



HAL
open science

A model-based approach to addressing energy demand in sustainable urban systems

Abdoulaye Gamatié, Thomas Leduc, Daniel Siret, Gilles Sassatelli, Michel
Robert

► **To cite this version:**

Abdoulaye Gamatié, Thomas Leduc, Daniel Siret, Gilles Sassatelli, Michel Robert. A model-based approach to addressing energy demand in sustainable urban systems. *Sustainable Computing: Informatics and Systems*, 2023, 37, pp.100844. 10.1016/j.suscom.2022.100844 . hal-03944160

HAL Id: hal-03944160

<https://hal.science/hal-03944160>

Submitted on 10 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NoDerivatives 4.0 International License

Highlights

A Model-Based Approach to Addressing Energy Demand in Sustainable Urban Systems

Abdoulaye Gamatié, Thomas Leduc, Daniel Siret, Gilles Sassatelli, Michel Robert

- A new simulation tool for studying urban deployments of the so-called GENESIS (generation of energy-synergistic information systems) technology, such as evaluating the impact of IT equipment usage on local energy consumption. GENESIS relies on distributed modules (or green nodes), each may combine a server, a renewable energy harvester, local energy storage, and custom programmable power electronics systems that handle energy transfer between interconnected modules. The simulation tool's typical output is a map of an urban area with estimates of energy production and consumption over time, taking into account the weather conditions in that area.
- A case study of potential GENESIS deployment scenarios in two different French cities: Nantes and Montpellier. Both cities differ in terms of urban layout, buildings typology and weather conditions. We investigate the conditions under which, from an energy standpoint, such deployments can favor sustainability in relation to typical urban activities in each city district we studied.
- A thorough discussion on the integration of the proposed innovative technology within sustainable urban systems w.r.t. some relevant perspectives.

A Model-Based Approach to Addressing Energy Demand in Sustainable Urban Systems[★]

Abdoulaye Gamatié^{a,*}, Thomas Leduc^b, Daniel Siret^b, Gilles Sassatelli^a and Michel Robert^a

^aLIRMM, Université de Montpellier & CNRS, 191 Rue Ada, Montpellier, 34095, France

^bNantes Université, ENSA Nantes, École Centrale Nantes, CNRS, AAU-CRENAU, UMR 1563, F-44000 Nantes, France

ARTICLE INFO

Keywords:

green edge computing
renewable energy
digital services
urban environment modeling
simulation tool

ABSTRACT

Sustainable cities have an integrated approach to addressing their citizens' environmental, social, economic, cultural, and institutional concerns. For a successful transformation of urban systems in this context, it is important to investigate key enabling technologies that can make such transformations possible. Furthermore, computer-based modeling approaches for analyzing the benefits of such technologies for urban organization and activities are highly desirable. Adopting a model-based approach, we introduce in this paper, a technique for addressing sustainable urban environments, that employs a green edge computing technology to meet the needs of urban systems for energy and digital services. Our study evaluates the deployment of this technology in various urban contexts and demonstrates how it can help satisfy the crucial demand for power from digital activity by utilizing the solar energy potential of urban roofs. The proposed modeling framework is a valuable analysis tool for assisting decision-making for improving the quality of urban systems due to its unique features.

1. Introduction

Efforts are made to create sustainable cities that actively address environmental, social, economic, cultural, and institutional challenges [11]. To improve the quality of life, they strongly encourage sobriety in all aspects of life, including the consumption of goods and services such as transportation systems. They advocate, in particular, the optimal use of natural and renewable resources, which significantly contributes to lower energy consumption and carbon emissions. This necessitates the involvement of public governance in charge of planning and implementing the city transformation required by the sustainability challenge [26].

One key enabler solution to support this transformation is digital technology. Examples include big data collected by multiple sensors in cities for processing and analysis using AI techniques [10]. Edge and fog computing paradigms [27, 1], on top of the Internet of Things (IoT), have emerged as suitable answers to the sustainable digital transformation of cities. They employ geographically distributed mini data centers to manage the data and workloads generated by urban information and communications technology (ICT) on a local level. The ability of these computing paradigms to process data locally, i.e., close to the data sources, significantly improves information privacy and security. This feature is compatible with the critical data governance and sovereignty requirements of data-driven decision-making in urban contexts [35]. While edge and fog computing paradigms are

excellent candidates for satisfying the local demand from urban systems [26], their effective deployment in cities must consider certain factors. In particular, they must be implemented in such a way that their corresponding energy consumption remains sustainable and economically acceptable, i.e. without causing a detrimental increase in the public governance's electricity bill.

The purpose of the present study is to tackle the challenge of how to support digital services in urban neighborhoods and address the essential energy demand associated with them. The rationale behind the proposed approach assumes the deployment of a cutting-edge technology, called GENESIS (*generation of energy-synergistic information systems*) and introduced in the sequel. In this technology, data/workload computations are combined with renewable energy harvesting and management to extend the principles of the edge and fog computing. The emphasis of this paper is on how its second feature can successfully fill urban digital energy demands.

1.1. Addressed problem

We concentrate on the difficult problem of developing a modeling and analysis framework for the effective integration of sustainable technologies in urban settings to meet the energy demand from digital services. This is addressed by two major research questions, RQ1 and RQ2, as follows:

(RQ1): how can we employ the principles of edge computing coupled with renewables to provide a technological answer that fills the needs of a sustainable urban system, as discussed in the previous section? The critical challenge here is to propose a solution that is resilient, scalable, energy-efficient, and environment friendly.

(RQ2): how can the effective integration of our proposal be modeled in the urban context to anticipate typical

[★]Funding: This work was supported by the French CNRS [project IDs ISORE and AVIDON].

*Corresponding author: Abdoulaye Gamatié

✉ abdoulaye.gamatie@lirmm.fr (A. Gamatié);

thomas.leduc@crenau.archi.fr (T. Leduc); daniel.siret@crenau.archi.fr (D. Siret); gilles.sassatelli@lirmm.fr (G. Sassatelli);

michel.robert@umontpellier.fr (M. Robert)

ORCID(s): 0000-0002-8326-3257 (A. Gamatié); 0000-0002-5728-9787 (T. Leduc); 0000-0003-0091-7129 (D. Siret); 0000-0002-6396-286X (G. Sassatelli); 0000-0001-6179-9977 (M. Robert)




Figure 1: A typical deployment of the GENESIS paradigm for sustainable urban environment development.

issues like the impact on energy costs or the organization of the urban systems? The challenge raised by this key question is the development of a computer-based model capable of properly supporting the detailed description of urban contexts that incorporate our proposed technology. The model can make use of geo-informatics concepts and public databases to objectively assess the potential impact of this integration on sustainable cities.

1.2. Our contribution

We propose a model-based framework for investigating urban system deployments for long-term sustainability. As an answer to (RQ1), our solution relies on a novel green heterogeneous computing paradigm referred to as GENESIS, validated via a small-scale demonstrator in the laboratory. It aims to meet urban systems' demand for energy and digital services. Figure 1 depicts its typical urban deployment.

The heterogeneous nature of our solution stems from the combination of a large number of *green nodes* with a small number of in-door *mini data centers* (1–10 racks of servers within 1–25 square meters compute space [34]) to define a city-level data center. These mini data centers (Figure 1 bottom-right) are housed within secure buildings. They are linked to the power grid.

The green nature of the solution comes from the usage of specific nodes. Such nodes may combine in the same module: a server, a renewable energy harvester, local energy storage, and custom programmable power electronics

systems handling energy transfer between interconnected green nodes [30]. In particular, we consider solar panels and batteries for energy harvesting and storage. As suggested in Figure 1 (top-right), all of the electronic equipment, including servers and batteries, is installed on the roof¹, just behind the solar panels, to take advantage of passive cooling via appropriate packaging that ensures secure isolation.

The data processing demands of urban areas are met using the most local compute resources, preferably on green nodes to take advantage of free renewable energy. The harvested energy excess unused by the deployed green nodes can be re-injected into the utility grid thereby generating some potential financial revenues. A graphical user interface (see Figure 1 middle-right) provides remote monitoring and control of system resources such as harvested energy, energy demands, and so on.

To investigate a large-scale city-wide deployment of GENESIS, we define and implement a simulation tool that allows one to evaluate, for example, the impact on local energy consumption. This aims to answer (RQ2). The defined tool is made up of a Python implementation of a variety of functionalities orchestrated via a workflow: processing of alphanumeric census data and topographic data, retrieval and processing of solar ephemeral required for irradiation

¹They can also be installed in-door if necessary. The system's design additionally allows for this deployment mode. It should be noted, however, that the energy cost associated with in-door server cooling is not negligible. Furthermore, depending on the distance between the solar panels and the in-door equipment, there will be more energy losses during the transfer.

calculations, and integration and mapping of energy consumption and supply models.

The typical outputs of our tool are urban area maps (e.g., see left-hand side of Figure 1) with corresponding energy production and consumption estimates over time, while taking their specific weather conditions into account. For the sake of relevance, the tool manipulates data from public city databases as input. The models captured in our tool are partly calibrated using data from our GENESIS demonstrator, like computing resource dimensioning, energy harvesting and storage resource sizing, and energy transfer efficiency factor between green nodes. This enables to consistently address sustainability issues in urban systems based on an innovative paradigm such as GENESIS. Our tool is therefore a promising decision-making support, which is essential in city governance policy.

As illustrative case studies, we investigate a few GENESIS deployment scenarios in two different French cities as proof-of-concept: Nantes and Montpellier. Both cities differ in terms of urban layout and buildings typology, as well as weather conditions. We analyze the conditions under which, from an energy standpoint, such deployments can favor sustainability in relation to typical digital activities in each city district we studied. Ultimately, we discuss several relevant perspectives opened up by the current work, especially regarding the integration of GENESIS-like technologies within sustainable urban systems.

1.3. Outline of the paper

The rest of this paper is structured as follows. Some related work is presented in Section 2. The GENESIS green computing paradigm is introduced in Section 3. Then, Section 4 describes the modeling and evaluation of GENESIS's urban context-based deployment. Section 5 examines two deployment scenarios in different French cities in terms of energy consumption. Section 6 discusses some open questions raised by this work. Finally, Section 7 provides some closing remarks and perspectives.

