

A Multi-Agent Framework for Smart Grid Simulations: Strategies for power-to-heat flexibility management in residential context

Original

A Multi-Agent Framework for Smart Grid Simulations: Strategies for power-to-heat flexibility management in residential context / Rando Mazzarino, Pietro; Macii, Alberto; Bottaccioli, Lorenzo; Patti, Edoardo. - In: SUSTAINABLE ENERGY, GRIDS AND NETWORKS. - ISSN 2352-4677. - ELETTRONICO. - 34:(2023). [10.1016/j.segan.2023.101072]

Availability:

This version is available at: 11583/2978675 since: 2023-05-22T10:47:09Z

Publisher:

Elsevier

Published

DOI:10.1016/j.segan.2023.101072

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

A Multi-Agent Framework for Smart Grid Simulations: strategies for power-to-heat flexibility management in residential context

Pietro Rando Mazzarino^a, Alberto Macii^a, Lorenzo Bottaccioli^b, Edoardo Patti^a

^a*Dept. of Control and Computer Engineering, Politecnico di Torino, corso Duca degli
Abruzzi 24, Turin, Italy*

^b*Interuniversity Dept. of Regional and Urban Studies and Planning, Politecnico di
Torino, corso Duca degli Abruzzi 24, Turin, Italy*

Abstract

The complexity of a system, such as the Smart Grid, comes from the interaction of different fields of expertise and actors cooperating together. As a result, co-simulation frameworks are emerging to meet the need for analysis from a broader perspective. The feature to better understand the Smart Grid environment is the ability to simulate its different components simultaneously. Therefore, we propose a multi-agent co-simulation framework that can serve as a test-bed for multiple Smart Grid strategies. In particular, we focus on Demand Response programs that exploit the thermal behavior of residential buildings. The proposed framework is modular to test Plug&Play models. Moreover, it is highly flexible and configurable to evaluate realistic scenarios. We tested the platform on a case study of 1,000 buildings, performing an analysis of the effects of micro temperature deviations in buildings on the primary grid substation balancing problem. The results show the flexibility of the platform in testing different strategies. Moreover, power imbalances could be mitigated simply by acting on the indoor temperature set-points with smaller deviations.

Keywords:

Demand Response, Co-simulation, Multi Agent Systems, Power Imbalances, Flexibility, Thermal Inertia

1. Introduction

Over the past decade, the distribution grid has seen an increase in small-scale producers of renewable energy sources (RES), which affects the accuracy of power demand forecasts. Inaccurate forecasting of distributed RES generation or small-scale residential sector demand could result in scheduled demand and generation that differ from those measured, causing power imbalances [1]. As a possible solution, flexibility resulting from demand response (DR) programs can be used. DR includes the modification of electricity consumption patterns by customers in response to different electricity prices over time, incentive payments, or when system reliability is jeopardized [2].

Residential DR programs can take advantage of the thermal flexibility of buildings. In fact, thermal loads account for the largest share of building consumption, providing ample potential flexibility. However, various elements, including the physical characteristics of the building and the thermal comfort of the user, can reduce this potential. To study DR strategies and for any kind of Smart Grid analysis, tools capable of coupling simulators and models covering different real-world aspects are needed.

Co-simulation tools play a key role in evaluating the effects of a given strategy in the broader perspective of the whole system. The study of DR strategies with stand-alone simulations may be lacking in understanding the consequences on more complex systems. This is the case, for example, with studies on exploiting building flexibility in the context of Smart Grids, among which only a few consider city energy supply systems [3]. In particular, technical modeling of the power grid is often overlooked, thus neglecting the effects of proposed DR strategies on the feasibility of resource reallocation and power exchanges. Thus, a horizontal perspective on the multitude of subsystems that make up the underlying environment on which particular strategies are tested is needed. Indeed, the Smart Grid context could be viewed as a Multi-Agent System (MAS) where, while testing the behavior of a specific physical agent, we need a good characterization of the external environment to understand its response. Moreover, the importance of co-simulation tools is directly related to their role as “common ground” for testing. Several proposed DR studies cannot be easily compared due to substantial differences in test conditions. Therefore, the only way to assess consistent comparison and benchmarking relies on tools that can serve as a test-bed. In addition, the flexibility of these tools is essential for simulating

different contexts and scenarios.

Another important aspect is the scalability of the frameworks; by analyzing existing urban energy modeling systems (UEMS), it is evident that this quality is closely related to the level of detail of the models chosen. Aoun et al. [4] present a nice representation of the relationship between building models and scalability. More detailed models usually allow for the simulation of a smaller number of elements, while simplified models, such as resistance-capacity (RC), allow for the simulation of a larger number of entities. Therefore, an appropriate goal for a co-simulation environment should be to allow and facilitate the evaluation of mixed levels of detailed perspectives.

On these premises, we propose a multi-agent co-simulation framework to evaluate flexibility strategies in Smart Grids. The key design points of the framework are i) Flexibility - i.e. agents and models do not conform to a fixed typology, but can be fully configured; ii) Plug&Play - i.e. the ability to easily attach and detach simulators and models; iii) Scalability - i.e. the ability to simulate a large set of entities while maintaining a good level of modeling detail; and iv) Modularity - i.e. the ability to easily exchange modules such as DR strategies. In particular, with this paper, we present an extension of our previous work [5] in which we tested two DR strategies exploiting building thermal flexibility while integrating a power grid model, a model for the building thermal behaviour and the HVAC subsystems models (extended from City Energy Analyst (CEA) [6]). In this work, we better describe the platform itself and stress its modularity and flexibility by integrating a new DR strategy based on a much more computational intensive approach - i.e. a Genetic Algorithm (GA). The novelty of this work is mainly represented by the platform of which potential is demonstrated by the proposed analyses. The possibility of easily swapping DR strategies allows the comparison between the DR strategies proposed in [5] and the new one based on GA, showing the usefulness of this tool.

For the proposed simulations we modelled three main types of agents: i) the DSO (Distribution System Operator), in charge of controlling the electrical grid; ii) the Aggregator, responsible for the DR management, thanks to the thermal flexibility volumes gathered from buildings; and iii) the Building Agent, encapsulating the thermal models, the HVAC models and the control system.

The rest of the paper is organised as follows. Section 2 reviews the solutions presented in literature providing an overview of existing tools. Section 3 presents the structure of the framework, introducing the agents and their be-

haviours. Section 4 illustrates the chosen scenarios, while Section 5 discusses our experimental results, comparing the adopted DR strategies. Finally, Section 7 draws the conclusions.

2. Related Works

In literature, as shown by Kathirgamanathan et al. [7], there is a wide range of solutions for managing and quantifying the energy flexibility of buildings. It is stated that energy flexibility could be addressed through several different perspectives but, in particular, this review focuses on control and management strategies (e.g., rule-based solutions and Model Predictive Control (MPC)). Several solutions for control and management exist and this breadth of methods requires a fair comparison with the same parameters of judgment. This is only the tip of the iceberg when analyzing the entire Smart Grid context in terms of co-simulation, management and other solutions. Therefore, the main shortcoming of the smart grid literature is the absence of an overall perspective when simulating or analyzing specific strategies. In particular, we report some papers related to the fields addressed by our work: thermal modeling of buildings, flexibility estimation, and co-simulation frameworks. We want to show that each of these major papers lacks some aspects that could be easily taken into account by exploiting our framework.

