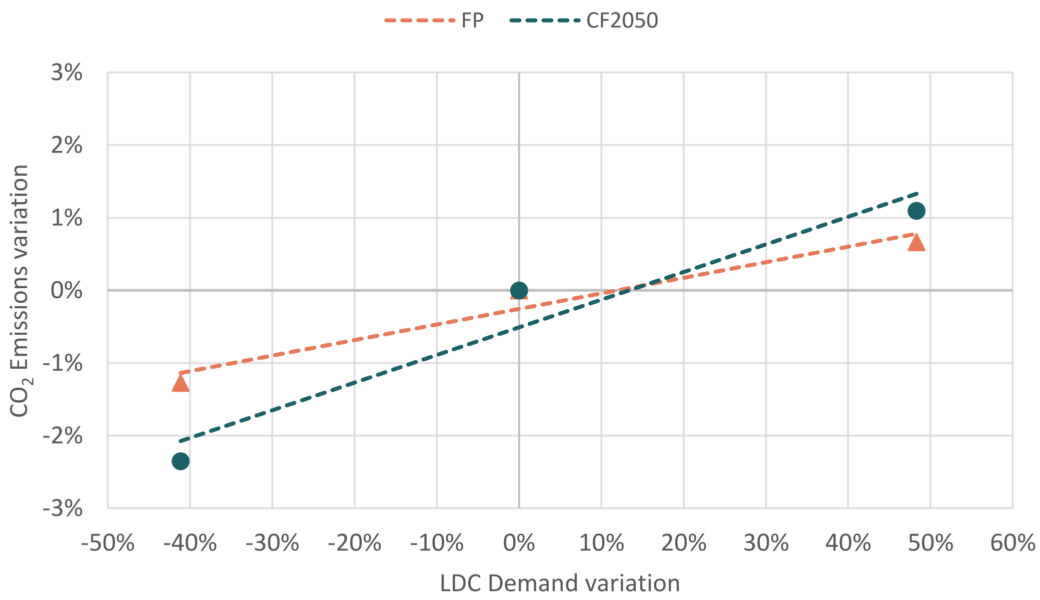


**Figure 59 –** Sensitivity analysis on $CO_2$ emissions in the period 2035-2050

## 6.3 Sensitivity on waste heat

Results from these additional simulations are also consistent with the findings in 5.3: in the Frozen Policy scenario it is always convenient to reuse entirely the waste heat (Figure 60), while in the Target scenario this is convenient only towards the end of the time horizon (Figure 61). However, when in the target scenario a higher amount of waste heat is available (such as in the “high demand” case), results show that it is already convenient in 2045 instead of 2050.