2. Related work

The literature has previously discussed the relevance of the edge and fog computing paradigms for sustainable cities. Perera et al. [26] typically presented a general survey on this topic, giving insights into how fog computing can be used on top of IoT infrastructure in smart cities. They also go over some general desirable features of fog computing platforms, such as security and privacy, in relation to a subset of use case scenarios like smart transportation, smart waste management, smart health and well-being, smart power grid, and smart water management. Pereira et al. also focus on key governance aspects in the context of sustainable smart cities in [35]. They propose a roadmap as a tool for addressing sustainability issues, including the integration of urban infrastructures and ICT.

Several surveys on sustainable cities have been conducted, with the importance of adopting disruptive digital technologies [11, 3, 2]. Estevez et al. cover a substantial body

of work on smart sustainable cities in [11]. This includes suggestions for local policymakers, government practitioners, and academic experts. Furthermore, the authors discuss some relevant initiatives from various countries, as well as a conceptual framework to help with the process of planning, developing, and evaluating these initiatives. Beyond modern technologies such as mobile phones, sensor networks, big data, grids, cloud services, and IoT, they believe that open-source software, web services, and geo-informatics should be emphasized. These technologies improve citizens' quality of life. Modeling and simulation tools that predict the behavior of urban systems and citizens are valuable for assisting policymakers in decision planning and implementation.

In [3, 2], Bibri et al. emphasize the critical need for new wave computing ICT to be seamlessly integrated with sustainable urban systems. They identify suitable technologies and applications for sustainable cities, such as the potential combination of big data analytics and context-aware computing. Indeed, big data applications and technologies combined with IoT technologies enable the exploitation of stakeholder information to better manage sustainable cities [17]. The volume, velocity, and heterogeneity of data that characterize big data, on the other hand, necessitate appropriate energy-efficient high-performance computing infrastructures. Context-aware computing makes it easier to adjust urban service demands based on the location of users, available hosts, and devices [31].

Kong et al. [21] propose a comprehensive study on green-energy-aware approaches. They discuss various job scheduling approaches aimed at achieving favorable conditions for lowering energy costs in data centers. The conditions include the cost of electricity, the level of carbon emissions, and the availability of renewable energy. Cioara et al. introduce a methodology for operation planning to maximize the use of local renewable energy harvested in data centers in [6]. They specify a mechanism for shifting the energy demand from low renewable energy production time slots to higher energy production time slots. This allows for a reduction in power consumption from the utility grid. The shifting mechanism is based on a prediction of energy production. Based on an in-house simulation framework, the authors demonstrated that their approach increases renewable energy usage by 12%.

By leveraging renewable energy to improve the energy-efficiency of computing infrastructures, our green heterogeneous computing paradigm shares goals with approaches like [21] and [6]. However, in comparison to these approaches, our GENESIS-oriented approach incorporates an additional dimension, namely energy migration beyond job migration, which contributes to lower overall energy costs and creates added value for residents.

The modeling of the solar potential of urban roofs has been extensively addressed [13], based on the IEA's pioneer work [19]. Some studies focused on the issue at different analysis levels, ranging from national to local [18]. The quality of architectural integration of panels and their acceptability by the inhabitants are covered in [12]. On the other

hand, the impact of urban compactness on the photovoltaic potential of cities is investigated in work [23]. Here, a major challenge is to overcome the solar masking issue induced by urban buildings. More recently in [36], machine learning-based techniques are used to integrate the constraints of mask, roof structure, efficiency of the device according to its temperature, etc. Beyond the usage of photovoltaic panels, complementary solutions such as wind turbines and battery banks are equally considered to optimize distributed site-specific electrical hubs [25]. For this purpose, multi-criteria heuristics are applied to determine Pareto front solutions.

While the aforementioned studies are mainly devoted to the estimation of the potential photovoltaic capacity of roof surfaces, very few have compared this green energy potential w.r.t. the effective demand from urban environments. Costanzo et al. [7] compare measured electricity consumption with simulated irradiation on some campus buildings' facades and roofs (the deployment of PV panels being a ratio of the total envelope surface). Groppi et al. [15] propose a comparative approach between the energy supply generated by PV panel installation on buildings and their thermal (heating and hot water) and electrical needs. The assessment of the demand is mainly based on a building typology (construction period, materials, constructional size, nature and type of openings, etc.), independent of the actual number of inhabitants or uses. Using an ISTAT database, it estimates energy consumption based on building features.

Contrarily to our modeling approach, [15] does not take into account any detail about important parameters such as the number of inhabitants and their activity modes or usages, which are necessary for addressing more realistic urban scenarios. On the other hand, [7] does not consider advanced technologies, such as GENESIS, which aims at combining the usual energy demand with the increasing demand in computing resources and data storage in urban environments.

3. A green edge-computing technology for urban context

As mentioned in the introduction section, we consider a novel technological solution, which accounts for renewable energies and favors a scalable deployment in urban contexts. The following sections provide an overview of the technology design, which includes both hardware and software components.

3.1. Hardware part

Our technology is built from the ground up to meet the stated goals of sustainable computing. Each green node is based on a compute blade (a motherboard with CPUs, memory modules, and optionally a GPU) similar to those found in data centers. This compute hardware is powered by custom software-controlled power electronics that can accept a variety of power sources, including solar, battery, and utility grid. Green nodes can also exchange energy, allowing a node with a depleted battery to continue operation by utilizing energy from one or more remote nodes.




Figure 2: Green node architecture (backside)

Figure 2 depicts the hardware system organization with two green nodes. Each green node is represented by a rooftop solar panel with additional hardware installed on the backside: a compute board, batteries, and the so-called “powerboard”, which acts as a smart energy gateway connected to the solar panel, batteries, and neighbor nodes via dedicated DC wires for energy sharing. Some green nodes (right node on Figure 2) are linked to the utility grid, either to supplement renewable energy if harvested energy is insufficient and batteries are empty, or to inject any surplus energy for monetary compensation.

In most cases, a mesh configuration is used, which means that there is no single route between any two nodes in the system, but rather many. This raises an interesting optimization question, the concepts of which are briefly introduced in section 3.2, which is essentially where map compute jobs should be executed and from which node (local or remote) the energy should be drawn.

To ensure the best possible utilization of renewables, the system is equipped with the ability to perform these reconfigurations online, i.e. without interrupting compute job execution or energy delivery to compute hardware. The former is made possible by the use of mature technologies found in data center software stacks (virtual machine migration), whereas the latter is the result of careful design of the green node power electronics hardware. This hardware is capable of fusing power streams and thus enabling progressive handover operations, i.e. transitions from one energy source to another.

A green node uses its own harvested or stored energy whenever possible. If this is not possible, its assigned data/workload is migrated to remote green nodes, or energy is fetched from other green nodes so that it can continue to operate. Power from the utility grid is only used as a last resort to ensure operational readiness. The energy cost of workload execution on green nodes is zero as it is derived from free renewable energy.

Figure 3 depicts a conceptual representation of a single green node cluster instance in which one compute job is mapped to a green node with a nearly-empty battery that must rely on energy from two other nodes (green arrows). The use of utility grid electricity or the ability to push

electricity back into the grid is depicted by red arrows. This energy transfer is assumed to serve the purpose of covering local household needs, partially powering a neighbor indoor mini data center, or being transferred to other neighbor housings should a proper powerline exist or specific agreements with the provided electricity exist in the specific case explored in this paper.




Figure 3: Conceptual representation of energy transfer

3.2. Software part

A dedicated software infrastructure is used to manage the aforementioned technology. This consists of a full software stack that enables the operation of the available green computing nodes in conjunction with mini data centers. The infrastructure’s ultimate goal is to optimize the overall system behavior in terms of operating costs and the quality of service provided to end users. Technically, this entails the implementation and integration of the following major components.

The monitoring component uses a database to retrieve information about the operation of the deployed hardware computing system. The optimization component exploits some heuristics for an efficient orchestration of workload execution and energy allocation in the system. In particular, it uses mixed integer linear programming (MILP) [32] in conjunction with a rolling horizon heuristic [4] to formulate and solve the resource allocation problem. For this purpose, it also considers a number of input information including system utilization statistics, meteorological prediction data, etc. The resulting optimization decisions are applied via a control component.

The integration of all these modules with the OpenStack infrastructure [29] allows for the smooth and transparent management of the virtual machines (VMs) running on the various green nodes and remote indoor mini data centers. OpenStack is a popular and active open-source cloud computing infrastructure software project. It enables us to

handle job migration across a heterogeneous system, such as between supported green nodes and mini data centers. The considered resource allocation problem in our target distributed system takes advantage of both data (or virtual machines) and energy migrations across green nodes. A major goal is to reduce the amount of energy drawn from the utility grid as much as possible.

The system is managed by the operators or end users via a graphical interface that allows them to allocate some given computer resources, such as virtual machines, to monitor the system, and to exercise some remote control over the allocated resources. As shown in Figure 1, we use the Grafana dashboard monitoring [5] to display the state evolution of the prototype and to send orders for some system actions, such as data or energy migration or node switching off/on.

3.3. A small-scale prototype

To validate the concept, a prototype of four green compute nodes was deployed (see Figure 4). The images in the middle and right-hand side of the figure show a side-view and a top-view of the prototype system, respectively. Each green node has a solar panel with a surface area of 1.59 m^2 . Each node has a battery capacity of 2 kWh^2 . A node’s computing server is an Intel Xeon E5 bi-socket with 10 cores and 64 GB of RAM memory. The image on the left side of Figure 4 zooms in on the various components integrated within a green node. We can see the line inverter, which allows the harvested energy to be injected into the utility grid, and the powerboard, which is in charge of energy routing between the components, such as from battery to server, or from solar panel to server or battery.

Using the prototype, we tested a variety of execution scenarios, such as solar panel-powered and/or battery-powered operation modes. In order to calibrate some large-scale simulation models featuring the proposed green computing nodes, we also performed some useful measurements, such as energy transfer loss between different nodes based on various scenarios. The four-node prototype works with an in-house server (which serves as a mini data center) to migrate VMs between the two platforms using OpenStack.