Authors of [8] and of [9] have presented how to exploit houses thermal mass to generate flexibility and analysed the effects on the indoor comfort of simple control strategies of the heating system. They have shown the effects on the grid but only from an economic point of view ignoring the physical constraints and the possible drawbacks. Authors of [10] proposed an economic MPC for assessing building flexibility potential, they included storage and solar panels, and the building thermal model is a black-box. This approach neglects the physical parameters of the building and is very tailored upon the study case, thus making it difficult to generalise and reuse in broader contexts. On the other hand, Zhang et al. [11] tested two control strategies, i.e. MPC and rule-based, relying upon a TRNSYS calibrated building model, which, being very detailed, does not scales-up or ensure building heterogeneity. In [12], building heterogeneity is taken into account and the flexibility potential is analysed over more than one building. However, this work only quantifies the effect of set points changes and no advanced control strategies are proposed. Other interesting models are presented in [13] and [14] but

they still focus on management techniques and on a single specific building without considering building heterogeneity and scalability.

Alongside specific building models and control strategies, the co-simulation approach has gained attention due to its capability to couple and reuse several domain-specific models developed with various tools. According to this, a second set of solutions have been analysed. In [15], the authors introduced a co-simulation framework that uses PandaPower [16] and TRNSYS to simulate the electric grid and the buildings. The model used for buildings represents a single-family Spanish house typology. Four scenarios with a 3 minutes time-step and an increasing penetration rate of heat pumps to determine the impact on the grid are analysed. Instead, Molitor et al. [17] present the MultiEnergy System COSimulator for City District Energy Systems, i.e. a framework that allows performing simulations of district scale energy systems. In [17], Neplan is used to simulate the electrical network, while a control algorithm sets the operation schedule of heating systems at fixed time intervals. Even in this case, a single family building model has been developed. On the basis of this, four different heating systems have been addressed. To scale-up and fasten simulations, parallelisation of processes have been used. The largest scenario includes 795 building energy systems. However, both [15] and [17] tend to generalise a single building model not allowing heterogeneous representation of the environment.

Wang et al. [18] designed an urban energy co-simulation framework based on standard Functional Mock-up Interface and CityGML semantic 3D city model. It analyses two different scenarios simulating one and six buildings coupling Nottingham Multi Agent Stochastic Simulation and EnergyPlus. This framework has been cited because it well represents the modularity and flexibility on the scenario creation and building models presented in our framework but it neglects the power grid integration.

In this work, we propose a Multi-Agent Framework for co-simulation in urban environments that will serve as a test-bed for DR strategies. It takes into consideration power grid and building modelling. In particular, we offer the possibility to model different types of buildings exploiting archetypes. In particular, the thermal model, the HVAC typology and its component, the occupancy patterns and the appliances can be chosen from a static database or manually configured thanks to the parameterized models. In addition, thanks to the framework performances and the reduced order models is possible to easily scale up the number of buildings and different sizes and typologies of power grid can be modelled. The framework is characterised

Table 1: Comparison among the proposed Framework and solutions presented in literature

Authors	Grid	Building heterogeneity	Number of buildings	Agent-based	Co-simulation	Advanced Control Strategy
Finck et al. [10]	✗	✗	1	✗	✗	ANN-MPC
Zhang et al. [11]	✗	✗	1	✗	✓	Hybrid Generalized Pattern Search Algorithm with Particle Swarm Optimization
Yin et al. [12]	✗	✓	large-scale group of different customers	✗	✗	✗
Le Dréau et al. [8]	economic	✗	1	✗	✗	✗
Masy et al. [9]	economic	✗	1	✗	✗	optimal predictive control
Wang et al. [18]	✗	✓	6	✓	✓	✗
Tardif et al. [15]	✓	✗	50	✗	✓	✗
Molitor et al. [17]	✓	✗	795	potentially	✓	potentially
Our previous work [5]	✓	✓	1000	✓	✓	✗
Our Framework	✓	✓	1000	✓	✓	configurable (e.g Genetic Algorithm)

by its scalability, modularity and high configurability. Therefore the objective is to show the easiness of testing different DR strategies over the same heterogeneous scenario. Thus, in this extension of our previous work [5], we introduce a more complex flexibility estimation method based on GA, which acts as a control strategy for the individual building. Table 1 lists and compares previous work to illustrate the differences and to highlight the novelties proposed by our framework. Looking at columns of Table 1, we have highlighted some of the features addressed or not by the analysed studies. The majority of analysed papers do not take into consideration the power grid and its modelling except for two cases in which the grid is considered only from an economic point of view [8][9]. Building heterogeneity and the number of simulated buildings are important aspects for a proper analysis of district energy performances, several studies focus on small set of units and with a reduced characterisation of them. Finally, for what concerns the involved ICT technologies, only one study, [18], coupled agent-based and co-simulation techniques but did not test any advanced control strategy. In conclusion, it can be said that our framework enhances simulation capabilities by combining the advantages of co-simulation and agent-based concepts. This integration facilitated the incorporation of multiple physical and operational aspects into a single simulation environment, that are typically not interconnected in current literature solutions.

3. Methodology

This section presents the methodology followed to address the design and implementation of our multi-agent framework. The main purpose of this work is the design and implementation of a Plug&Play multi-agent co-simulation

environment that will serve as a test-bed for different management strategies in urban distribution-level power grids. The overall system architecture in Figure 1 consists of three main layers: i) the *Data-source layer*, to manage the input data-structures and the archetypes; ii) the *Co-simulation layer*, the core of our framework; iii) the *Application layer*, to allow end-users interactions with the co-simulation process. This last layer gives access to the databases allowing customisation, scenario selection through a GIS-based interface and visualisation of data and results. In the following sections, a more in depth description of both Data-source and Co-simulation layers is given.

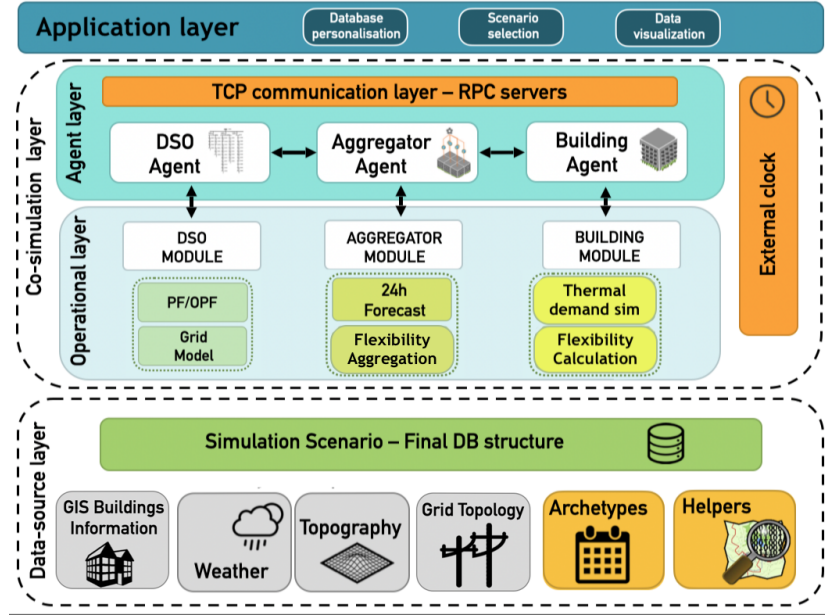


Figure 1: Schema of the proposed framework.

3.1. Data-source layer

The Data-source layer performs data-management and processing to obtain a comprehensive description of the urban scenario under analysis. As shown in Figure 1, grey blocks represent raw inputs, in particular: i) *GIS Buildings Information* (e.g. geometries, ages of construction, height and typology of use); ii) *Weather* information for the chosen simulation period and the geographic area (e.g. outdoor temperature, radiation, etc.); iii) *Topography* information about the elevation of the terrain; iv) *Grid topology* of

the power grid. In case the user do not provide some of these data, *Helpers* modules allows to retrieve some of these information from external open-source databases (e.g. OpenStreetMap for retrieving the case study zone information).