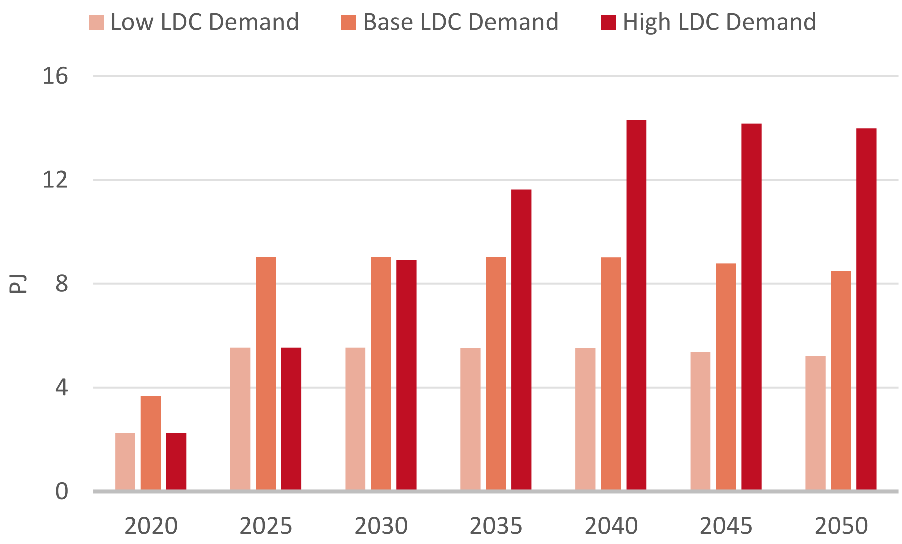


**Figure 60 –** Waste heat reuse in FP scenario with different LDC demands

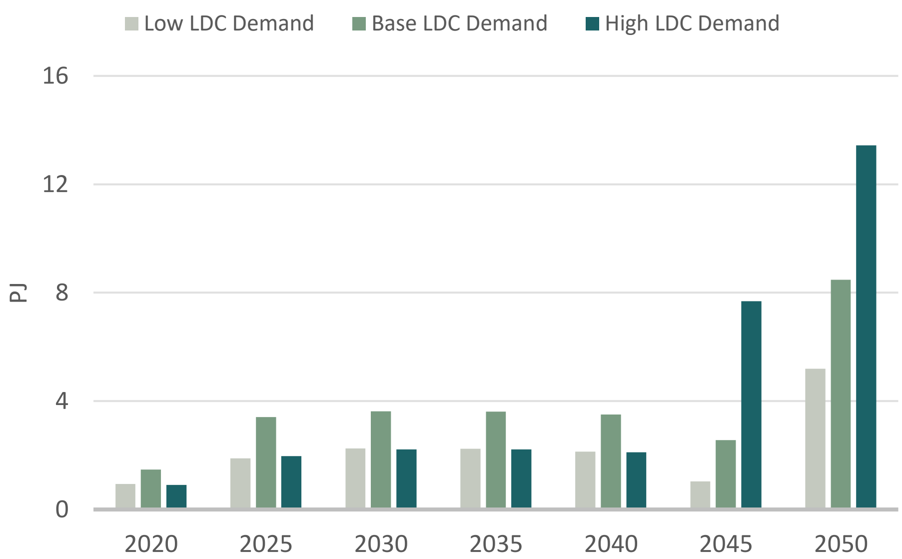


**Figure 61 –** Waste heat reuse in CF2050 scenario with different LDC demands

The total waste heat reused in the time horizon is directly proportional to large data centers' power demand in the FP scenario, while it is slightly below direct proportionality in the Target scenario, as witnessed by Figure 62. In fact, in the FP scenario its recovery is always fully exploited because it is always cheaper than electric ambient air heat pumps. In the Target scenario, instead, the conditions favorable to a full waste heat recovery are: a higher efficiency of heat pumps (especially if coupled with a higher availability of waste heat) and a boundary on $CO_2$ emissions (which necessary excludes fossil fuel technologies). These circumstances are only verified in the end of the time horizon.