3.4. Exploitation modes of GENESIS

The deployment of GENESIS green nodes considered in the scope of this study assumes that each node includes at least one solar panel as well as energy transfer logics that allow interaction with the rest of the system. It is optional to integrate batteries or compute servers (with associated data transfer logics). The GENESIS paradigm’s adaptability naturally leads to a variety of functional modes for exploitation in urban contexts.

The first mode primarily involves the use of its energy harvesting capacity on a local scale via the deployed solar panels. It enables direct use of harvested energy to compensate for the energy demand from urban systems. It can

²This battery capacity is fully customizable, i.e. it can be either increased or decreased. In our current prototype, the choice is simply economical.



Figure 4: Four-node prototype system deployment

also be injected into the utility grid to properly balance the overall energy distribution in a city. The second mode usually includes energy storage facilities, such as batteries, as well as energy harvesters. It enables compensating for energy demand through more flexible allocation of collected energy during critical situations, such as at night or in adverse weather conditions. The third mode makes use of all of GENESIS’s functionalities, including computing and energy management components. Urban activities, services and residents can use the computing systems integrated into GENESIS green nodes here. Typically, they can function as proximity data centers, processing or hosting local user data.

GENESIS’ computing systems, on the other hand, can be used as a Platform-as-a-Service (PaaS) cloud computing infrastructure. Potential customers would be able to develop, execute, and manage various applications without having to build and maintain the infrastructure themselves. A service provider is required in this case to manage the deployed GENESIS computing infrastructure. Then, a possible business model is similar to that of Amazon Elastic Compute Cloud³ (EC2). Indeed, EC2 offers users virtual CPUs and a memory budget for a few dollars per usage hour. The income from such a business for the rental of their roofs could benefit city dwellers.

4. Modeling of GENESIS urban deployment

Considering the hardware and software capabilities of GENESIS, we can now assess its relevance in actual urban areas. Our objective is to develop a model that compares the combined supply of digital and energy on the one hand with the demand for digital in buildings on the other. It is therefore imperative, first, to quantify the territory’s capacity to host the GENESIS system, and consequently, the amount of energy that can be collected and accumulated for digital needs. Additionally, we want to measure the amount of energy consumed by the residents of the territory to run their IT devices.

To spatialize our model, that is, to anchor it on a given territory, we need geospecific or geotypical data. These data

³<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/concepts.html> (Accessed March 2022).

are of three kinds. They must first allow us to characterize the solar deposit of the surfaces or roofs on which the photovoltaic panels of the system will be installed. Additionally, they must allow categorizing the population within the premises based on age. Here, we assume that the way digital technology is used varies based on which group a person belongs to. Lastly, the data should make it possible for household IT devices to be quantified by type. This number is based on data from a French national household survey.

4.1. Spatial calibration

To highlight the interest of the model and the resulting simulation tool, we need to identify an areal unit. The latter should be of an appropriate size, enabling us to not only carry out a set of calculations but also to provide a cartographic representation that is intuitive not only from the standpoint of a single urban block but also from the viewpoint of an entire building (since this is the scale at which the GENESIS green nodes are mounted). For these various reasons, the district’s scale seems appropriate.

Thus, our model works similarly to a Matryoshka doll, aggregating data at various levels. The GENESIS green node is an approximately one and a half square meter rectangular panel. Based on the building footprint and orientation, this green node may or may not be mounted on a building. Next, the buildings are grouped together at the city block level. Here, we assume that this geographic proximity will facilitate the pooling and redistribution of the energy accumulated in the green nodes. Lastly, we compare the potential of each block at the urban district level. We will see later in the article that we need to add an additional level to this stack. We must, indeed, disaggregate a collection of data contained within a regular grid.

4.2. Potentials for deployment on roofs

As a first step, we used standard topographic data (according to the European INSPIRE directive) to characterize the cities based on the heights of polygonal footprints. Based on this simplistic model of a city, every building in it is seen as a collection of right prisms, i.e. because there is no data immediately available, roofs are assumed as flat roofs, rather than being modeled according to their diversity. We must



Figure 5: Above the well-aligned facade is a 3.2 m deep roof ribbon (yellow shaded polygon) on which a GENESIS green node may be potentially installed. To avoid the sun shading effect, we removed any part of the roof ribbon that protruded beyond the party wall. Red polygons correspond to building footprints in the terrain data (BD Topo®).

adopt an alternative deployment strategy due to this general flattening of potentially sloped roofs (which does not include chimneys, air vents, and other skylights which can impede GENESIS panel installation).

To overcome the lack of fine modeling of urban roofs, we focused on detecting south-facing facades, assuming that we could deploy the panels in a 3.2-meter-deep ribbon (see 5). In practice, given the size of the panels we use, this depth allows two rows of green nodes to be inserted (see the right side of the figure 4). This is an input parameter of the model and can of course be modified. This deployment at depths of less than four meters (if arbitrary) appears to answer assumptions about roof structures in practice, and the acceptability and ease of deployment in morphologically constrained urban environments.

Geoprocessing used to assess the potential for installing GENESIS green nodes consists of three steps. The first step is to remove the shared walls by applying a geometric dissolution operation, which consists of a spatial union of the input geometries. It provides the resulting geometry without overlapping. The second step allows only well-exposed and sufficiently long sections to be selected from the outlines of various building plans. In practice, this treatment can leave only south-facing elevations, i.e. elevations whose azimuth is contained within 45° of the full south. The third and final step simply consist in paving each facade linearly from the previous filter in ribbons less than four meters deep (see previous discussion). This potential placement of green nodes along each facade is a function of the specific dimensions of the panels and mounting ribbon.

Note that our model does not account for all the complex effects of solar shading between buildings. However, these effects depend not only on the 3D geometry of the building, but also on the position of the sun over time. To overcome this limitation (and possibly the roof constraints), we incorporate as a parameter an arbitrary average percentage of

coverage that depends on the general morphology of the target urban area.

4.3. Population and IT usage estimation

The data from the population census is gridded to ensure individual anonymity. The disadvantage of this gridding in this study is that it requires us to disaggregate the census data $C \in \mathbb{R}^m$ in order to assign population estimates by building. C is a vector data of dimension m (from now on and throughout the article, m refers to the size of the census space) because it embeds various dimensions relating to the various age groups, but also to the number of households. In a preliminary version of the model, we assume that the population distribution $C(b) \in \mathbb{R}^m$ per building footprint b is simply a function of its floor area relative to the total floor area of the study area, as shown in Eq. 1, as follows:

$$C(b) = \frac{\text{area}(b)}{\sum_{b' \in B} \text{area}(b')} \times C \quad (1)$$

The set of building footprints in the census area is represented by B in this equation. Of course, this estimate is completely arbitrary. It is only a statistical representation of the premises' occupancy.

In our simplified model, the quantity of population $P(b, t)$ actually present in buildings and likely to use IT equipment is a function of both this disaggregated census data $C(b)$ and time t . We assume that the occupation of the dwelling varies according to age groups during the day, depending on the nature of the day (weekday, weekend day, etc.) and even the context (business as usual period, lockdown, generalized partial teleworking, etc.). This variable quantity is a function that, at any time $t \in \mathbb{R}$, associates the value of the selected IT equipment usage scenario $S(t) \in \mathbb{R}^m$. By scenario, we mean in this model the usage of IT equipment in the household, by age group and for a given date. An example of usage is described in Table 4. It is an arbitrary linear combination of time-indicator functions representing a pattern of usage.

Let \odot denote vector pairwise multiplication. The population actually present in the building b at time t and likely to use IT equipment, denoted by $P(b, t) \in \mathbb{R}^m$, can be written as the pairwise product of $S(t) \in \mathbb{R}^m$ and $C(b) \in \mathbb{R}^m$ as exhibited in Eq. 2, as follows:

$$P(b, t) = S(t) \odot C(b) \quad (2)$$

The above uniform distribution of age groups across buildings is based solely on the floor area criterion. It is purposefully kept simple for the sake of proof of concept in our current tool implementation. By taking into account household composition, a less trivial settlement model could be developed.

4.4. Energy balance related to digital activity

We will refer to the number of IT equipment to be evaluated in our study as n from now on and throughout the rest of the paper. The goal of this section is to introduce $D \in$

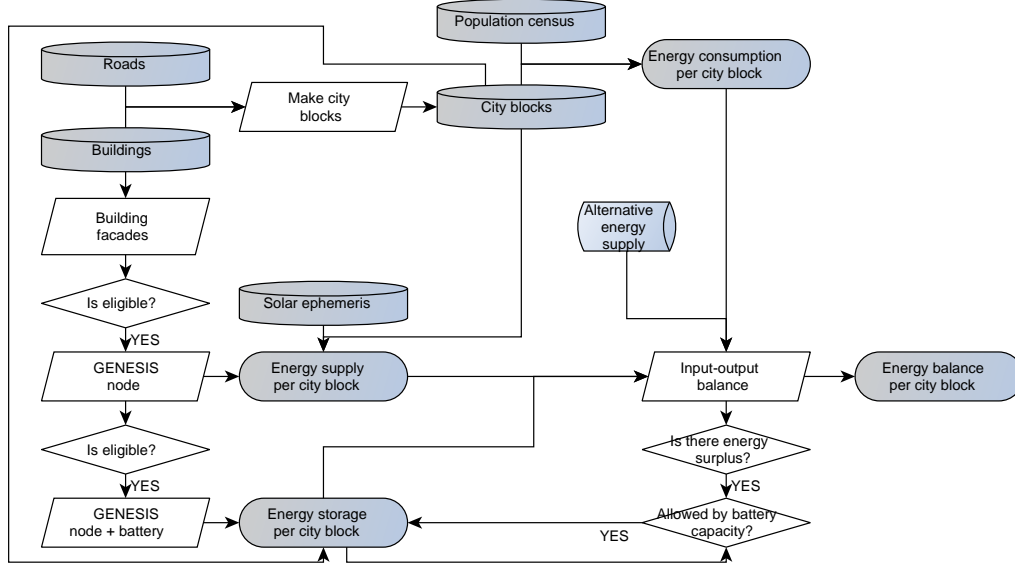


Figure 6: Presentation of the whole geoprocessing workflow.

