



HAL
open science

Short-Scale Stochastic Solar Energy Models: A Datacenter Use Case

Sara Alouf, Alain Jean-Marie

► **To cite this version:**

Sara Alouf, Alain Jean-Marie. Short-Scale Stochastic Solar Energy Models: A Datacenter Use Case. Mathematics , 2020, Queue and Stochastic Models for Operations Research, 8 (12), pp.1–26. 10.3390/math8122127 . hal-03046453

HAL Id: hal-03046453

<https://inria.hal.science/hal-03046453v1>

Submitted on 8 Dec 2020

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 4.0 International License

Short-Scale Stochastic Solar Energy Models: a Datacenter Use Case

Sara Alouf ^{1*} and Alain Jean-Marie ^{2*}

¹ Inria, Université Côte d’Azur, France

² Inria, LIRMM, Univ Montpellier, CNRS, France

Correspondence: name.surname@inria.fr

The authors contributed equally to this work.

Author version

Abstract

Modeling the amount of solar energy received by a photovoltaic panel is an essential part of Green IT research. The specific motivation of this work is the management of the energy consumption of large datacenters. We propose a new stochastic model for the solar irradiance, featuring minute-scale variations and therefore suitable for short-term control of performances. Departing from previous models, we use a weather-oriented classification of days obtained from past observations, to parameterize the solar source. We demonstrate through extensive simulations, using real workloads, that our model outperforms the existing ones in predicting performance metrics related to energy storage.

Keywords solar energy; stochastic modelling; performance evaluation

1 Introduction

It is now a well-known fact that the energy footprint of information technologies has become a sizeable proportion of the global energy consumption, and that the operations of datacenters are responsible for a large fraction of it. As argued in [15], major companies operating huge datacenters are interested to take advantage of “green” renewable energy resources, both in the sake of saving money and participating to the CO2 emission reductions. Powering datacenters with (at least some) locally produced solar energy is an attractive possibility. On the other hand, it is also well-known that renewable energy resources, such as wind or sun, feature a large variability, which adds up to the variability of the workload. In these conditions, the management of datacenters involves much flexibility in *e.g.* scheduling of jobs and equipment, pricing of services [5, 16, 10]. The models allowing to compute these policies require accurate models of the stochastic solar energy input.

1.1 Need for accurate models of solar irradiance

In this paper, we propose a new model for stochastic sources of solar light, and advocate its use when studying the performance, static and dynamic optimization or dimensioning of datacenters.

First, as we shall demonstrate in our experiments, existing models fail to accurately predict the basic performance metrics for energy storage: battery shortage and energy spillover. We depart from the usual approach by not considering long-duration, multi-day models. Our point of view is that stationarity scarcely occurs, even within one day, and that discussing the transitions from one day to the next is unnecessary.

We base this point of view on the following use case scenario. The datacenter manager (DM) plans the resources (computing units in service, purchase of energy on the market, ...) one or several days ahead. He/she will use simulations to investigate several scenarios and optimize some criteria. A model is needed for the solar energy input that next day (also for the workload of the datacenter but this is outside the scope of this paper). As anybody in the world, the DM has access to weather info about next day. Given the financial stake, it can be assumed that the DM has the means to buy a precise weather outlook. Datacenter managers also have the means to obtain astronomical and geographical data such as the time of sunrise and sunset every day, the value of the clear-sky irradiance (see Section 2). In summary, what is needed here is a one-day stochastic model, parametrized by some prediction about the kind of weather available some time ahead.

Second, we believe that a correct way to assess the quality of a solar source model, is through a “queueing evaluation”. We mean by this that instead of focusing on the arrival process of energy with methods of the signal processing literature (as we have ourselves done in [21]), it is necessary to evaluate the quality of a source through its performance in a finite queue. We believe this is relevant for several reasons. On the one hand, as already argued, there is no clear stationarity in the processes, so that spectral analysis may not be relevant, or is difficult to interpret. On the other hand, queueing systems (especially with finite capacity) are not linear systems, so that the dependence of their response on the spectrum of a signal is not immediate to see. Finally, many queueing models of datacenter consumption and optimization are being developed in the literature ([5, 16, 15, 7, 18, 8] to quote just a few): they need a realistic energy source model, as well as a realistic consumption model. We devote this work to the first element, the second having been taken from real (therefore realistic) traces.

1.2 Related work

Models of solar irradiance and of the transformation of irradiance into electric current, are not new in the literature. As we will argue shortly, we will not be interested in the detailed modeling of transformation and we concentrate on irradiance itself. Models in the literature differ in their purpose, and on the data they use. Some models aim at capturing the accumulated sunlight at the scale of days or weeks, which is relevant for predicting the response of vegetation. Others consider smaller time scales: days, hours, or minutes depending on the application envisioned.

To set the context for this review, we introduce first the relevant vocabulary. The *global solar irradiance* refers to the amount of solar energy reaching a specific area and is expressed in W/m^2 . It varies over the course of the day and from day to another, we will therefore denote it as $I_G(d, m)$ where d stands for “day” and m stands for “minute”. The ideal value of this irradiance would occur when no cloud or obstacle is present: this *clear sky irradiance* is denoted with $I_{CS}(d, m)$. The global solar irradiance $I_G(d, m)$ is seen as the result of applying a multiplicative “noise” to the clear sky irradiance $I_{CS}(d, m)$. This corresponds to the decomposition:

$$I_G(d, m) = I_{CS}(d, m) \times \alpha(d, m) \tag{1}$$

where we have used the notation $\alpha(d, m)$ to refer to the *clear sky index* which is how the multiplicative noise is known in the literature.