Once collected this raw information, a preliminary characterisation of the case study is already provided. Then a further step is performed in order to enhance the data description: CEA’s archetypes [6](see *Archetypes* block in Figure 1) are exploited in order to correlate building inputs with more complex information. Archetypes are ideal examples of building systems that are used to assume the unknown characteristics of the building under analysis from known information. Therefore, specific properties of the envelope are associated with each building by the year of construction, and RC models are parameterized for each specific case. The HVAC system and related sub-systems (distribution and generation/transformation) are assigned to each building based on the age of construction, but the user can easily assign the preferred technology from those in the static databases that compose the archetypes data structure. This allows different solutions to be tested and possible retrofits of older buildings to be considered. In addition, knowledge of the building’s type of use allows occupancy patterns and programmed temperature set-points to be associated.

Once completed the correlation of buildings and CEA’s archetypes, the scenario is fully characterised and each information is accessible and ready for customisation. Furthermore, a solar radiation calculation is performed in order to retrieve data about the external gains on the buildings. The complete simulation scenario is saved in a final database structure containing all the information: *Simulation Scenario - Final DB structure* in Figure 1.

3.2. Co-simulation layer

The core layer of the proposed framework is the *Co-simulation layer*. It is composed by two main sub-layers, the *Agent layer* and the *Operational layer*. To gain modularity, the Agent and Operational layers have been kept separate, maintaining the difference between the simulator wrappers (the agents) and the physical or operational models to be instantiated (the operational modules). We intend the agents as wrappers because they have a fixed structure that act as an interface to the simulated environment for the models or simulators that they encapsulate. In fact, as shown in Figure 1, the operational capabilities of each agent are implemented in the Operational

layer in the form of Plug&Play modules, while the communication capabilities are implemented in the Agent layer. This architecture resembles the idea behind co-simulation frameworks, in which there are several simulators instantiating different model instances. In fact, we designed a co-simulation environment by exploiting the concepts of the MAS world, the agents become the simulators that wrap the models and communicate with each other. The main difference is that in our case the structure is fully distributed, without the need for a centralized orchestrator, and through proper implementation the agents can wrap any type of software actor. Building a co-simulation environment using the MAS architecture allows for easy deployment and parallelization, as well as paving the way for more complex data exchange workflows, enabling the integration of time- and event-based simulators without the cumbersome design of an orchestrator.

Section 3.2.1 explains the types of agents, the communication infrastructure, and their role in the co-simulation environment, while Section 3.2.2 briefly explains the built-in models.

3.2.1. Agent layer

The agent layer is the fixed backbone of the platform and consists of agents with a standardized input/output interface that can communicate with each other and be generated at will. By exploiting the AIOMAS [19] python library, based on the asynchronous programming library ASYNCIO [20], the communication infrastructure has been implemented as a distributed architecture (see *TCP communication layer - RPC servers* block in Figure 1). Agents “live” in different containers that run as separate processes. This allows the framework to be parallelized and run on different computers and/or servers. To ensure that messages are exchanged between them, communication links are implemented on the TCP/IP protocol stack thanks to the Remote Procedure Calls (RPC) features built into the containers. Each container has its own RPC server and a specific TCP address in order to act as a gateway for its member agents. Time synchronization is through external clocks that regularly update the container clocks, which automatically set the agents’ schedules. The agents, however, are operational model wrappers (which can also be external simulators); their independence allows high configurability both on the choice of modules to be integrated and on the interactions to be taken between them. In our implementation, as shown in Figure 1, we identified three main agents and their specific roles:

(i) ***DSO Agent*** starts the workflow (see Section 3.3) and it is in charge

of balancing and voltage acceptability of the power grid. It interacts with Aggregators by asking them the flexibility reserve. Its main responsibility is to manage the power unbalances at the primary substation while maintaining stable the portion of grid within its competence. To accomplish these tasks, it models the grid and performs an optimal calculation for the resources dispatch by means of its block modules in the operational layer (see Section 3.2.2).

(ii) **Aggregator Agents** are the service providers allowing the implementation of DR strategies of aggregate management of resources. They interact with both *Building Agents* and the *DSO Agent*, acting as intermediaries. Aggregators forecast the load demand for the buildings they coordinate and communicate it to the DSO. Later, when the DSO needs flexibility to balance the network, it is asked to estimate the reserve amount and report it. This is done by virtually aggregating a number of buildings. The number of aggregators, as well as the number or which buildings they coordinate, is fully configurable, allowing the platform to test different levels of aggregation and the possible effects on the chosen strategy.

(iii) **Buildings Agents** represent the actual buildings. They are grid connected, meaning that their actual power demand is directly sensed by the grid. Due to the proposed configuration this is translated into a direct communication between them and the DSO (which in this framework takes both the role of the utility and the power grid). Beside the modeling of the building physical behaviour they also instantiate a module for estimating and controlling real-time flexibility (see Section. 3.2.2).

3.2.2. Operational layer

In the previous sections, it was mentioned that agents serve as a wrapper for physical and operational models. In fact, each agent instantiates specific models to simulate some specific features. This section discusses the main modules tested, but one of the main innovations of this work is its modularity. Thus, each of the models in the operational layer can be easily interchanged with others, and even in the same simulation, each agent can instantiate a different model.

Starting from the models used by the *DSO Agent* we have the *Grid Model* block and the *PF/OPF* block as in Figure 1. They both exploits the PANDAPOWER [16] python library. The former allows the representation of the power grid and all its components (e.g., lines, buses, transformers, loads, generators, storage), while the latter allows the calculation of Power Flow (PF)

or Optimal Power Flow (OPF). In particular, the network model can read a specific case file, allowing dynamic changes to it. For example, loads and generators variables can be updated also at runtime and the infrastructure can be modified, making the power grid fully configurable. Dynamic updating of loads, generators and storage is used at each simulation time-step to update building loads with an associated minimum/maximum limit to express flexibility. Instead, focusing on the PF/OPF block, the power flow is calculated to estimate the losses in the system and thus to obtain a real value of the power drawn at the primary substation. This is necessary because the DSO decides to balance or not to balance the grid only on the basis of the power detected at the primary substation, so if the difference between the programmed and actual power is greater than a configurable threshold the DSO activates the balancing strategy. The balancing strategy takes advantage of the flexibility offered by aggregators in terms of more or less load power per building. In order to optimally redistribute resources, an optimal power flow is calculated taking into account the flexible loads communicated by the aggregators. It is worth noting that in the following formulation the exploitation of flexibility with respect to grid withdrawn is prioritized thanks to the realistic assumption of much lower prices. The used OPF formulation and solving method are taken from PANDAPOWER [16]. This model take into consideration: i) AC load flow equations; ii) branch constraints (maximum line loading); iii) bus constraint (maximum and minimum voltages and angles requirements); iv) transformer constraints (maximum loading); v) operational constraints, for each load, generator and storage a maximum and minimum for active and reactive power can be specified (this is used to express the flexibility rather than operational limits); vi) costs, the cost function to be minimized is expressed as a piece-wise linear function for all the element in the grid. For a more in depth discussion on the formulation and related technicalities we refer to [16] [21] [22] [23]

The aggregator modules we have implemented in this work are rather simple. The first is the *24h forecast* block, which returns the day-ahead energy consumption to the DSO. Since the actual forecast is outside the scope of this work, this module is simply represented by noise added to the actual consumption profile estimated from the aggregation of building loads. Each aggregator estimates this value for its portion of the buildings. The main idea of this module is to distribute the prediction task to groups of buildings in order to reduce the prediction error that can result from aggregating too many different buildings. This has not been addressed in this work, but the

modularity of the framework allows easy integration of real and sophisticated forecasting tools without affecting the rest of the framework. The other block instantiated by the aggregator is the *Flexibility Aggregation*, again the model simply acts as an aggregator of information from buildings, but it could easily be replaced by a resource redistribution or balancing module for the aggregator subgroup. In fact, in this work we want to evaluate the effect of individual buildings flexibility (chosen by the building management system) on the whole system. Some strategies from the aggregator’s point of view could be added in the future (e.g. peer to peer markets in case of energy communities).