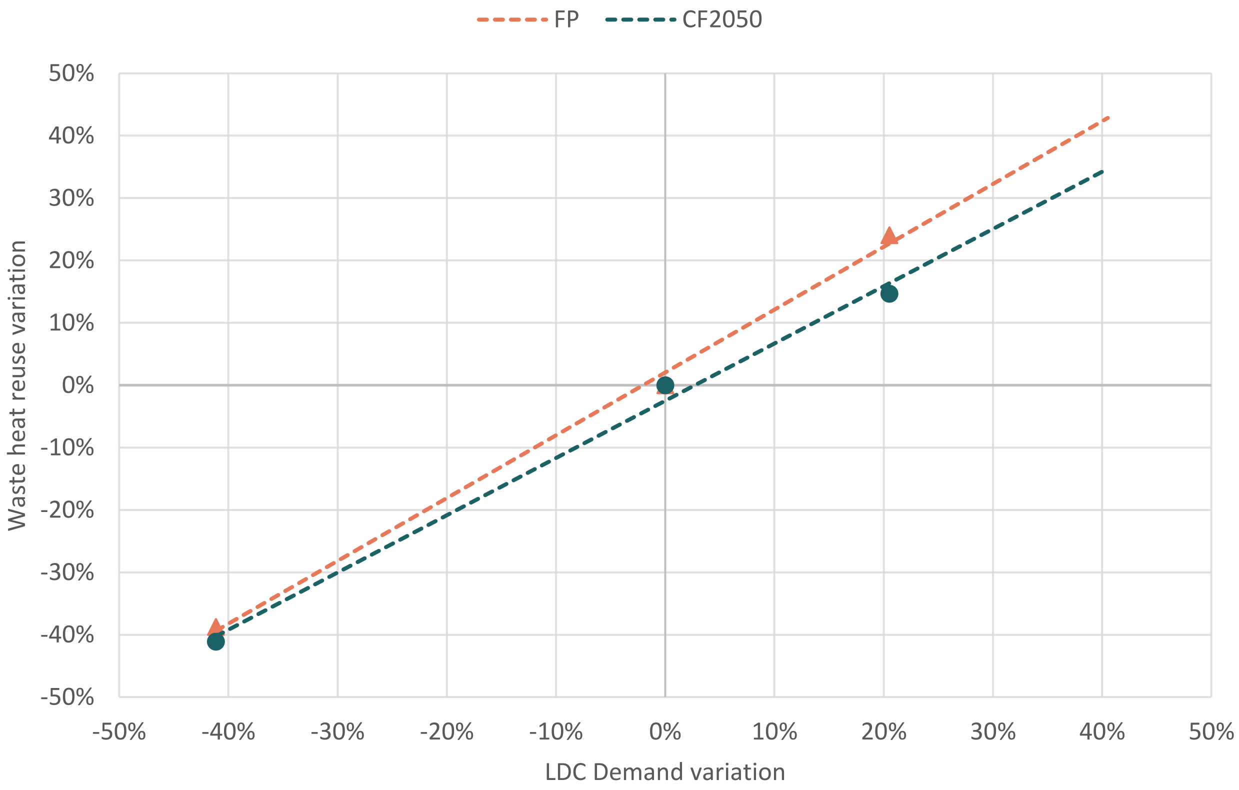


**Figure 62 –** Sensitivity analysis on waste heat reuse in the period 2020-2050

## 6.4 Sensitivity on marginal prices

As already discussed in 5.4, marginal prices of electricity and district heat depend on the technologies used to supply those commodities (i.e. on the configuration of their respective sectors) and on the demand they have to satisfy. Sections 6.1 and 6.3 proved that a change in large data centers' power demand does not affect the way in which the model chooses the technologies to deploy, but it influences only their activity. Consequently, marginal prices are not very different from the "base" case. In particular, district heat marginal prices remain practically unchanged because large data centers' waste heat is not enough (or too low) to let the model remove (or install) completely a different technology. Also electricity marginal prices are not particularly affected, but the trend of their variation is rather interesting as shown by Figure 63 and Figure 64.