$\mathbb{R}^n \times \mathbb{R}^m$. In practice, D is a matrix of n rows by m columns that assigns a possession ratio D_{ij} to each type of equipment and age group. Because they are derived from national level censuses, these equipment ratios are not geospecific.

From a high-level perspective, the energy required to operate the IT equipment of a given territory is determined not only by the number of inhabitants and equipment, but also by the actual consumption of each device. The latter can even be variable depending on the specified date (in the case study to follow, we will typically decide to charge the smartphones only once a day at 7 pm). Let $E(t) \in \mathbb{R}^n$ be the function that models power consumption over time by IT equipment type. The electricity consumption $\mathcal{C}(b, t)$ – related to IT equipment usage – of building b at date t can then be written as the inner product of $D \times P(b, t)$ by $E(t)$, see Eq. 3 as follows:

$$\mathcal{C}(b, t) = \langle D \times P(b, t), E(t) \rangle \quad (3)$$

We can see that by associating a pattern of presence with each age group on the one hand and a ratio of ownership of various IT devices on the other, the above model assumes that the same individual uses all of his/her devices at the same time. This assumption is not always realistic.

4.5. Mapping of GENESIS to digital energy demand

All of the processing described in the preceding sections is grouped into the workflow depicted in Figure 6, which is written in Python. This relies on the integration of various libraries, including Pandas [28] for processing alphanumeric census data; shapeLy [14] and GeoPandas [20] for processing topographic data; PySolar [33] for retrieving and processing the solar ephemeris required for irradiation calculations, and t4gpd [22] for integrating and mapping the energy consumption and supply models. It is also worth mentioning the use of

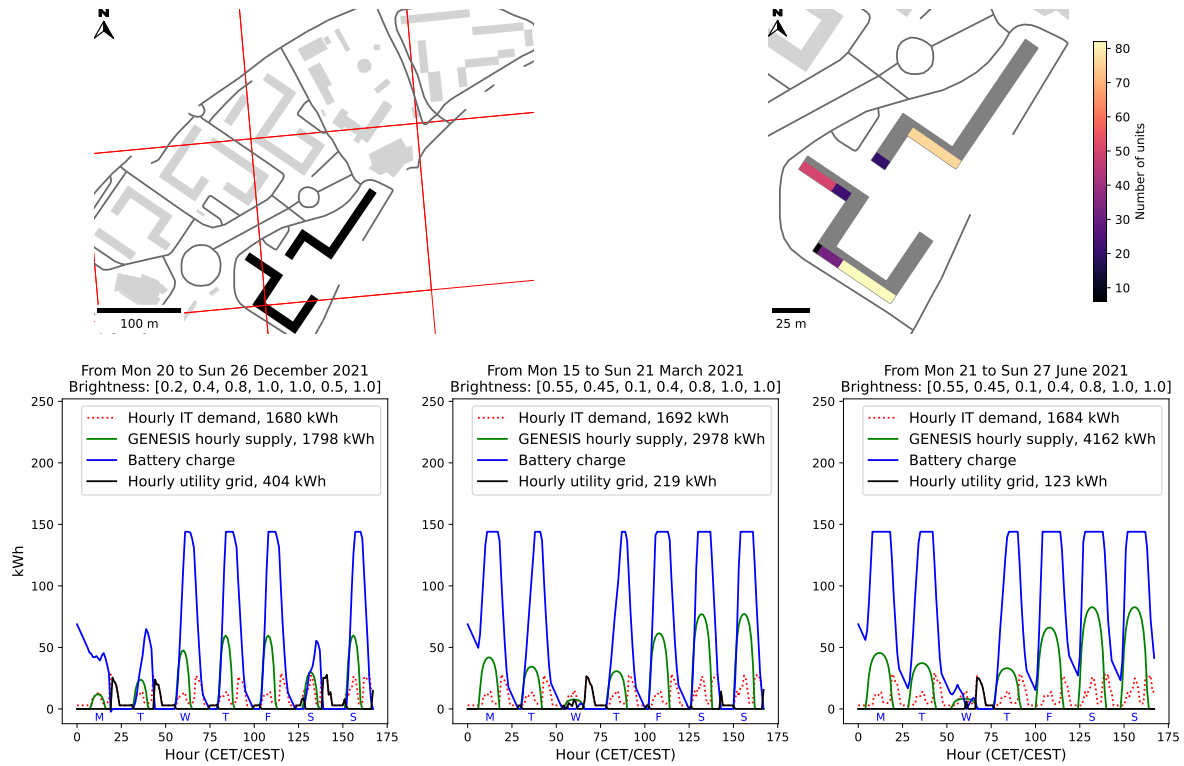
the NumPy [16] library which, thanks to its array programming capabilities, makes the various energy consumption calculations very efficient. In our simulator, for example, we have translated Eq. 3 as follows: `dot(matmul(D, multiply(S, C)), E)`, with the functionalities `dot`, `matmul`, and `multiply` being directly imported from NumPy.

The Direct Normal Irradiance (DNI) is required to calculate the solar deposit. The DNI parameter represents the amount of solar radiation received per unit area by a surface in an instant. This surface is preferably considered perpendicular (or normal) to the rays emanating from the sun's current position in the sky. We derive the DNI value from topographic data extraction by calculating the apparent solar path associated with the latitude and longitude of a given urban site. Because this value estimation assumes a clear sky, we have chosen to weight the daily solar contribution by a sky brightness coefficient: the closer this coefficient is to zero, the more cloudy it is. However, the closer this coefficient is to one, the clearer the sky.

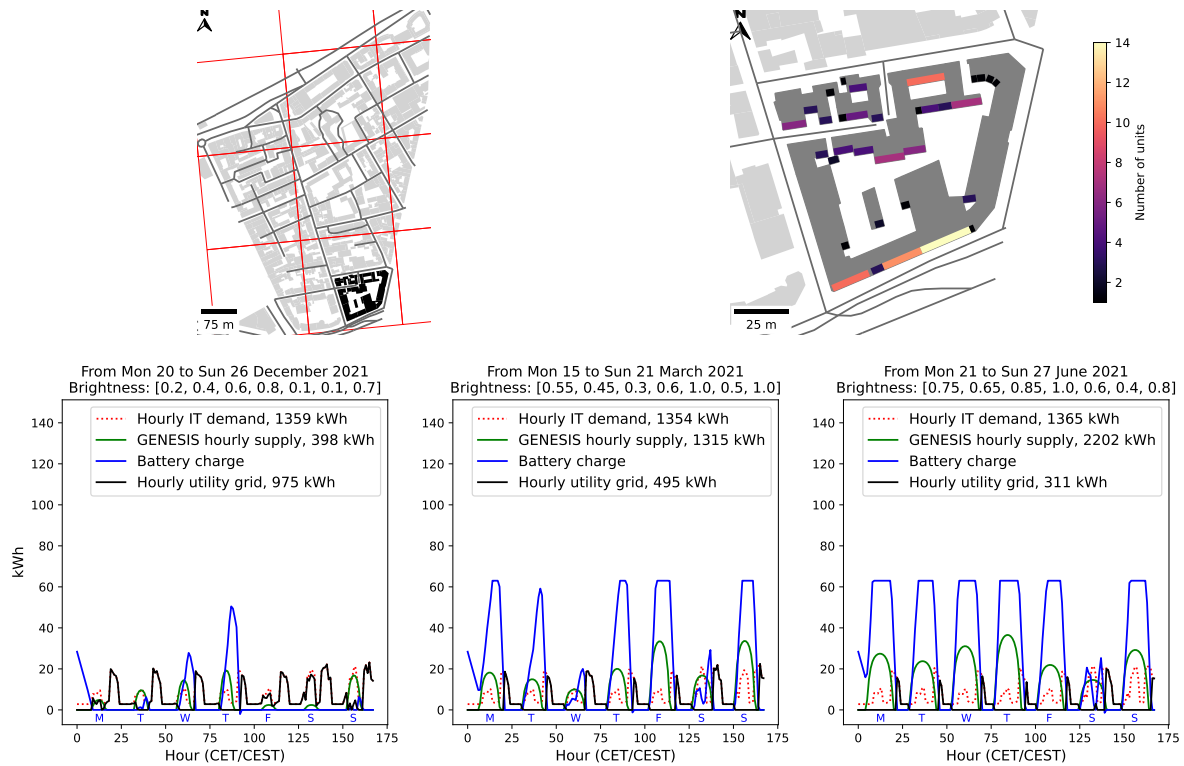
Figures 7a and 7b depict typical outputs (maps and curves) directly generated by our GENESIS deployment simulator in two French urban areas. As discussed in Section 5, we implemented the presented model on the “Pas du Loup” district of Montpellier city (43.6° N, 3.8° E) and the “Madeleine” district of Nantes city (47.2° N, 1.5° W). We note here that Montpellier has a sunny Mediterranean climate unlike Nantes which has a more often overcast oceanic climate. The simulator allows to combine actual topographical data (such as building footprints), IT equipment usage data, and solar energy harvesting data at various times of the year. It also easily handles various organic urban morphologies, as demonstrated by the two examples.

To run these simulations, a number of parameters, including the aforementioned ones, must be specified, such

A model-based framework for addressing the energy demand in sustainable urban systems



(a) Estimate for a given city block, district of "Pas du Loup" (Montpellier). In this case, the total energy collected by the 288 green nodes exceeds the storage capacity of the 48 battery-equipped nodes. To provide free harvested energy to the city block's 293 residents while avoiding strain on the utility grid, storage capacity must be increased.



(b) Estimate for a city block, district of "Madeleine" (Nantes), 186 residents, 128 green nodes, 21 battery-equipped nodes. The specific urban constraints impose a one-third reduction in the ribbons' coverage rate with GENESIS in this case. The difference in energy demand from the utility grid between winter and summer demonstrates GENESIS's role in reducing this demand for IT equipment by two-thirds.

Figure 7: GENESIS deployment to answer energy demands from local IT equipment in Montpellier and Nantes cities.

Table 1

Parameters and corresponding values for Figure 7.