Finally, the models instantiated by the Building Agent represent one of the focal points of this work, in addition to the framework itself. These are the blocks *Thermal demand sim* and *Flexibility calculation*, as in Figure 1. The former simulates the physical behavior of the building. It includes several subsystems and calculations. The model we present is an improvement of the CEA [6] dynamic demand forecasting tool, which has been modified to perform demand calculations at each time-step separately, for the desired time period, and with the desired time resolution. In fact, the original tool was able to estimate building loads with a one-shot calculation by exploiting time series data, with our modifications it can be included in a step-wise simulation and control actions on the simulated systems are possible thanks to the implementation of inverse methods for each calculation step. First, the thermal behavior is simulated by means of an RC model conforming to ISO-13790, each building is treated as a single thermal zone. This model is depicted in Figure 2 and is composed of 6 resistances and 1 capacitance as in the European and Swiss standard SIA 2044, it is used to calculate the sensible heat taking into consideration the storage effect of the thermal mass and the indoor air, while modelling the temperature exchange among occupancy, outdoor air and solar radiation. The calculated sensible heat needed by the thermal zone to reach specified temperature set-point is then used to model the heat fluxes of the subsystem involved: emission, distribution and generation/transfer systems. For these subsystems several technologies are present in the CEA static databases and all of them are modelled in order to calculate losses and auxiliary power incurred. The chosen technologies for our case study are presented in Section 4.

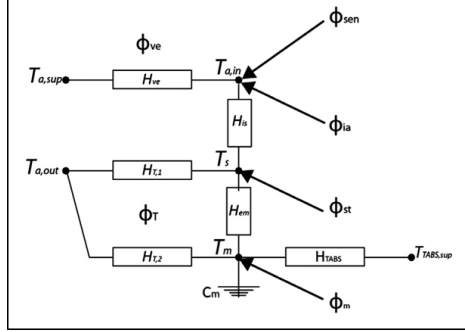


Figure 2: The used RC model [6]. Where $T_{a,in}$ is the indoor air temperature, T_s is the mean between indoor air temperature and mean radiant temperature, T_m is the equivalent thermal mass temperature, $T_{a,out}$ is the external temperature, $T_{TAB,sup}$ is the supply temperature from slab embedded HVAC system (when existing) and $T_{a,sup}$ is the supply temperature from ventilation. Then, ϕ_{ia} , ϕ_{st} , ϕ_m represent internal gains and solar heat sources, ϕ_m is the heat flux from the thermal mass of the building, ϕ_{ven} the ventilation heat flux, ϕ_T the heat flux from external environment and finally ϕ_{sen} is the sensible heat to be calculated. Furthermore, we have the thermal mass capacity and the resistances representing all the heat transfer delays.

3.3. Co-simulation workflow

The analysis conducted in this paper, using the proposed platform, mainly focuses on the study of a DR strategy by exploiting the energy flexibility of buildings. In this work, flexibility is generated by the thermal behavior and heating loads of residential buildings. The DSO exploits this flexibility to optimally reallocate resources to meet imbalances between scheduled power profiles in the day-ahead and measured power demand. Figure 3 shows the interactions that occur during the simulation process. The workflow can be divided into day-ahead and intra-day operations.

The Day-ahead operations are the following. i) The *DSO Agent* collects the day-ahead power schedule from the *Aggregators Agent*. ii) The DSO performs power flow calculations computing the value of the active power at the primary substation for the next 24-hours.

Instead, the interactions taking places at each time-step during real-time operations are the following. i) The *Building Agents* calculate their load from heating system and the *DSO Agent* retrieves them. ii) The *DSO Agent*, exploiting this information, performs power flow calculations and compares the obtained active power value at primary substation with the scheduled one. The *DSO Agent* compares the deviation between these two values with a de-

fined threshold. If this deviation falls within the power threshold, the *DSO Agent* is able to cover all the unbalance by itself, meaning that it no longer exploits flexibility from users. Instead, if the power threshold is exceeded, a message asking for flexibility is sent to the *Aggregator Agents*. iii) *Aggregator Agents* forward this message to each *Building Agent* in their premises, which perform the calculation of the flexibility reserve, estimating how much flexibility they are willing to offer for that time-step. This information is sent back to the *Aggregator Agents* and then to the *DSO Agent*. iv) The *DSO Agent* computes the optimal power flow with the new information, changing the resource allocation to obtain the desired adjustment. v) The iteration is over when results of the OPF calculation are spread back to the *Aggregator Agents*, disaggregate the information and tell each *Building Agents* how much to modulate their HVAC systems.

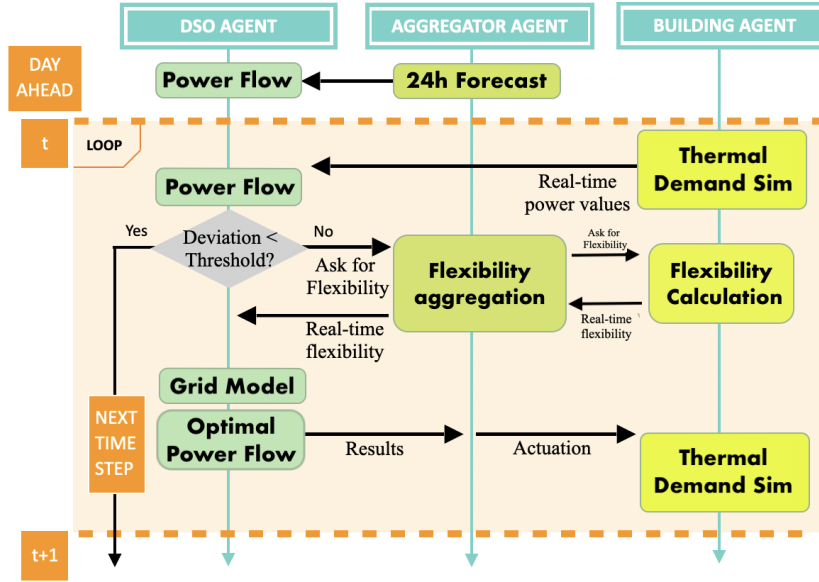


Figure 3: Interactions taking place during a simulation time-step

4. Experimental Set-up

Our tests and analyses have been carried out defining an urban scenario composed of 1000 buildings in a city district of Turin (Italy) as shown in Figure 4(b). The majority of buildings are multi-residential (apartment buildings) with some exceptions of mixed usages (e.g. 80% residential, 20% com-

mercial), this is used to take into consideration different occupancy schedules. Each building was modeled as a single thermal zone, and the envelope differs according to the specific age of construction, retrieved in the scenario creation phase. In addition, the buildings are equipped with a heating system consisting of a ground/water heat pump, a distribution system, and radiators as emission terminals. The heat pumps are ideally modeled without regard to operational constraints, so the control is done as modulation of the electrical power used for conversion. The distribution system is sized with respect to the size of the building, so that a reasonable amount of leakage is considered. Radiators are sized with respect to the size of the thermal zone. These are the technologies considered, however, the flexibility of the framework allows for testing various other solutions (e.g., different generations and different emissions). Also, in this simulation we used the same technologies for each building (appropriately sized for the specific buildings), but heterogeneous building technologies could be considered. The period chosen for the simulation is January in a typical weather year in the North of Italy, thus considering only the heating season. We have used a Medium Voltage power grid consisting of a primary substation with three bus-bars operating at a nominal voltage of 22 kV, which on their turn supply 51 substations equipped with Medium Voltage/Low Voltage (22 kV/400 V) transformers as shown in Figure 4(a). The buildings are evenly distributed, to be supplied, on the 51 substations.