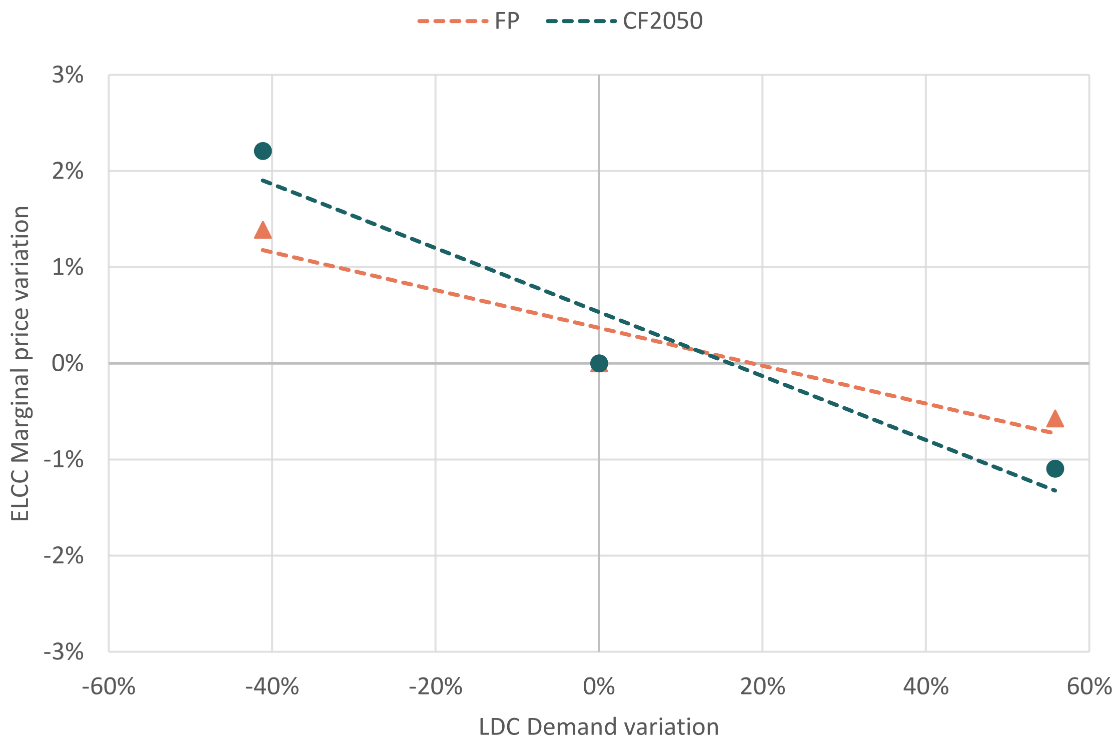


**Figure 63 –** Sensitivity analysis on electricity marginal price in 2050

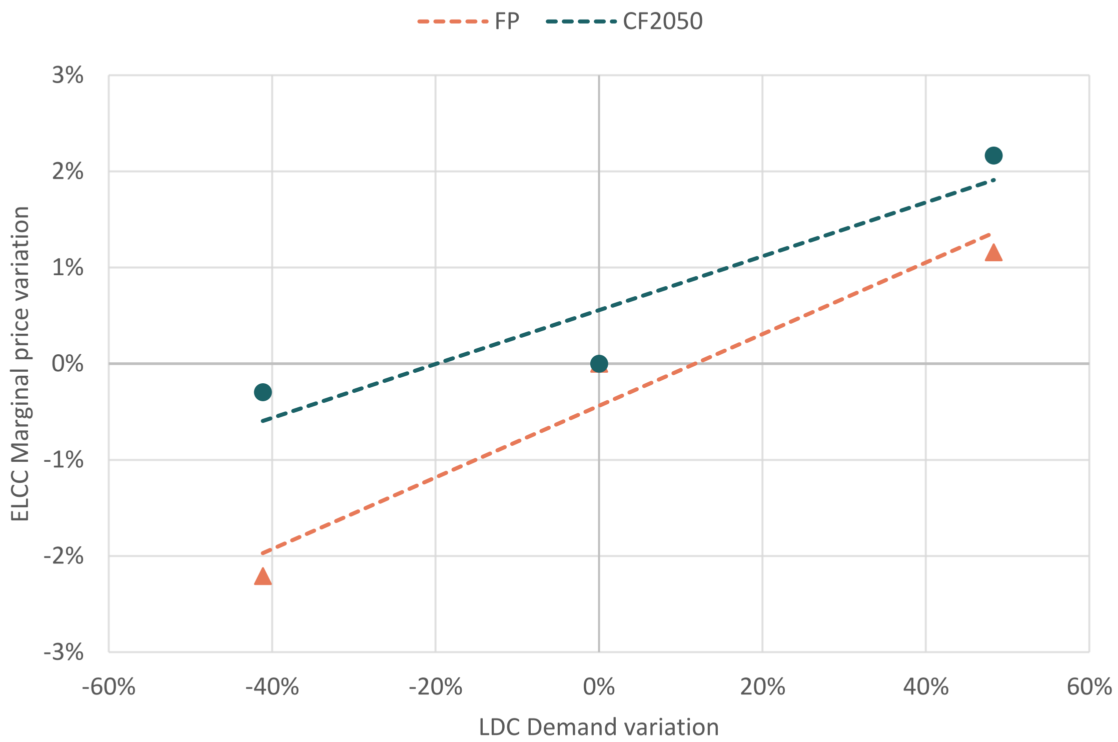


**Figure 64 –** Sensitivity analysis on electricity marginal price in the period 2035-2050

In 2050 electricity marginal prices are inversely proportional to the additional power demand caused by large data centers. This is due to the evolution of the system power capacity in the previous years: results show that a higher additional power demand induces the model to build new capacity at a pace higher than the dismission rate of the old capacity. Consequently, when higher additional power demands are inserted in the system, there is more capacity to sustain it without increasing excessively marginal prices at the end of the time horizon. In fact, when looking at the average electricity marginal price in the period 2035-2050 it increases with higher additional power demands because of larger investments in new capacity.

## 6.5 Total quantities throughout the model horizon

This section is intended to give a quick idea of the impact that large data centers may have on the Danish energy system analyzing three extremely synthetical quantities, obtained as a sum over the entire model horizon:

- Total new installed capacity;
- Cumulated $CO_2$ emissions;
- System cost.

Results regarding the Frozen Policy and the Target scenarios with the presence of large data centers are also endowed with the uncertainties due to the cases “low LDC demand” and “high LDC demand”. The former is represented by the lower end of the error bar, while the latter is represented by its higher end.