Parameter	Montpellier	Nantes
GENESIS coverage	50%	30%
Ribbon width	3.2 m	3.2 m
GENESIS batt. ratio	1:6	1:6
Battery capacity	3 kWh	3 kWh
PV panel size	1.6m × 1m	1.6m × 1m
PV panel efficiency	19%	19%
Loss rate during (dis)charge	90%	90%
Min. battery charge level	15%	15%
Time step	1h	1h
South azim. range	270 ± 45°	270 ± 45°
Winter brightness	see Fig. 7a	see Fig. 7b
Spring brightness	see Fig. 7a	see Fig. 7b
Summer brightness	see Fig. 7a	see Fig. 7b

as the brightness coefficient, as indicated in Table 1. The GENESIS nodes equipped with batteries are used to store the surplus harvested solar energy, i.e. when it exceeds the energy demand from a given city block (see parameter named GENESIS batt. ratio in Table 1). When storing the excess energy in the battery, we apply a Loss rate during (dis)charge (also known as round-trip efficiency, see Table 1). Similarly, when the harvested solar energy supply to GENESIS nodes is insufficient and the battery charge exceeds a threshold (Min. battery charge level), the discharge is amplified by the same loss rate.

In addition, we consider a parameter related to the Battery capacity, which plays a significant role in whether or not the energy demand of IT equipment can be met without using energy from the utility grid.

5. Case study

We consider the “Pas du Loup” and “Madeleine” districts in Montpellier and Nantes, respectively, to demonstrate the relevance of our simulation tool. We investigate various GENESIS deployment scenarios at the district level in order to determine the most advantageous configuration in terms of energy consumption from the utility grid (which is costly and not necessarily green). It should be noted that the PaaS exploitation mode of GENESIS (see Section 3.4) is not taken into account in this study due to the absence of public databases on the cloud data processing demand per residents in urban environments.

5.1. Gathering data to build simulation models

The urban topographic data that we used are freely available online at the French IGN Géoportail site⁴.

We used the data tables from a technical report [8] to calculate the distribution of IT equipment per inhabitant. The smartphone ownership ratios by age group are reported for the year 2019 in [8], Table 4, page 35. The ratios of computer equipment by age group are given for the same year in [8], Table 6, page 42. Similarly, the ratios of tablet ownership are reported in [8], Table 8, page 46. The ratios of smart speaker

⁴See <https://geoservices.ign.fr/bdtopo> (Accessed March 2022).

Table 2

Assumed distribution of IT equipment rate (in %) by age groups. In the current study, this distribution is a matrix D , composed of $n = 7$ rows (IT equipment categories: gwy = gateway, ltp = laptop, tbl = tablet, sph = smartphone, sob = smart object, spk = smart speaker, tv = television) and $m = 8$ columns (hld = household, and age groups). For instance, 91% of individuals in 11 – 17 age group hold a laptop, while only 58% of 65 – 79 age group hold a laptop.

		11	18	25	40	55	65	
	hld	-	-	-	-	-	-	80+
		17	24	39	54	64	79	
gwy	95	0	0.0	0.0	0.0	0.0	0.0	0.0
ltp	0	91	82.0	73.0	83.0	76.0	58.0	0.0
tbl	0	37	36.0	49.0	45.0	47.0	29.0	0.0
sph	0	86	98.0	95.0	80.0	62.0	44.0	0.0
sob	0	31	25.0	21.0	17.0	9.0	5.0	0.0
spk	0	15	14.0	1.0	9.0	7.0	5.0	0.0
tv	0	0	79.2	93.1	95.3	97.9	97.9	0.0

Table 3

Assumed average energy consumption (in Wh) per IT device. In this study, this energy consumption is a function of time t . It thus associates with any $t \in \mathbb{R}$, a vector $E(t) \in \mathbb{R}^n$ with $n = 7$ components (IT equipment, see legend of Table 2). Let h be the hourly component of t (obtained via the $t.hour$ attribute as stored in the Python *datetime* instance). The fourth component of $E(t)$ is based on a Dirac unit impulse, which takes the value 1 only at 7 pm (it is assumed that each smartphone is recharged between 7 pm and 8 pm) and is zero otherwise. The other components are all time independent constants.

	gwy	ltp	tbl	sph	sob	spk	tv
Energy	26	75	4	10× δ (h-19)	1	5	100

are reported in [8], Table 9, page 49. The ratios of connected smart objects are reported in [8], Table 11, p. 54. We added up the ratios of connections by ADSL or Optic Fiber from [8], graph 55, p. 80, for the connection of households to the Internet. Finally, information on TV equipment ratios by age group is extracted from the French INSEE “*Equipelement des ménages*” survey, available online⁵.

In practice, and in order to specify each of the model’s parameters, we have grouped all of these parameters in the matrix variable D shown in Table 2. We have grouped the estimated average consumption data per IT equipment in the $E(t)$ vector shown in Table 3. These data are rough estimates that can naturally be modified. Finally, Table 4 shows an example of a typical usage scenario for IT equipment based on the day and time of day, as well as the type of equipment and the user’s age group. The data in this table correspond to our model’s $S(t)$ matrix (see Eq. 2).

Table 4 describes a typical scenario of IT equipment usage. This scenario is completely arbitrary and is referred to as *standard usage* (i.e. business as usual period) in the

⁵See Fig. 1 at <https://www.insee.fr/fr/statistiques/4277714> (Accessed March 2022).

Table 4

Assumption of IT equipment usage ratios (in %) according to age group, baseline scenario $S(t)$ during standard usage, i.e. no lockdown. Note that $S(t) \in \mathbb{R}^m$, with $m=8$ (rows of the table), is a function of two integer variables h ($t.hour$ attribute as stored in the Python *datetime* instance) and w ($t.weekday()$, the day of the week starting with 0 for Monday). The function $\Pi(t_0, t_1, t)$ corresponds to a modified indicator function. It is equal to 1 on the semi-open interval $[t_0, t_1[$ and 0 elsewhere. Φ_x denotes here a normal continuous random variable $\Phi_x \sim \mathcal{N}(x, \frac{x}{10})$.

Age groups	IT equipment usage ratios
household	100
11-17	$\Pi(0,5,w) \times (\Pi(12,14,h) \times \Phi_{20} + \Pi(18,20,h) \times \Phi_{70} + \Pi(20,23,h) \times \Phi_{50}) + \Pi(5,7,w) \times (\Pi(8,13,h) \times \Phi_{90} + \Pi(18,20,h) \times \Phi_{70} + \Pi(20,22,h) \times \Phi_{90})$
18-24	$\Pi(0,5,w) \times (\Pi(19,21,h) \times \Phi_{70} + \Pi(21,24,h) \times \Phi_{60}) + \Pi(5,7,w) \times (\Pi(10,15,h) \times \Phi_{50} + \Pi(21,24,h) \times \Phi_{70})$
25-39	$\Pi(0,5,w) \times (\Pi(9,12,h) \times \Phi_{20} + \Pi(12,14,h) \times \Phi_{30} + \Pi(19,21,h) \times \Phi_{60} + \Pi(21,24,h) \times \Phi_{70}) + \Pi(5,7,w) \times (\Pi(10,15,h) \times \Phi_{90} + \Pi(18,24,h) \times \Phi_{90})$
40-54	$\Pi(0,5,w) \times (\Pi(9,12,h) \times \Phi_{30} + \Pi(12,14,h) \times \Phi_{35} + \Pi(19,21,h) \times \Phi_{80} + \Pi(21,23,h) \times \Phi_{70}) + \Pi(5,7,w) \times (\Pi(8,10,h) \times \Phi_{50} + \Pi(12,14,h) \times \Phi_{60} + \Pi(18,22,h) \times \Phi_{70})$
55-64	$\Pi(0,5,w) \times (\Pi(9,12,h) \times \Phi_{40} + \Pi(12,14,h) \times \Phi_{35} + \Pi(19,21,h) \times \Phi_{80}) + \Pi(5,7,w) \times (\Pi(7,9,h) \times \Phi_{90} + \Pi(12,14,h) \times \Phi_{60} + \Pi(18,22,h) \times \Phi_{70})$
65-79	$\Pi(7,9,h) \times \Phi_{20} + \Pi(9,12,h) \times \Phi_{40} + \Pi(12,15,h) \times \Phi_{90} + \Pi(15,17,h) \times \Phi_{50} + \Pi(17,22,h) \times \Phi_{70}$
80+	$\Pi(9,13,h) \times \Phi_{60} + \Pi(14,22,h) \times \Phi_{40}$

sense that it does not refer to either partial or full lockdown scenarios as experienced during the recent Covid-19 issue, where the usage of IT equipment has significantly evolved. The function $\Pi(t_0, t_1, t)$ in this table corresponds to a modified indicator function, which is equal to 1 on the semi-open interval $[t_0, t_1[$ and 0 elsewhere. Φ_x here denotes a normal continuous random variable $\Phi_x \sim \mathcal{N}(x, \frac{x}{10})$. It allows for the expression of the IT equipment usage model for each age group. Typically, for the 18-24 age group, the following expression:


$$\Pi(0,5,w) \times (\Pi(19,21,h) \times \Phi_{70} + \Pi(21,24,h) \times \Phi_{60}) + \Pi(5,7,w) \times (\Pi(10,15,h) \times \Phi_{50} + \Pi(21,24,h) \times \Phi_{70})$$

means that, for any weekday, i.e. when the variable w has a value between 0 and 4, there is approximately 70% of the IT equipment used between 7pm and 9pm and approximately 60% of the IT equipment used between 7pm and midnight. On the other hand, on weekends, i.e. when the variable w has a value between 5 and 6, the usage decreases to approximately 50% between 10am and 3pm, and increases to approximately 70% between 9pm and midnight.


Figure 8 shows an instantiation of the IT equipment usage scenario (see Table 4) by age group for a typical weekday and a typical weekend day. This representation highlights the peaks in usage over a day.