The hierarchy configuration of the agents consists of an individual DSO and 11 aggregators clustering several buildings as shown in Table 2.

Table 2: Case study hierarchy configuration

AggregatorID	# of Buildings	# of Buses	AggregatorID	# of Buildings	# of Buses
AGG_0	71	5	AGG_6	103	5
AGG_1	91	5	AGG_7	89	5
AGG_2	109	5	AGG_8	91	5
AGG_3	100	5	AGG_9	96	5
AGG_4	99	5	AGG_10	55	3
AGG_5	96	5			

On top of this simulated environment, several control strategies and analysis could be carried out. In order to test the flexibility and modularity of our platform, we decided to test three different solutions. The possibility of switching from one strategy to another allows to easily compare and evalu-

for the three scenarios is the temperature tolerance, which is used as a constraint for the internal temperature. The values chosen for the tolerance are in accordance with the ASHRAE [24] standard on indoor temperature fluctuations to avoid user discomfort. Two different time-step resolutions, namely 1 hour and 15 minutes, were evaluated to emphasize the limitations of the RC formulation. It is worth to be noted that time resolution is completely configurable, but of course it must be compliant with the physical significance of the involved models. The framework was tested on a single server machine, spawning agents across its cores, ensuring both parallelization and concurrency. The proposed scenarios are better explained in the following sections.

4.1. Fixed deviation Scenario

The fixed deviation scenario is the more idealised, in which all the buildings always participate in the DR program. In particular, every time it is requested by the aggregators, buildings will always offer the maximum possible flexibility while staying within the given range of temperature tolerance. The reserve is expressed as positive and negative deviation from the needed power values to achieve scheduled temperature set-points in each time-step. Then, the DSO will exploit the communicated reserve completely or not depending on the performed optimisation calculations. In this work, the temperature tolerance was assigned the same for each building, but the platform allows different values to be tested for each individual building, more details on this strategy are given in [5].

4.2. Random deviation Scenario

The random deviation scenario is the most unpredictable one, in which buildings adhere to the DR program completely randomly. Specifically, whenever requested by aggregators, buildings randomly choose the amount of flexibility reserve they are willing to offer. The choice is made by randomly choosing an amount of deviation from a normal distribution bounded by “no deviation” and a maximum value that coincides with the power associated with the temperature tolerance. Also, in this case, the reserve is a positive and negative deviation to the scheduled power values, thus the building response is heterogeneous contemplating the different possible involvement of the buildings, more details on this strategy are given in [5].

4.3. Genetic deviation scenario

As a main novelty with respect to [5], we introduced the genetic deviation scenario described below. This scenario has been chosen to prove that the presented framework could support more intelligent management strategies at the building level. In contrast with the previous, it works on a wider time horizon (e.g. 3 hours) and performs optimisation in order to assess the maximum building flexibility for a specific time period. Building Agents must solve, regularly, a scheduling problem for a configured time horizon. The objective function in Equation (1) maximises the value of flexibility over the whole time horizon and it is soft-constrained by the indoor comfort of users thanks to a penalty factor. The flexibility is expressed as a deviation between power consumption (P_t) and the pre-scheduled power profile ($P_{0,t}$), which is the needed load for achieving the indoor temperature set-points (T_{set}) in buildings. Instead, the penalty factor is calculated by multiplying the absolute value of power deviation ($|P_{0,t} - P_t|$) times a configurable constant (c), when the temperature constraints (T_{tol}) are violated. Thus, the penalty increases with the magnitude of occurred violation and if constraints are respected it is equal to zero. In a nutshell, the problem consists of finding the best indoor temperature trend over the time horizon that allows the building to offer the maximum flexibility to the grid without affecting the users comfort too much. The values of power and temperature are strictly related to the models of the HVAC and the thermal behaviour of the building. Therefore the sub-optimal solution of the problem is composed of the indoor temperature (T_{res}) schedule and the corresponding power profile. The latter is communicated to the Aggregator in form of boundaries of the flexibility reserve (e.g. min values of the power reduction). It is worth noting that the previous strategies only focused on the specific time-slots without having knowledge of the behaviour on a longer period, therefore the thermal inertia is contemplated only as a passive effect and not used as an active flexibility contribution.

$$\max \sum_{t=1}^{t+hor} (P_{0,t} - P_t) - penalty \quad (1)$$

$$penalty = \begin{cases} 0 & \text{if } T_{res} \geq T_{set} - T_{tol} \\ 0 & \text{if } T_{res} \leq T_{set} + T_{tol} \\ c(|P_{0,t} - P_t|) & \text{elsewhere} \end{cases}$$

Algorithm 1 Pseudo-code for the Genetic algorithm with used hyper-parameters values

```

1: procedure GENETIC ALGORITHM( $T_{set}$ ,  $n_{hor}$ ,  $f(P)$ ,  $N_{it}$ ,  $S_{pop,init}$ ,  $\alpha$ ,  $\beta$ ,
    $\gamma$ ,  $X_{type}$ )
2:    $\triangleright T_{set}$  is a vector of temperature set-points with dimension  $(1 \times n_{hor})$ 
3:    $\triangleright n_{hor}$  is the horizon: number of evaluated time-steps = 3hours
4:    $\triangleright f(P)$  is the objective function to optimise = Equation 1
5:    $\triangleright N_{it}$  is the maximum number of iterations = 100
6:    $\triangleright S_{pop,init}$  is the initial size of population = 30
7:    $\triangleright \alpha$  is the mutation probability = 0.2
8:    $\triangleright \beta$  is the parents portion = 0.3
9:    $\triangleright \gamma$  is crossover probability = 0.5
10:   $\triangleright X_{type}$  is crossover typology = uniform
11:   $\triangleright$  Initialization
12:   $S_{pop,init} \triangleright$  Random choice of initial population, set of  $T_{set}$  vectors
13:   $n = 0 \triangleright$  iteration number
14:   $\triangleright$  Execution
15:  while convergence not reached or  $n < N_{it}$  do
16:    Evaluation of fitness  $\triangleright$  run the demand module for each  $T_{set,i}$ 
17:    Selection of best performing
18:    Crossover  $\triangleright X_{type}, \gamma$ 
19:    Mutation  $\triangleright \alpha$ 
20:    Composition of new population  $\triangleright \beta$ 
21:     $n \leftarrow n + 1 \triangleright$  Update of iteration number
22:  end while
23:  return  $T_{set,opt}$  and  $P_{min/max,t,opt} \triangleright$  Optimized solution is returned
24: end procedure

```

The pseudo-code in Algorithm 1 explains the procedural steps of the GA algorithm used. To estimate the building flexibility reserve through the optimal scheduling of indoor temperature set-points, we used a heuristic method


that allows us to encapsulate the entire building HVAC model. Therefore, the search is performed from a pool of solutions of indoor temperature set-points programmed for a given horizon, then, for each iteration of the genetic algorithm, the entire thermal demand simulation model is computed to evaluate for each possible set-point solution the respective power value. In this way, the suitability of the solutions can be evaluated on the objective function (Equation 1), that depends on the power values, even though the decision variables are the temperature set-points. The pseudo-code of the Algorithm 1 also shows the hyper-parameters used by the GA.

5. Experimental Results


In this section, we present the experimental results of all the performed simulations. The balancing dynamic is shown in Figure 5 which reports results of the Fixed deviation scenario using two different time resolutions. Both Figure 5a and Figure 5b depict a snapshot of 24 hours of the power profiles at the primary substation. The difference between 15 min time resolution and 1h is tangible and consists of the finer trends obtainable using smaller time-steps. The presented profiles are: i) the day-ahead scheduled trend (light-blue line), ii) the predefined power threshold (light-grey area), iii) the real-time values collected during every time-step (red dashed line) and iv) the adjusted curve (green dashed line) - i.e. the new power profile taking into account flexibility contribution.