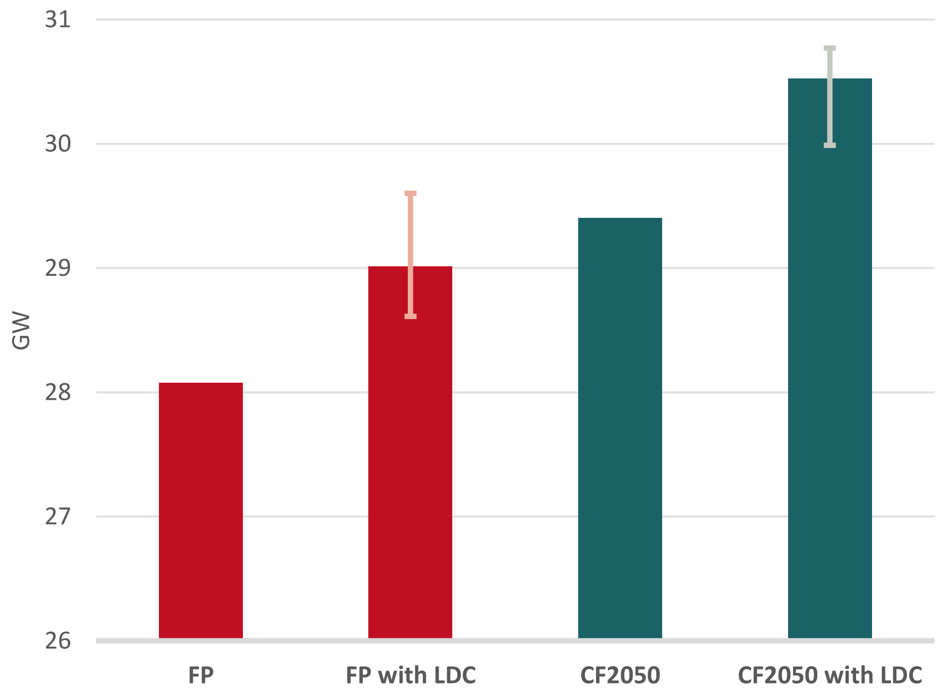


**Figure 65 –** Cumulated new power capacity in the period 2020-2050 with uncertainties

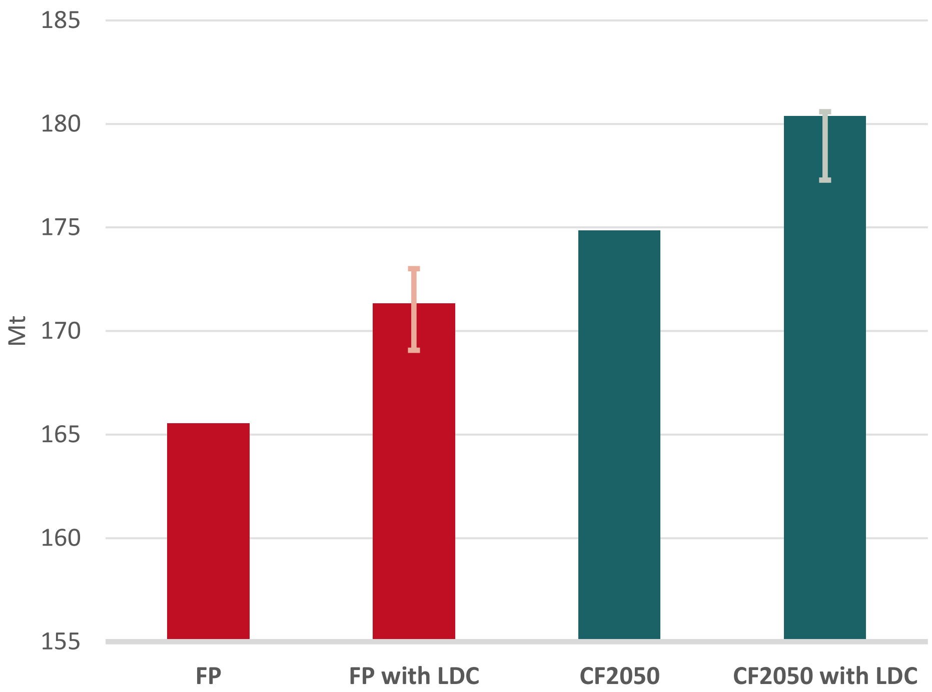


**Figure 66 –** Cumulated $CO_2$ emissions in the period 2020-2050 with uncertainties

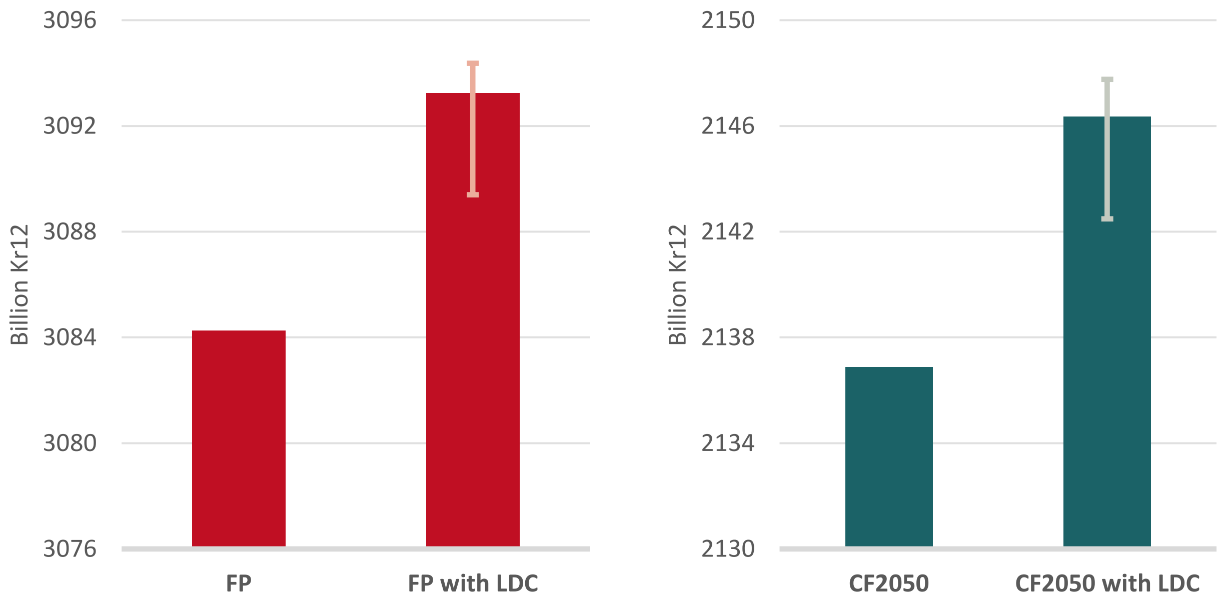


**Figure 67 –** Total system cost

# Chapter 7

## Further developments

Modelling large data centers in the framework of long-term energy system analyses represents a novel approach. As such, it bears limitations and further progresses can be achieved. In addition, TIMES-DK itself has room for improvements, particularly regarding its temporal and spatial resolution [19].

Developments involving large data centers can be identified in the following aspects. First, a better knowledge of their power consumption should be possessed at least at a daily temporal resolution. In my opinion, this aspect could be addressed in two ways: on the one hand, it could be enough to know the nominal power and a real workload profile of specific large data centers within the energy system under study (i.e. improving the method used in this thesis). On the other hand, it could be also possible to retrieve their power consumption modelling the electric behavior of large data centers' elementary units (servers, UPS, PDUs, etc.). The latter approach has a relevant advantage: implementing better efficiencies and different costs for future equipment, which can then be chosen by the model (similarly to the solution adopted for residential appliances in TIMES-DK [19]). However, this requires estimates about the number of installed servers and about their operational configurations, information which is hardly shared by the companies that possess large data centers. A partial way out consists in abandoning the model of the facilities and considering the total number of servers (and other equipment) installed in the country, as suggested by [9], [10]. However, this further solution relies on market surveys, which are expensive, and their future projections have only a 5-year time horizon.

Then, further developments are strictly correlated to the temporal and spatial definition of the energy system model. If a finer temporal scale can be adopted it may be convenient improving the modelling of the cooling system and the consequent waste heat recovery, to grasp more precisely hourly variations that may influence the model's choices regarding capacity installations. Moreover, large data centers have a well-defined geographical position, thus a finer spatial scale allows evaluating more accurately the feasibility of waste heat recovery. Notice also that a further improvement about this aspect could be reached in the model presented in this thesis by assessing exactly if the considered large data centers are within district heating areas or not, thus implementing costs for additional piping if necessary.

Finally, it could be interesting to add the possibility of electricity storage. In fact, large data centers may give flexibility to the entire power system at an hourly level. In addition, their location seems large enough to host several storage technologies. Therefore, a tool such as TIMES-DK may evaluate whether this solution is convenient from a systemic point of view.