5.2. Exploration of GENESIS urban deployment

We study the performance of possible GENESIS deployments at the scale of an urban district as a function of: the number of green nodes equipped with solar panels on the one hand, and the number of green nodes additionally equipped with initially half-full batteries on the other hand. The performance indicator we chose corresponds to the ratio of energy requested from the utility grid to total energy demand for the district's IT equipment over each evaluated week period. The 3D surfaces shown in Figures 9a and 9b demonstrate that yields are naturally a function of time and



(a) Weekday.



(b) Weekend day.

Figure 8: Graphic translation of the IT equipment usage scenario presented in Table 4 for two typical days.

that they deteriorate in winter due to sunlight conditions (less irradiation, increased cloudiness).

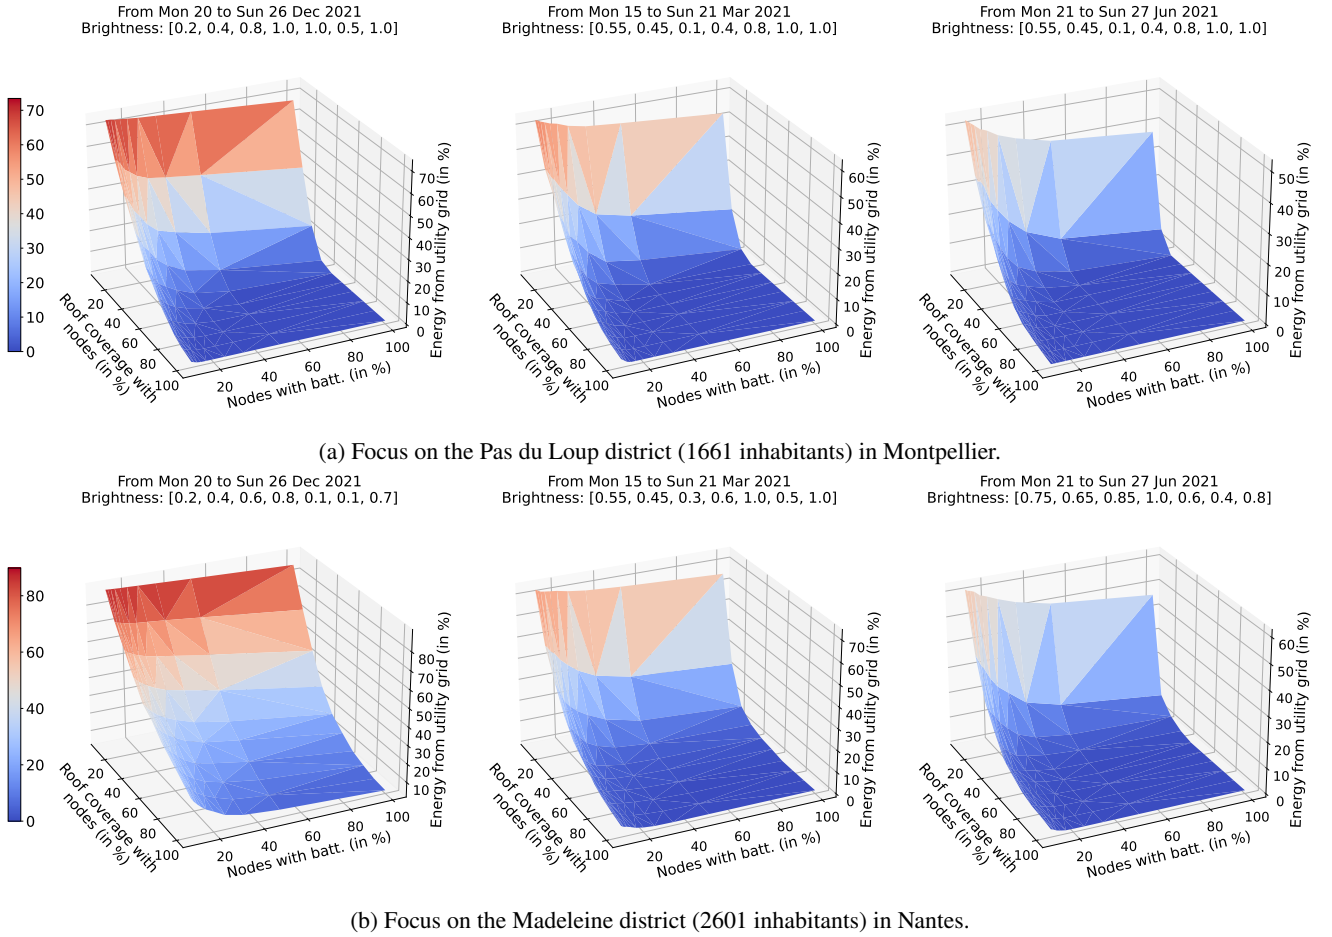


Figure 9: Exploration of different deployment conditions of GENESIS w.r.t. energy required from the utility grid. Horizontal axes denote i) the percentage of roof surface coverage with green nodes, and ii) the percentage nodes equipped with batteries out of the nodes deployed on roofs. The vertical axis denotes the percentage of energy from the utility grid required for satisfying the local IT equipment demand, in addition to the harvested energy.

In Figure 9a, if we consider that 50% is a reasonable ribbons' coverage rate (in terms of financial investment and technical feasibility) for the rooftop configuration in the Montpellier district, we get a total of 2294 GENESIS nodes. We can then vary the ratio of nodes including batteries after we have fixed this parameter. The situation is much better than in Nantes because, despite having less than 25% of the nodes equipped with batteries, the GENESIS solution meets more than 90% of the energy needs for IT equipment in the worst of the three configurations. With such a battery ratio, one can almost compensate the energy consumption of IT equipment over the whole night for the configuration given at the spring equinox or summer solstice.

In Figure 9b, if we assume that 30% is a realistic ribbons' coverage rate for the Nantes district's roof configuration, we get a total of 2566 GENESIS nodes. We can then vary the ratio of nodes carrying batteries after we've fixed this parameter. If 33% of the nodes have a storage capacity of 3kWh each, the GENESIS solution will meet 95% of the energy needs for IT equipment at the summer solstice. In contrast, at the winter solstice, the solution degrades significantly and only 42% of the needs are met. It should be noted

that in the intermediate regime, at the spring equinox (and under the weather conditions specified in the curves' titles), more than 83% of the needs are met by the solution. For the same ribbons' coverage rate in GENESIS nodes (30%, i.e. 2566 nodes), we can see that the solution never allows us to exceed the optimum yield of 96.5%, whatever the battery ratio studied. The tool clearly shows that increasing the storage ratio is not very useful; instead, it is preferable to try to increase the number of green nodes, and thus the harvesting capacity.

The results shown in Figure 9 are very useful for a quick and global analysis of the decision space for GENESIS deployment in urban contexts. They have been obtained using our simulator within 166 seconds in the case of the Pas du Loup district (8 city blocks in Montpellier), while 188 seconds were required for the Madeleine district (33 city blocks in Nantes). The simulations were run on a Dell Precision T1700 PC, Intel® Xeon® CPU E3-1246 v3 at 3.50GHz, with 16 Go RAM. Based on the estimated performance, stakeholders can assess the feasibility of the corresponding investments in terms of system component cost in relation to local weather conditions. In general, we can see from

Figure 9 that the best performance level (i.e. the least amount of energy required from the utility grid) is mostly sensitive to the proportion of solar panel-equipped green nodes determining roof coverage, as opposed to the proportion of battery-equipped green nodes. This observation, of course, is heavily reliant on the arbitrary input data used in the current case study.

5.3. Fine-grain analysis of deployments

The previous global analysis can be refined by focusing on specific GENESIS deployments to evaluate in detail the evolution over time of the energy demand from utility grid in relation to three relevant variables: energy supplied by GENESIS solar panels, accumulated battery charge, and energy demands for IT equipment.

Figure 10 demonstrates two specific GENESIS deployment options at the district scale in the cities of Montpellier and Nantes. The simulation times are respectively 214 seconds and 255 seconds for the Pas du Loup (8 city blocks in Montpellier) and Madeleine (33 city blocks in Nantes) districts, using the same the desktop PC as previously. In the case of Pas du Loup district, the scenario consists of 2294 green nodes that provide 50% ribbons' coverage. Among these nodes, 1147 modules each have a 3 kWh battery. In the Madeleine district, the proposed deployment uses 2566 green nodes to provide 30% ribbons' coverage. Among these nodes, 843 modules contain each 3 kWh battery.

The temporal evolution of the different curves shown in Figure 10a suggests that the considered deployment setup is adequate to meet the energy demand for IT equipment in the Pas du Loup district over different emblematic weeks. There is no need for energy from the utility grid, so there is no black curve visible in the plots. Given the global evaluation depicted in Figure 9a, this configuration is even a near-optimal choice.

As for the Madeleine district, the results shown in Figure 10b indicate that the considered deployment setup cannot avoid utility grid energy usage. This is most noticeable in December and March, when the black curve is clearly visible. Given this knowledge, stakeholders may wish to increase the number of green nodes to cover more than 30% of the available ribbons' surface. Unfortunately, due to architectural constraints and roof superstructures in the building, this is not possible. As a result, the configuration depicted in Figure 10b is the near-optimal choice.

5.4. Some preliminary insights

The above case study demonstrates that the deployment of a solution at the crossroad of edge computing and smart grid like GENESIS, can theoretically fill the energy demand for the residents digital activities within two very different districts of the French cities of Montpellier and Nantes. In particular, our simulation tool, which relies on a geographic information system (GIS), maps the digital power demand of a neighborhood with the energy harvesting capabilities of the distributed GENESIS modules, deployed within the same area.

5.4.1. On the key role of energy transfer capability

We note that the energy transfer mechanism of GENESIS plays a crucial role in the mutualization of the harvested energy in an urban area. Typically, the energy budget unused by the residents of some buildings are migrated towards some other buildings where the local residents have a demand. This feature, among others, distinguishes our GENESIS-oriented solution from similar approaches like [7] [15] as discussed in Section 2. The same observation can be made on the computation side when GENESIS is used in its PaaS exploitation mode. Indeed, the possibility of migrating data or workloads between different green computing nodes and mini data centers enables to relieve overloaded servers by offloading part of the jobs on idle servers, within a neighborhood. This promotes the concept of short energy and digital channels, and creates a direct and tangible link between city dwellers and the energy and computation infrastructures that meet their digital needs.