The scheduled trend represents the forecast for the next 24 hours of what will be the withdrawn power at primary substation. Instead, at each time-step the measured value of power demand, i.e. real-time trend, is retrieved. If the power is outside the threshold range (which in our case is $\pm 5\%$) the DSO activates a balancing strategy. The result of applying the DR strategy is represented by the adjusted trend, which is the attempt to bring back the power withdrawn into the range of tolerance and to overlap as much as possible the new power profile with the Day-ahead Scheduled trend. Figure 5a and Figure 5b highlight the effect of the different time-resolutions corresponding to different regulation frequencies. Thus, we can see that more frequent control actions (e.g. scenario with a 15 min resolution) result in more scattered behaviours and a finer description of the trends.

Table 3 presents the results for each of the performed simulation. For both time resolutions, 1 hour and 15 minutes time-steps, respectively, we have performed three simulations per scenario carried out using three dif-



(a) 1 hour resolution



(b) 15 minutes resolution

Figure 5: Snapshot of the *Fixed* deviation scenario with tolerance of 0.5°C

ferent ranges of temperature tolerance, i.e. $\pm 0.5^{\circ}\text{C}$, $\pm 1.0^{\circ}\text{C}$, $\pm 2.0^{\circ}\text{C}$, to analyse the impact of the indoor comfort constraint on the flexibility reserve. The metrics proposed to compare results are : i) *% of success*, i.e. the total number of times in which the adjustment was successful to bring back the power withdrawn into the threshold range; ii) *% covered by flex*, which represents, in average, how much power unbalance has been fulfilled exploiting power from the flexibility reserve during the simulation period; iii) *Root Mean Square Error* (RMSE) in MW of the adjusted trend converging on the scheduled trend (e.g. green and blue lines in Figure 5, respectively); iv) *T deviation*, representing the mean deviation in $^{\circ}\text{C}$ from temperature set-points registered during the whole simulation by all buildings.


Table 3: Results of all the simulations

Time-step	Scenario	Tolerance [°C]	% of success	% covered by flex	RMSE [MW]	T deviation [°C]
1 hour	Fixed deviation	± 0.5	68.95	88.67	0.273	0.42
		± 1	96.64	99.78	0.050	0.46
		± 2	88.02	99.82	0.273	0.14
	Random deviation	± 0.5	25.54	94.51	0.457	0.25
		± 1	64.78	88.04	0.275	0.42
		± 2	91.67	99.68	0.102	0.47
	Genetic deviation	± 0.5	31.25	89.91	0.503	0.20
		± 1	72.32	90.70	0.424	0.34
		± 2	88.52	99.55	0.179	0.42
15 min	Fixed deviation	± 0.5	69.90	90.16	0.201	0.39
		± 1	89.38	99.95	0.079	0.36
		± 2	85.23	98.40	0.260	0.10
	Random deviation	± 0.5	26.69	97.56	0.374	0.23
		± 1	69.72	90.17	0.206	0.40
		± 2	87.82	99.94	0.127	0.35
	Genetic deviation	± 0.5	25.00	94.27	0.447	0.20
		± 1	53.64	88.27	0.316	0.33
		± 2	86.97	99.11	0.198	0.46


By observing column *% of success* in Table 3, we can notice that for both genetic and Random deviation scenarios this value increases when the temperature tolerance range increases. This effect is completely foreseeable since an increase of the temperature tolerance deviation results in an increase of the offered flexibility (i.e. the higher the temperature deviation is the higher the possible power reduction will be). An interesting observation comes out for the Fixed deviation scenario in which the direct correlation between tolerance range and power flexibility no longer seems to be true. In fact, the *% of success* for the Fixed deviation case from ± 0.5 to ± 1 °C increases and in contrast it decreases when using ± 2 °C of tolerance. This phenomenon does not invalidate the assumption of proportionality between temperature tolerance ranges and the amount of flexibility. Thus, in this case, having a larger potential flexibility allows some buildings to highly deviate their indoor temperature from the set-points. This, at certain time-steps, causes the over-exploiting of the buildings flexibility reserve. When this happens for a certain period and for all the buildings in the scenario, the shared over-exploitation will make buildings running out of reserve for the next time-steps, in order to avoid constraints violations. Besides the performances of success, we can see that the values of *% covered by flex* are all above the 88%. This means that even if not always, it is possible to lead

back the power profile into the threshold range.


In general, the majority of the unbalances have been covered thanks to the DR approach and the reallocation of resources. From the very small values of $RMSE$, it is possible to understand that during the simulated month (January of a typical year in the North of Italy) the adjusted trend has mostly been coinciding with the scheduled one.



(a) *Fixed* deviation scenario



(b) *Random* deviation scenario



(c) *Genetic* deviation scenario

Figure 6: Snapshot of the power profiles at primary substation of the three scenarios with a temperature tolerance of $\pm 1^\circ\text{C}$ and 15 min of time-resolution

Figure 6 shows a comparison of power balancing at primary substation

for the three scenarios in the very same conditions: $\pm 1^\circ\text{C}$ of temperature tolerance and 15 min time-resolution.


Qualitative considerations on the differences among scenarios could be drawn by looking at these graphs. First of all, we can observe that a tolerance of $\pm 1^\circ\text{C}$ in the Fixed scenario (see Figure 6a) is in general enough to succeed and completely balance the grid. From the graph, we can actually notice that the adjusted trend (i.e. green dashed line) coincides with the forecast trend (i.e. light blue line) almost for the whole day. This does not happen for both Random deviation scenario (see Figure 6b) and Genetic deviation scenario (see Figure 6c), because they are much more “individualistic”. Individualistic means that, in these scenarios buildings do not always offer the maximum of flexibility associated with the temperature deviation as, on the contrary, it happens in the Fixed deviation scenario. In the Random deviation scenario buildings participate randomly and with random reserves.

In the Genetic deviation scenario when and how to participate is chosen by each building individually, this occurs even if the optimisation objective is to offer flexibility to the grid. For this reason, both Figure 6b and Figure 6c highlight that, even if the balancing strategies behave correctly by leading back the power profile, the $\pm 1^\circ\text{C}$ temperature tolerance does not unlock enough flexibility when systems behave without coordination among final users.


Finally, the Genetic deviation scenario is the more conservative in terms of indoor comfort and as a consequence when compared with the others from the grid balancing side of the problem perform differently having a different set of priorities. This comparison shows the effectiveness of the platform in enabling the swapping of different strategies even with different objectives.

Only for the Fixed deviation scenario with $\pm 2^\circ\text{C}$ of temperature tolerance, the values of *T deviation* in Table 3 seem to disprove the correlation among temperature tolerance and actual deviation from set-points. Indeed, in this case, we have smaller values of deviation even if the temperature tolerance is larger. This demonstrates a peculiarity of the Fixed deviation scenario with larger temperature tolerances than $\pm 1^\circ\text{C}$. Generally, taking a single building both the genetic and the Random deviation strategies will be more conservative in terms of indoor comfort with respect to the Fixed deviation scenario during the whole simulation period. But, looking at the whole system, if all actors have large flexibility reserves (e.g. Fixed deviation scenario), the DSO will be able to completely balance the grid by performing a fairer resource redistribution among the buildings. Thus, in the case in


which there is already enough flexibility for the balance, the Fixed deviation strategy allows fulfilling the common goal by sacrificing less. In general, making an exception for the above mentioned Fixed deviation scenario, the T deviation values increase with the increase of the temperature tolerance. This comes directly with the fact that the temperature tolerance act as a boundary on how much flexibility the building is willing to offer.