# Chapter 8

## Conclusions

The rising power demand required by new hyperscale data centers is posing some issues in the countries that are physically hosting them. Among these countries, Denmark is probably the only one to possess an extensive model for its entire energy system (TIMES-DK), which aims at supporting policy makers through long-term scenario analyses. This thesis exploited this unique opportunity to create a model for large data centers in the context of a broader energy system model. Consequently, it was possible to evaluate the implications of their estimated power consumption from a systemic perspective.

Firstly, a major effort was made in understating their behavior from an energy point of view. It was discovered that variations in their power consumption are only due to their workload profile and the cooling system design. In addition, they produce an elevated amount of low-temperature waste heat, which can be recovered. Therefore, this brought to the creation of a thermodynamic model for a small sample data hall, whose primary focus is estimating the daily variations in the power consumed by the cooling system (added to the power consumption of the IT equipment) and the quantity of available waste heat.

Secondly, outcomes from the thermodynamic model (having an hourly temporal resolution) were aggregated in time and space for the implementation in TIMES-DK, which divides a year in 32 time-slices and has to represent the whole bunch of large data centers instead of a single data hall. In addition, the model was given the possibility to choose whether to recover their waste heat or not.

Successively, results were analyzed by means of two benchmark scenarios: Frozen Policy and Target scenario. Both of them agree that that the power system is able to satisfy the additional power demand thanks to 0.9-1.2 GW of new wind offshore plants in the long term. However, the transition towards 2050 is not straightforward: wind offshore plants are still expensive for further immediate installation and the same applies to wind onshore and PV, which in addition face capacity bounds due to spatial constraints. Consequently, in the first years the model suggests an extended use of fossil-fueled CHP plants, which cause a peak of $CO_2$ emissions in the FP scenario in 2025 (+9%) and a general increase of the cumulated emissions through the entire time horizon in both scenarios (+2%). Moreover, in the

Target scenario the presence of large data centers modifies the power demand of two other sectors: transports and industry. The former is particularly affected and faces a slower electrification both in freight and passengers transport. Nevertheless, it should be noticed that these results are valid under the hypothesis of net zero annual electricity exports, which was added to avoid that uncertainties in import and export prices influenced excessively the outcomes.

Then, results also showed that a waste heat reuse for district heating (after a quality upgrade) may not be always a workable solution from a systemic perspective until 2045. In fact, only the Frozen Policy scenario suggests a complete recovery, while in the Target scenario it is recommended only in the last period of the time horizon. However, due to their well-defined location, a more specific assessment should be made based on the possibilities offered by the hosting municipalities.

Finally, sensitivity analyses (obtained varying the annual power consumption of large data centers) confirmed the tendencies described in the results.

# References


[1] "2050 Energy strategy - European Commission." [Online]. Available: https://ec.europa.eu/energy/en/topics/energy-strategy-and-energy-union/2050-energy-strategy. [Accessed: 23-Feb-2018].

[2] "Independent from Fossil Fuels by 2050 -The official website of Denmark." [Online]. Available: http://denmark.dk/en/green-living/strategies-and-policies/independent-from-fossil-fuels-by-2050. [Accessed: 23-Feb-2018].

[3] V. Depoorter, E. Oró, and J. Salom, "The location as an energy efficiency and renewable energy supply measure for data centres in Europe," *Appl. Energy*, vol. 140, pp. 338–349, 2015.

[4] O. Gozcu, "Worldwide Energy Analysis of Major Free Cooling Methods for Data Centers."

[5] Cisco, "Cisco Global Cloud Index : Forecast and Methodology , 2014–2019," *White Pap.*, pp. 1–41, 2014.

[6] "The Copenhagen Post," *http://www.cphpost.dk/*.

[7] J. G. Koomey, "Estimating total power consumption by servers in the US and the world," *CIRP Ann. - Manuf. Technol.*, vol. 53, no. 1, pp. 151–154, 2007.

[8] J. G. Koomey, "Worldwide electricity used in data centers," *Environ. Res. Lett.*, vol. 3, no. 3, p. 34008, 2008.

[9] W. L. Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, Jonathan Koomey, Eric Masanet, Nathaniel Horner, Inês Azevedo, "United States Data Center Energy Usage Report," no. September, 2016.

[10] L. Stobbe, M. Proske, H. Zedel, R. Hintemann, J. Clausen, and S. Beucker, "Entwicklung des IKT - bedingten Strombedarfs in Deutschland – Kurzdarstellung – Die Prognose im Überblick," no. 29, 2015.