5.4.2. On the scalability of our simulator

The simulation time required for obtaining the results shown in Figure 9 is very reasonable. Even though, our case study only concentrates on district level, we can extrapolate the simulation time at upper levels. Roughly speaking, assuming in average 3 minutes simulation time per district, it would take approximately 4 hours and 25 minutes to evaluate a larger scenario at the level of Montpellier city, composed of 88 districts. For the city of Nantes, comprising 97 districts, the city-wide evaluation would require 4 hours and 50 minutes. For a large city in France such as Paris, which contains 992 districts, the extrapolated simulation time would be about 50 hours.

Pushing further this extrapolation to the department level, which consists of several cities, the simulation time could reach up to 25 hours for the Loire-Atlantique department, comprising Nantes city, at a total of 503 districts. Now, considering the Hérault department, including Montpellier city, at a total of 578 districts, the extrapolated simulation time would be about 29 hours.


It is important to notice that the above extrapolation does not consider any parallelization implementation of our tool. Such an evolution will undoubtedly improve the scalability for addressing very large scale scenarios, e.g. country level. The metropolitan France and Corsica cover 48,589 districts.

6. Discussion and open questions


The problem addressed in the present work, like all problems related to urban systems, is not only technical, but also economic, social, and cultural. The discussion in the sequel is therefore divided into two parts: the feasibility of an urban deployment of GENESIS on the one hand, and its integration into urban ecosystems on the other.

6.1. Feasibility of urban deployment for GENESIS

Is GENESIS feasible in theory, but not in practice? This question has no single answer and depends on a variety of local factors unique to the urban system under consideration.



(a) Focus on the Pas du Loup district: 1661 residents, 2294 green nodes (50% coverage) out of which 1147 include each a battery of 3 kWh.



(b) Focus on the Madeleine district: 2601 residents, 2566 green nodes (30% coverage) out of which 843 include each a battery of 3 kWh.

Figure 10: Evolution of the estimated energy collected by all GENESIS nodes (in green), the estimated Wh consumed by IT equipment (in dotted red), the charging rate of the batteries embedded in the nodes (in blue), and the demand for additional energy from the utility grid (in black) over a given week (in winter, spring, and summer).

6.1.1. Main feasibility factors

The first and most obvious factor is related to local climatic conditions. The operation of GENESIS is based on adequate and consistent solar gain. The case studies presented in the preceding section of this article demonstrate that the solution is feasible for two French cities with vastly different climatic conditions. GENESIS is estimated to be viable in all cities where solar collection is feasible.

The availability and accessibility of roof surfaces that are well exposed to the sun and capable of supporting an additional equipment load is the second factor. The effect of solar shading can reduce the potential solar exposure of rooftops in very dense urban structures such as city centers. Building accessibility may be difficult and costly, and existing roofs may be unable to support additional loads.

Another consideration is the possibility of regulatory constraints, such as heritage protection in certain historical zones, which may preclude the addition of equipment visible from public space. Another critical factor is the number of owners in the area in relation to the number of buildings,

i.e. the number of interlocutors who must be persuaded to install collective equipment such as GENESIS. The number of interlocutors in residential condominiums is generally much greater than the number of buildings, which obviously complicates the decision to implement collective equipment. The number of interlocutors in suburban areas is roughly equal to the number of buildings (individual houses). In business or commercial areas, the number of interlocutors may be much lower than the number of buildings, as in the case of an investor who owns a group of buildings rented to companies or stores. Negotiations for the implementation of GENESIS-type equipment can be simplified in this case.

On the other hand, the population living or working in the district, as well as their specific digital needs on both a personal and professional level, is a factor to consider. The goal of promoting short digital channels leads to a preference for areas with high demand for computing or digital storage. These include, for example, service companies such as banks, third places, local government services or transportation companies, universities, and so on.

Finally, while not strictly necessary, note that coupling of green nodes with in-door mini data centers (see Figure 1) strongly enhances the acceptability of GENESIS in urban contexts. Indeed, because green nodes are deployed on roofs, data security may appear sub-optimal at times, e.g. potentially vulnerable to vandalism. In this case, the green nodes could be used as *memoryless* compute nodes (that is, they process data without storing it in their local memory), while the in-house mini data centers would collect the results produced by the remote green nodes.

6.1.2. Characterizing the technical feasibility

We can characterize the technical feasibility of a GENESIS deployment based on the type of neighborhood considered by combining the aforementioned factors. A dense, central urban neighborhood has a large population and a high demand for digital services, but the rooftops may be difficult to access, regulatory barriers may exist, and the number of parties involved may be too large to reach an agreement. Only public-sector incentives are likely to result in the integration of a GENESIS-type solution in this type of neighborhood.

Private or public multi-family housing developments (old or new apartment buildings in a residential area) appear to be more suitable candidates. They typically have a large number of easily accessible flat roofs, fewer regulatory obstacles, a large population and thus a high demand for digital services, and a smaller number of discussion partners, particularly for social housing managed by a single social landlord. The difficulty here is related to the homogeneity of digital demand, which may result in overloads at certain times (in the evening or on weekends) and vacancies in IT activity demands during the rest of the time. One solution could be that the GENESIS computing capacities put in place by the co-ownership are rented to other economic actors in the area during periods of low demand. In France, for instance, the Qarnot⁶ Company charges in average 0.25 € per CPU.hour. A GENESIS prototype green node is currently equipped with a 16-core server with a peak power of 400 W. The energy consumed by all 16 cores for an hour is then 400 Wh. Accordingly, an estimated energy surplus of 1.217 MWh (obtained within a week in June when simulating the GENESIS deployment in the Pas du Loup district, see Figure 10a), would generate about 48,680 CPU.hours per week, corresponding to 12,170 € as financial return per week. Though this is a rough approximation, it suggests that the computation components of GENESIS green nodes could contribute to refund the investment costs.

Other potential candidates for the installation of GENESIS solutions include commercial and industrial zones on the outskirts of cities. The structures are numerous, easily accessible, and have large, well-exposed roof surfaces. Regulatory barriers are generally low, and the demand for digital services can be extremely high. Furthermore, the number of partners can be reduced, making installation and the search for a viable business model to finance the solution easier.

6.2. Impact of GENESIS on urban ecosystems

Assuming that GENESIS finds a concrete application in an urban system, it will be necessary to assess its integration into the urban ecosystem, which is defined here as the city's or district's ecological, social, economic, and cultural environment. A few critical factors must be considered, as discussed further below.

The first factor is the GENESIS hardware system's acceptability in a dense urban environment. It is obvious that integrating the system as close to living areas as possible is only possible if the green nodes are sufficiently quiet. This is determined by the integrated computing system. Our small-scale demonstrator currently includes a compute blade with GPU, similar to what is found in a typical data center. The sound emission of the compute blade depends on the executed workload intensity, which can result in the activation of the system's integrated fans in extreme heat conditions. In general, the computing portion of a green node uses a generic integration interface that supports any computer motherboard, such as those found in server-class systems, personal computers, or embedded compute boards, to be connected. The sound emission level varies depending on the options chosen, but it must be reasonable enough to allow GENESIS to be used in a manned environment.

When deploying a technology like GENESIS, one should consider the visual transformation of the target urban ecosystems in addition to the sound emission. On the one hand, the current prototype module, dubbed the green node, is quite small and resembles a typical photovoltaic panel. On the other hand, we believe that the visual presence of these modules raises awareness of the existence of infrastructures related to city dwellers' digital activity and their environmental impacts. From this viewpoint, it appears to be worthwhile to make the system visible, just as the infrastructures for energy distribution (electricity, gas), waste management, and so on are visible in the city.

Another critical aspect of integrating GENESIS into an urban ecosystem is the heat emitted by the system's photovoltaic panels, computing, and energy management systems. The produced heat is dissipated into the ambient air in the current small-scale prototype, with the help of a system of controlled fans when necessary. As a result, the module may appear to be a thermal polluter comparable, for example, to heat pumps installed for air conditioning. This heat can be recovered in a variety of ways. Based on the model of the digital heaters promoted by the Qarnot Company [24], we can imagine using the heat dissipated by GENESIS to contribute to building heating, assuming the evolution of GENESIS's technical infrastructure. Another avenue to pursue is the incorporation of GENESIS modules into urban agriculture structures installed on the roofs of city buildings. In this case, the heat dissipated by GENESIS could be directly used for crop cultivation.

⁶<https://qarnot.com/> (Accessed March 2022).

7. Conclusion and perspectives

We presented a model-based framework for the long-term sustainability of urban systems. In our solution, we proposed an expressive and flexible simulation tool that captures the deployment of an emerging green technology, named GENESIS, in existing urban neighborhoods. Using it, we can study typical deployment conditions under which urban IT equipment can be powered by renewable energy. Our study analyzed the deployment of GENESIS in two cities in France, Nantes and Montpellier, as a case study for the relevance of our proposal.

This work contributes to the emergence of more sustainable urban systems by addressing the issues raised by the rapid development of city dwellers' digital needs. In theory, a digital technology combining edge computing and smart grid paradigms appears to be relevant. Nevertheless, as with any innovation, the adoption of GENESIS necessitates a change in dwellers' habits. Our work has demonstrated the technology's viability and the importance of incorporating it into existing urban systems.

Future studies will include expanding our simulation tool to capture more detailed deployment scenarios, such as the effects of solar masking between buildings in an urban context, more advanced distribution models of IT equipment among houses, and equipment usage to reflect actual situations, among other things. In urban environments, additional GENESIS exploitation modes, such as PaaS cloud computing infrastructure, can be investigated. This comprises an economic analysis of such investment profitability, as in [9].