(a) *Fixed* deviation scenario



(b) *Random* deviation scenario



(c) *Genetic* deviation scenario

Figure 7: Snapshot of the power profiles at primary substation of the three scenarios with temperature tolerance of 2°C and 15 min of time resolution

Comparing Figure 7 with Figure 6, it is possible to notice how much balancing results get better when temperature tolerance is raised by one degree. In particular, even the Genetic deviation scenario successfully brings back the power profile into the threshold (see Figure 7c). On the other side both Fixed deviation and Random deviation scenarios fulfil an almost perfect regulation with a tolerance of $\pm 2^{\circ}\text{C}$ (see Figure 7a and Figure 7b).

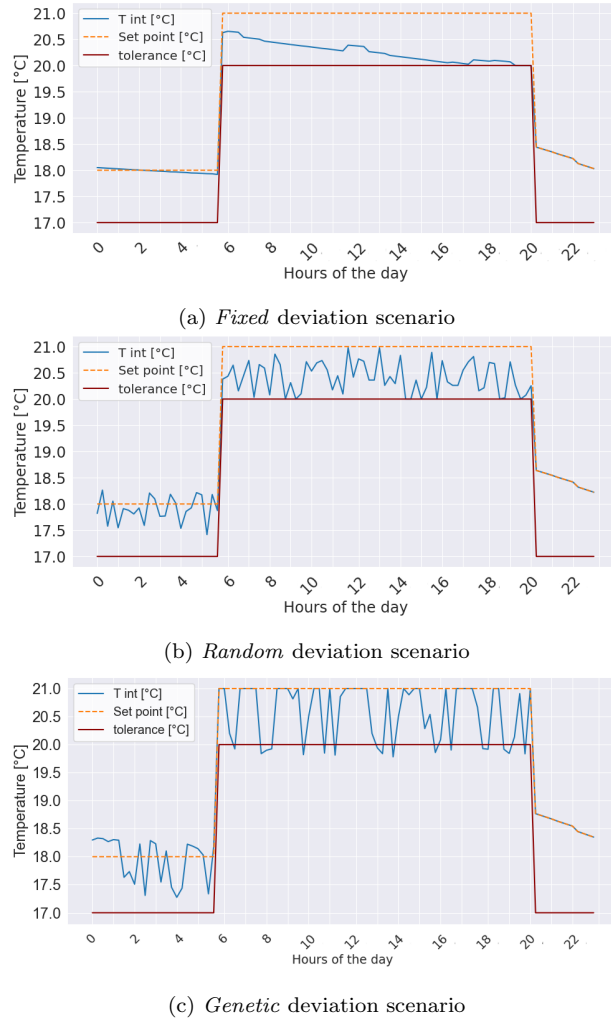


Figure 8: Indoor temperature of a single building for the three scenarios with 1°C tolerance and 15 min of time resolution

Figure 8 shows the effects of the three scenarios on the indoor temperature

profile of an individual building. The chosen conditions are $\pm 1^\circ\text{C}$ tolerance and 15 min time-resolution. As a first consideration, we can observe the smoother variation of temperature for the Fixed deviation scenarios (see Figure 8a) with respect to the more scattered behaviours obtained in both Random and Genetic deviation scenarios (see Figure 8b and Figure 8c).

In the Fixed deviation scenario this happens because buildings are always trying to offer maximum flexibility. In particular, looking between 5:00 am and 9:00 pm in Figure 8a, we can notice that the indoor temperature is slowly declining to the tolerance limit. This happens because, as previously mentioned, in the Fixed deviation scenario the DSO does not need to always exploit the whole deviation offered by a single building but can draw smaller amount of flexibility from all the buildings. Instead, Figure 9 shows the behaviour of the Fixed deviation scenario with $\pm 0.5^\circ\text{C}$ of tolerance in which the indoor temperature always reached the tolerance. Hence, meaning the complete exploitation of offered flexibility by the building.




Figure 9: Indoor temperature of a building in Fixed deviation scenario with 0.5°C tolerance and 15 min of time resolution

Figure 8c highlights that before any relevant flexibility exploitation the building rises its indoor temperature to give the maximum flexibility right after. This is the result of the temporal optimisation that is performed in the Genetic deviation scenario. In addition in the Genetic deviation scenario, the temperature tolerance is used as a soft constraint by looking at the small violations that occurred in Figure 8c. The scattered behaviour of both Random and Genetic deviation scenarios could seem unacceptable fluctuations

of the indoor temperature, but they actually do not affect the indoor comfort level being in a very short range of degrees.

6. Preliminary assessment on CPU timings

All the strategies presented in the previous section 5 have been tested in a server environment equipped with an Intel Xeon Processor with a frequency of 2.4 GHz, 32 cores and 128 Gb of RAM. To provide a preliminary assessment of the flexibility of our framework of distributing computational resources, we repeated our tests first by exploiting all the 32 cores and the reducing the number of used cores to 5. This gives us a preliminary analysis on the CPU timings, as shown in Table 4. Moreover, to evaluate the scalability of our solution in dealing with a variable number of buildings we run the tests on the three scenarios with 55 and 1000 buildings respectively. A first consideration should be made about the results associated with the scenarios integrating the genetic algorithm, the low time performance of the *Genetic deviation* scenarios, shown in Table 4, depends on two factors: (i) the genetic algorithm module has not been optimized; (ii) the distribution of entities executing the genetic algorithm along the cores has been done by agent type, which means that multiple aggregators performing the most computationally intensive tasks are operated on the same core, creating bottlenecks. Both of these problems can be easily solved, the former with an optimized module, which was out of the scope of this work, and the latter by simply isolating the most expensive tasks on dedicated cores or servers. In addition, looking at the performance of the other two scenarios, namely *Fixed deviation* and *Random deviation*, we can see that the timing performance is well over 1 minute per time-step, ensuring real-time behaviour for most of the operational and physical phenomena involved in urban energy systems. Studying the computational complexity of a distributed framework can be very challenging, and we aim to perform it in much more detail in future works. However, a first look at the capability of the framework is well represented by the two simplest scenarios. In fact, the temporal performance and computational complexity of these types of tools are strictly dependent on the integrated modules and the distribution of computational resources. Therefore, by looking at the simplest scenarios, in which the integrated modules do not create bottlenecks, we can get an idea of the scalability performance of the framework. Taking the *Fixed deviation* scenario with 55 buildings as an example, the average execution time is 1.86 seconds, while increasing the number of

	scenarios	avg time-step duration [s]	avg time-step duration [s]	avg overall duration [s]	avg overall duration [s]
		32 cores	5 core	32 cores	5 core
55 buildings	Fixed deviation	1.86	2.10	44.64	50.40
	Random deviation	1.94	2.13	46.56	51.12
	Genetic deviation	40.62	142.00	974.88	3042.00
1000 buildings	Fixed deviation	23.26	25.37	558.24	609.11
	Random deviation	22.20	24.74	532.8	593.76
	Genetic deviation	1067.21	4787.34	25613.20	114869.32

Table 4: Timings of simulations compared with a previous case with only 55 buildings

entities to 1000, where the average time is 23.26 seconds, showing a nonlinear increase. Looking at the effect of scaling the simulations on the server cores, it is possible to understand that the benefits introduced by the parallelization increase with the complexity of the involved modules. Indeed, when scaling from 5 to 32 cores the scenario that gain the most is the *Genetic deviation*, the others, being simple, do not need that much computational nodes and thus does not gain that much when increasing the number of cores.