[11] S. W. Ham, M. H. Kim, B. N. Choi, and J. W. Jeong, "Energy saving potential of various air-side

economizers in a modular data center," *Appl. Energy*, vol. 138, pp. 258–275, 2015.

[12] S. W. Ham, M. H. Kim, B. N. Choi, and J. W. Jeong, "Simplified server model to simulate data center cooling energy consumption," *Energy Build.*, vol. 86, pp. 328–339, 2015.

[13] S. Pelley, D. Meisner, T. F. Wenisch, and J. W. VanGilder, "Understanding and Abstracting Total Data Center Power," *WEED Work. Energy Effic. Des.*, 2009.

[14] P. R. Wang, M. Scharling, and K. P. Nielsen, "Teknisk Rapport 12-17: 2001 – 2010 Design Reference Year for Denmark," pp. 1–31, 2012.

[15] Energinet, "REPORT 2017 Energinet's analysis assumptions," no. 16, 2017.

[16] N. Van Beeck, "Classification of Energy Models," *Tilbg. Univ. Eindhoven Univ. Technol.*, no. May, p. 25, 1999.

[17] R. Loulou, "Documentation for the TIMES Model - Part I," no. July, pp. 1–151, 2016.

[18] J. Tattini, "Modeling of the Norwegian power system and analysis of the power trade in the Nordic countries," no. March, 2015.

[19] O. Balyk *et al.*, "TIMES-DK: technology-rich multi-sectoral optimisation model of the Danish energy system," *Energy Strateg. Rev. [under Rev.*, pp. 1–18.

[20] S. N. Petrović and K. B. Karlsson, "Residential heat pumps in the future Danish energy system," *Energy*, vol. 114, pp. 787–797, 2016.

[21] K. B. Karlsson, S. N. Petrović, and R. Næraa, "Heat supply planning for the ecological housing community Munksøgård," *Energy*, vol. 115, pp. 1733–1747, 2016.

[22] J. Tattini, M. Gargiulo, and K. Karlsson, "Reaching carbon neutral transport sector in Denmark – Evidence from the incorporation of modal shift into the TIMES energy system modeling framework," *Energy Policy*, vol. 113, no. February 2017, pp. 571–583, 2018.

[23] L. Barroso and U. Hölzle, "The Datacenter as a Computer," *Synth. Lect. Comput. Archit.*, vol. 24, p. 156, 2013.

[24] N. L. Shrestha *et al.*, "RenewIT – Deliverable D4.5 – Catalogue of advanced technical concepts

for Net Zero Energy Data Centres," no. 608679, 2015.

[25] M. Dayarathna, Y. Wen, and R. Fan, "Data Center Energy Consumption Modeling: A Survey," *IEEE Commun. Surv. Tutorials*, vol. 18, no. 1, pp. 732–794, 2016.

[26] "Our Footprint | Facebook." [Online]. Available: https://sustainability.fb.com/our-footprint/. [Accessed: 01-Dec-2017].

[27] Google Inc., "Google Environmental Report," pp. 1–68, 2016.

[28] K. Ebrahimi, G. F. Jones, and A. S. Fleischer, "A review of data center cooling technology, operating conditions and the corresponding low-grade waste heat recovery opportunities," *Renew. Sustain. Energy Rev.*, vol. 31, 2014.

[29] Viborg Kommune, "VVM-tilladelse til datacenter og ny højspændingsstation ved Foulum," p. 6, 2016.

[30] "Odense Data Center - Home | Facebook." [Online]. Available: https://it-it.facebook.com/OdenseDataCenter/. [Accessed: 22-Feb-2018].

[31] Danish Energy Agency, "Technology Data for Energy Plants Updated chapters, August 2016," no. August, p. 183, 2016.

[32] M. Avgerinou, P. Bertoldi, and L. Castellazzi, "Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency," *Energies*, vol. 10, no. 10, p. 1470, 2017.

[33] "Colocation Denmark - Data Centers." [Online]. Available: http://www.datacentermap.com/denmark/. [Accessed: 22-Feb-2018].

[34] "Statistics Denmark." [Online]. Available: https://www.dst.dk/en. [Accessed: 22-Feb-2018].

[35] K. Schultz and Advisor, "Global competition in the datacenter industry – why Denmark has a strong competitive position."

[36] R. A. Steinbrecher and R. Schmidt, "Data center environments ASHRAE's evolving thermal guidelines," *ASHRAE J.*, vol. 53, no. 12, pp. 42–49, 2011.

[37] S. Dockweiler, Personal communication, January 2018.