References

- [1] Hany F. Atlam, Robert J. Walters, and Gary B. Wills. Fog computing and the internet of things: A review. *Big Data and Cognitive Computing*, 2(2), 2018.
- [2] Simon Elias Bibri. On the sustainability and unsustainability of smart and smarter urbanism and related big data technology, analytics, and application. *Big Data Science and Analytics for Smart Sustainable Urbanism*, 2019.
- [3] Simon Elias Bibri and John Krogstie. ICT of the new wave of computing for sustainable urban forms: Their big data and context-aware augmented typologies and design concepts. *Sustainable Cities and Society*, 32:449–474, 2017.
- [4] Aldo Bischì, Leonardo Taccari, Emanuele Martelli, Edoardo Amaldi, Giampaolo Manzolini, Paolo Silva, Stefano Campanari, and Ennio Macchi. A rolling-horizon optimization algorithm for the long term operational scheduling of cogeneration systems. *Energy*, 184:73–90, 2019.
- [5] Mainak Chakraborty and Ajit Pratap Kundan. Grafana. In *Monitoring Cloud-Native Applications*, pages 187–240. Springer, 2021.
- [6] Tudor Cioara, Ionut Anghel, Marcel Antal, Sebastian Crisan, and Ioan Salomie. Data center optimization methodology to maximize the usage of locally produced renewable energy. In *2015 Sustainable Internet and ICT for Sustainability (SustainIT)*, pages 1–8. Institute of Electrical and Electronics Engineers (IEEE), 2015.
- [7] V. Costanzo, R. Yao, E. Essah, L. Shao, M. Shahrestani, A.C. Oliveira, M. Araz, A. Hepbasli, and E. Biyik. A method of strategic evaluation of energy performance of building integrated photovoltaic in the urban context. *Journal of Cleaner Production*, 184:82–91, 2018.
- [8] Patricia Crouette and Victor Prieur. Baromètre du numérique : Enquête sur la diffusion des technologies de l'information et de la communication dans la société française en 2019. Technical report, Conseil Général de l'Economie, de l'Industrie, de l'Energie et des Technologies (CGE), Autorité de Régulation des Communications Electroniques et des Postes (ARCEP) et Agence du numérique, 2020. DOI: 10.5281/zenodo.4332292, Tech. report (in French), <https://www.credoc.fr/publications/barometre-du-numerique-2019-1>.
- [9] Federica Cucchiella, Idiano D'Adamo, and S.C. Lenny Koh. Environmental and economic analysis of building integrated photovoltaic systems in italian regions. *Journal of Cleaner Production*, 98:241–252, 2015. Special Volume: Support your future today! Turn environmental challenges into opportunities.
- [10] Shuiguang Deng, Hailiang Zhao, Weijia Fang, Jianwei Yin, Schahram Dustdar, and Albert Y. Zomaya. Edge intelligence: The confluence of edge computing and artificial intelligence. *IEEE Internet of Things Journal*, 7(8):7457–7469, 2020.
- [11] Elsa Estevez, Nuno Vasco Lopes, and Tomasz Janowski. Smart sustainable cities : Reconnaissance study. 2016.
- [12] Pietro Florio, Giuseppe Peronato, A.T.D. Perera, Anthony Di Blasi, Kin Ho Poon, and Jérôme H. Kämpf. Designing and assessing solar energy neighborhoods from visual impact. *Sustainable Cities and Society*, 71:102959, aug 2021.
- [13] S. Freitas, C. Catita, P. Redweik, and M.C. Brito. Modelling solar potential in the urban environment: State-of-the-art review. *Renewable and Sustainable Energy Reviews*, 41:915–931, jan 2015.
- [14] Sean Gillies, Mike Taves, Joshua Arnott, Oliver Tonnhofer, Joris Van den Bossche, Jacob Wasserman, Aron Bierbaum, Tom Caruso, Allan Adair, Johannes Schönberger, Hannes, Thibault Deutsch, George Ouzounoudis, Ewsterrenburg, Phil Elson, Pierre PACI, Chad Hawkins, Howard Butler, Jeremiah England, Kevin Wurster, Gabi Davar, Dave Collins, Denis Rykov, Hugo van Kemenade, James Gaboardi, Juan Luis Cano Rodríguez, Konstantin Veretennicov, Bart Broere, Bas Couwenberg, and Brad Hards. Toblerity/Shapely: Shapely 1.8.0, 2021. DOI: 10.5281/zenodo.5597139, publisher: Zenodo, <https://doi.org/10.5281/zenodo.5597139>.
- [15] Daniele Groppi, Livio de Santoli, Fabrizio Cumo, and Davide Astiaso Garcia. A gis-based model to assess buildings energy consumption and usable solar energy potential in urban areas. *Sustainable Cities and Society*, 40:546–558, 2018.
- [16] Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, sep 2020.
- [17] Ibrahim Abaker Targio Hashem, Victor Chang, Nor Badrul Anuar, Kayode Adewole, Ibrar Yaqoob, Abdullah Gani, Ejaz Ahmed, and Haruna Chiroma. The role of big data in smart city. *International Journal of Information Management*, 36(5):748–758, 2016.
- [18] Jaroslav Hofierka and Ján Kaňuk. Assessment of photovoltaic potential in urban areas using open-source solar radiation tools. *Renewable Energy*, 34(10):2206–2214, oct 2009.
- [19] IEA. Potential for building integrated photovoltaics. Technical report, 2002. https://iea-pvps.org/wp-content/uploads/2020/01/rep7_04.pdf.
- [20] Kelsey Jordahl, Joris Van den Bossche, Martin Fleischmann, James McBride, Jacob Wasserman, Adrian Garcia Badaracco, Jeffrey Gerard, Alan D. Snow, Jeff Tratner, Matthew Perry, Carson Farmer, Geir Arne Hjelle, Micah Cochran, Sean Gillies, Lucas Culbertson, Matt Bartos, Brendan Ward, Giacomo Caria, Mike Taves, Nick Eubank, Sangarshanan, John Flavin, Matt Richards, Sergio Rey, Maxalbert, Aleksey Bilogur, Christopher Ren, Dani Arribas-Bel, Daniel Mesejo-León, and Leah Wasser. geopandas/geopandas: v0.10.2, 2021. DOI: 10.5281/zenodo.5573592, publisher: Zenodo, <https://doi.org/10.5281/zenodo.5573592>.
- [21] Fanxin Kong and Xue Liu. A survey on green-energy-aware power management for datacenters. *ACM Computing Surveys*, 47(2):1–38, 2015.

- [22] Thomas Leduc. t4gpd, 2021. DOI: 10.5281/zenodo.5771916, publisher: Zenodo, <https://zenodo.org/record/5782942>.
- [23] Nahid Mohajeri, Govinda Upadhyay, Agust Gudmundsson, Dan As-souline, Jérôme Kämpf, and Jean-Louis Scartezzini. Effects of urban compactness on solar energy potential. *Renewable Energy*, 93:469–482, aug 2016.
- [24] Yanik Ngoko, Nicolas Saintherant, Christophe Cerin, and Denis Trystram. Invited paper: How future buildings could redefine distributed computing. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 1232–1240, 2018.
- [25] A.T.D. Perera, Wahid M Nik, Dasaraden Mauree, and Jean-louis Scartezzini. An integrated approach to design site specific distributed electrical hubs combining optimization, multi-criterion assessment and decision making. *Energy*, 134:103–120, sep 2017.
- [26] Charith Perera, Yongrui Qin, Julio C. Estrella, Stephan Reiff-Marganiec, and Athanasios V. Vasilakos. Fog computing for sustainable smart cities: A survey. *ACM Comput. Surv.*, 50(3), jun 2017.
- [27] Brian Ramprasad, Alexandre da Silva Veith, Moshe Gabel, and Eyal de Lara. Sustainable computing on the edge: A system dynamics perspective. In *Proceedings of the 22nd International Workshop on Mobile Computing Systems and Applications, HotMobile '21*, page 64–70, New York, NY, USA, 2021. Association for Computing Machinery.
- [28] Jeff Reback, Jbrockmendel, Wes McKinney, Joris Van den Bossche, Tom Augspurger, Phillip Cloud, Simon Hawkins, Matthew Roeschke, Gfyoung, Sinhrks, Adam Klein, Patrick Hoefler, Terji Petersen, Jeff Tratner, Chang She, William Ayd, Shahar Naveh, JHM Darbyshire, Marc Garcia, Richard Shadrach, Jeremy Schendel, Andy Hayden, Daniel Saxton, Marco Edward Gorelli, Fangchen Li, Matthew Zeitlin, Vytautas Jancauskas, Ali McMaster, Pietro Battiston, and Skipper Seabold. pandas-dev/pandas: Pandas 1.4.1, 2022. DOI: 10.5281/zenodo.6053272, publisher: Zenodo, <https://doi.org/10.5281/zenodo.6053272>.
- [29] Tiago Rosado and Jorge Bernardino. An overview of openstack architecture. In *Proceedings of the 18th International Database Engineering & Applications Symposium, IDEAS '14*, page 366–367, New York, NY, USA, 2014. Association for Computing Machinery.
- [30] Gilles Sassatelli, Abdoulaye Gamatié, and Michel Robert. Data Processing System with Energy Transfer, October 2018. Patent num. WO2017178571, <https://patentscope.wipo.int/search/en/detail.jsf?docId=W02017178571>.
- [31] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *1994 First Workshop on Mobile Computing Systems and Applications*, pages 85–90, 1994.
- [32] Gerard Sierksma and Yori Zwols. *Linear and integer optimization : theory and practice*. Advances in applied mathematics. Chapman & Hall/CRC, 3rd edition edition, 2015.
- [33] Brandon Stafford. pysolar, 2020. DOI: 10.5281/zenodo.5518074, publisher: Zenodo, <https://doi.org/10.5281/zenodo.5518074>.
- [34] Yevgeniy Sverdlik. How is a Mega Data Center Different from a Massive One? <https://www.datacenterknowledge.com/archives/2014/10/15/how-is-a-mega-data-center-different-from-a-massive-one>. accessed 20 Sep. 2021, 2014.
- [35] Gabriela Viale Pereira and Luiza Schuch de Azambuja. Smart sustainable city roadmap as a tool for addressing sustainability challenges and building governance capacity. *Sustainability*, 14(1), 2022.
- [36] Alina Walch, Roberto Castello, Nahid Mohajeri, and Jean-louis Scartezzini. Big data mining for the estimation of hourly rooftop photovoltaic potential and its uncertainty. *Applied Energy*, 262(November 2019):114404, mar 2020.