7. Conclusion

The main gap addressed in this work is the lack in the literature of simulation frameworks that bring together the ability to test flexibility strategies on a realistic and highly configurable environment, giving the possibility to analyze the system response from different perspectives. We designed a MAS co-simulation framework in which the thermal model of buildings, physical model of the power grid, HVAC models, agent-based communication infrastructure, and demand response strategies were integrated. All of this, ensuring the features that allow for maximum flexibility of the tool in terms of use: modularity, Plug&Play, and scalability. The design and implementation followed the concept that the smart grid can be viewed as a multi-agent system, thus enabling complete infrastructure distribution and decomposition of the overall problem. The idea is to leverage MAS as a co-simulation platform, which is tailored to the specific field and requires a little more effort in implementation, but allows an entire system to be represented without the need for a coordinator, as is the case in the real world. The proposed platform has been able to simulate a thousand of buildings in a city district with different time-resolutions. We have integrated the main actors of the electric power exchange at the distribution level, by creating a parametrized scenario based on real data and building or technological archetypes. The way the platform creates urban scenarios is highly configurable allowing the

user to choose the desired level of detail. The multi-agent arrangement and the modularity of the platform embrace the possibility of extensions in multiple directions. Thus, external modules may be added, new agents may be integrated and different scenarios and strategies may be tested.

The second problem addressed is the studying of different power flexibility strategies and the possibility to compare them. Therefore, a primary need is to have an effective test-bed for testing the solution in the same environment. As an initial step, we focused our attention on electrically activated thermal loads. Thus we tested the platform for simulating three different strategies acting on heating loads of building premises. The strategies range from the simplest ones to more complex solutions and have been compared. It has been possible to analyse macro behaviours at the grid level (e.g. the balancing effects at primary substation) as well as internal micro behaviours at the building level (e.g temperature profiles). Results from the presented case study demonstrate that, even in absence of additional flexibility sources (e.g. storage), good balancing results could be obtained by simply applying minor indoor temperature deviations in the buildings. In conclusion, we presented a test-bed that can be used for co-simulation or agent-based simulations to evaluate different flexibility strategies on a common, configurable physical environment. Beyond the scenarios tested, which nonetheless demonstrate the feasibility of simple flexibility strategies, the main objective was to show the framework's potential to scale up to 1,000 buildings while maintaining a granular physical description of the technologies involved, such as heat pumps, building envelopes, home distribution and emission systems, and the power grid with attached components.

Acknowledgements

This work is part of the project NODES which has received funding from the MUR-M4C2 1.5 of PNRR with grant agreement no. ECS00000036

References

- [1] A. Estebasari, E. Patti, L. Bottaccioli, Real-time control of power exchange at primary substations: An opf-based solution, in: in Proc. of 2020 IEEE EEEIC / I CPS Europe, 2020, pp. 1–6. doi:10.1109/EEEIC/ICPSEurope49358.2020.9160765.

- [2] P. Siano, Demand response and smart grids—a survey, *Renew. Sust. Energ. Rev.* 30 (2014) 461–478. doi:<https://doi.org/10.1016/j.rser.2013.10.022>. URL <https://www.sciencedirect.com/science/article/pii/S1364032113007211>
- [3] A. Sola, C. Corchero, J. Salom, M. Sanmarti, Multi-domain urban-scale energy modelling tools: A review, *Sustainable Cities and Society* 54 (2020) 101872. doi:<https://doi.org/10.1016/j.scs.2019.101872>.
- [4] N. Aoun, R. Bavière, M. Vallée, A. Brun, G. Sandou, Dynamic simulation of residential buildings supporting the development of flexible control in district heating systems, in: in *Proc. of 13th International Modelica Conference*, Regensburg, Germany, March 4–6, 2019, no. 157, Linköping University Electronic Press, 2019.
- [5] P. Rando Mazzarino, C. De Vizia, E. Macii, E. Patti, L. Bottaccioli, An agent-based framework for smart grid balancing exploiting thermal flexibility of residential buildings, in: *2021 IEEE EEEIC/I&CPS Europe*, IEEE, 2021, pp. 1–6.
- [6] J. A. Fonseca, A. Schlueter, Integrated model for characterization of spatiotemporal building energy consumption patterns in neighborhoods and city districts, *Appl. Energy* 142 (2015) 247–265. doi:<https://doi.org/10.1016/j.apenergy.2014.12.068>.
- [7] A. Kathirgamanathan, M. De Rosa, E. Mangina, D. P. Finn, Data-driven predictive control for unlocking building energy flexibility: A review, *Renew. Sust. Energ. Rev.* 135 (2021) 110120.
- [8] J. Le Dréau, P. Heiselberg, Energy flexibility of residential buildings using short term heat storage in the thermal mass, *Energy* 111 (2016) 991–1002.
- [9] G. Masy, E. Georges, C. Verhelst, V. Lemort, P. André, Smart grid energy flexible buildings through the use of heat pumps and building thermal mass as energy storage in the belgian context, *Science and Technology for the Built Environment* 21 (6) (2015) 800–811.
- [10] C. Finck, R. Li, W. Zeiler, Economic model predictive control for demand flexibility of a residential building, *Energy* 176 (2019) 365–379.

- [11] K. Zhang, M. Kummert, Evaluating the impact of thermostat control strategies on the energy flexibility of residential buildings for space heating, in: in Proc. of Building Simulation, Springer, 2021, pp. 1–14.
- [12] R. Yin, E. Kara, Y. Li, N. DeForest, K. Wang, T. Yong, M. Stadler, Quantifying flexibility of commercial and residential loads for demand response using setpoint changes, *Appl. Energy* 177 (2016) 149–164.
- [13] H. Golmohamadi, K. G. Larsen, P. G. Jensen, I. R. Hasrat, Optimization of power-to-heat flexibility for residential buildings in response to day-ahead electricity price, *Energy and Buildings* 232 (2021) 110665.
- [14] G. Bianchini, M. Casini, A. Vicino, D. Zarrilli, Demand-response in building heating systems: A model predictive control approach, *Appl. Energy* 168 (2016) 159–170.
- [15] J. Tardif, F. DiazGonzález, M. Kegel, A. Sola, J. Salom, A. de Besòs, A co-simulation framework for assessing the interaction between heat pumps and the low voltage grid on a district scale, *IBPSA* (2020).
- [16] L. Thurner, A. Scheidler, F. Schäfer, J. Menke, J. Dollichon, F. Meier, S. Meinecke, M. Braun, pandapower— an open-source python tool for convenient modeling, analysis, and optimization of electric power systems, *IEEE Trans. Power Syst.* 33 (6) (2018) 6510–6521. doi:10.1109/TPWRS.2018.2829021.
- [17] C. Molitor, S. Groß, J. Zeitz, A. Monti, Mescos—a multienergy system cosimulator for city district energy systems, *IEEE Trans. Ind. Inform.* 10 (4) (2014) 2247–2256. doi:10.1109/TII.2014.2334058.
- [18] K. Wang, P.-O. Siebers, D. Robinson, Towards generalized co-simulation of urban energy systems, *Procedia Eng.* 198 (2017) 366–374.
- [19] Aiomas, <https://aiomas.readthedocs.io/en/latest/>, accessed:2019-11-11.
- [20] ASYNCIO, Asyncio documentation, accessed: 2020-01-02. URL <https://docs.python.org/3/library/asyncio.html>
- [21] H. Wang, On the computation and application of multi-period security-constrained optimal power flow for real-time electricity market operations (2007).

- [22] H. Wang, C. E. Murillo-Sanchez, R. D. Zimmerman, R. J. Thomas, On computational issues of market-based optimal power flow, *IEEE Transactions on Power Systems* 22 (3) (2007) 1185–1193.
- [23] R. D. Zimmerman, C. E. Murillo-Sánchez, R. J. Thomas, Matpower: Steady-state operations, planning, and analysis tools for power systems research and education, *IEEE Transactions on power systems* 26 (1) (2010) 12–19.
- [24] ASHRAE, Standards and guidelines, accessed: 2020-02-24.
URL <https://www.ashrae.org/technical-resources/standards-and-guidelines>