The literature features many models for the (deterministic) clear sky irradiance function $I_{CS}(d, m)$, as surveyed by Dave, Halpern and Myers in [6] or by Bird and Hulstrom in [3]. These authors propose themselves a model in [4], which we later refer to as “Bird’s model”, which involves many physical parameters and consequently require the knowledge or measurements of many quantities. Other models such as [20] are simpler to use.

The stochastic models which aim at representing the variations of I_{CS} with respect to its ideal value, use two different approaches. One direct statistical approach is to consider the function I_G as a time series and characterize its distribution. This is the approach of [25], in which the non-stationary process $I_G(t)$ is classically decomposed as

$$I_G(t) = m(t) + s(t) + Y(t)$$

where $m(t)$ is a slowly varying trend, $s(t)$ the seasonal (here, daily) component, and $Y(t)$ some ARMA process. We do not believe the stochastic component is likely to have the same characteristics throughout a year, so that some month-dependent is likely to be better. It is also dubious that such a model would spontaneously generate zero values at night, as it should. In addition, the fitting of one ARMA model on real data involves, according to [25], some exploration of the combinations of parameters, and some data-dependent tuning of learning parameters. We have chosen not to try this model in a first approach. The direct generation of $I_G(t)$ is also the topic of Miozzo et al. in [17], who base their model on semi-Markov processes. We shall present this model in Section 2 and compare it with our proposal.

An alternative approach is to use the decomposition (1) and find a model for the random process $\alpha(d, m)$. There is a body of literature devoted to the characterization of the distribution of α at a given time. Examples are Suehrcke and McCormick [24] and Jurado et al. in [13], who argue that this distribution can be seen as a mixture of two normal distributions. As reported by Ghiassi et al. in [9], the time-dependent process is then modeled as a Levy process, but this modeling is dismissed there as generated trajectories are very different in shape from actual ones. We shall not consider this model further in our study. Statistics on the temporal behavior of the process α are obtained by Gu et al. in [11]: they demonstrate that several temporal scales in the correlations of this process at the minute level. Despite this global understanding of the phenomenon, no model seems to emerge for generating synthetic traces of it, at the minute level. In [21], we have proposed such a

model as a semi-Markov process with a finite number of states, to which some random noise is added. Although promising, this approach lacked the dependency on seasons and weather, and also produced trajectories with excessive variations.

Finally, note that the question to decide whether a solar irradiance model is “good” or not is not entirely solved, as discussed in [12]. We consider that the use-case is what decides: a given model may prove accurate in situations and not in others. We favor a validation through the response of a storage system.

As conclusion of this review, we argue that there is still a need for a) models at the scale of the minute, and b) convincing validation of the accuracy of these models in scenarios involving datacenters.

Contributions The main contributions of this paper are:

- A forecast-based solar energy model that captures the small-scale fluctuations that the global irradiance exhibits. As will be shown in the paper, our model can be used to generate trajectories that closely mimic measured data, a feature unseen in state-of-the-art models.
- A practical use case scenario: we simulate a datacenter fed by solar energy through photovoltaic panels. The captured solar energy is stored in batteries. Through simulations we show that our energy model achieves the same performance as benchmarks: the distribution of the battery level is sensibly the same. None of the state-of-the-art models that we tested achieves as matching results as our model.

1.3 Organization of the paper

The rest of the paper is structured as follows. Section 2 presents the problem of modeling the solar energy and its challenges, and reviews some of the models found in the literature. Section 3 introduces our model of a datacenter fed by a dual energy supply, solar energy and grid, and discusses our main assumptions on it. Section 4 details our methodology to construct an energy source and presents the energy consumers that we will use in the experiments. Section 5 presents the setup of our experiments and the results that we obtained. Section 6 concludes the paper and presents some perspectives.

2 Solar Energy Models

In this section, we present the models we will use as comparison in the experimental evaluation. We begin with a discussion on the challenges of modeling the clear sky index $\alpha(\cdot)$.

2.1 Challenges

To illustrate the challenges faced in the process of modelling the solar energy, we depict in Figures 1 and 2 the global irradiance on selected days as measured by [26] and [2] in the cities of Portland (Oregon, US) and Golden (Colorado, US) respectively. Each row of Figures 1 and 2 displays the data relative to a specific day, but one year apart. Each subfigure shows the global and clear sky irradiance at the left and the clear sky index at the right. We can make several observations: First, each day has its own clear sky irradiance curve, that depends on astronomical data like sunset and sunrise but also on total ozone, surface pressure, ground albedo, to name a few of the impacting factors. Second, very large daily fluctuations can be observed in the global irradiance, even for short intervals as seen for instance in Figures 1d and 2a. Third, none of the displayed data appear to be stationary, not even the clear sky index.

2.2 Models for the clear sky irradiance

As mentioned in Section 1.2, the literature features many models of $I_{CS}(\cdot)$, with varying degrees of complexity. Whatever the model, it can be assumed that datacenter managers have the financial and technical means to obtain this function at the location of their facility.

It is more difficult for us to do so, so that we will rely on the data publicly available that features the computation of Bird’s model in addition to measurements. In addition, for some of our simulations, we will use a simplified sinusoidal model: we assume that the curve $I_{CS}(d, m)$ of every day d is a pure sine function

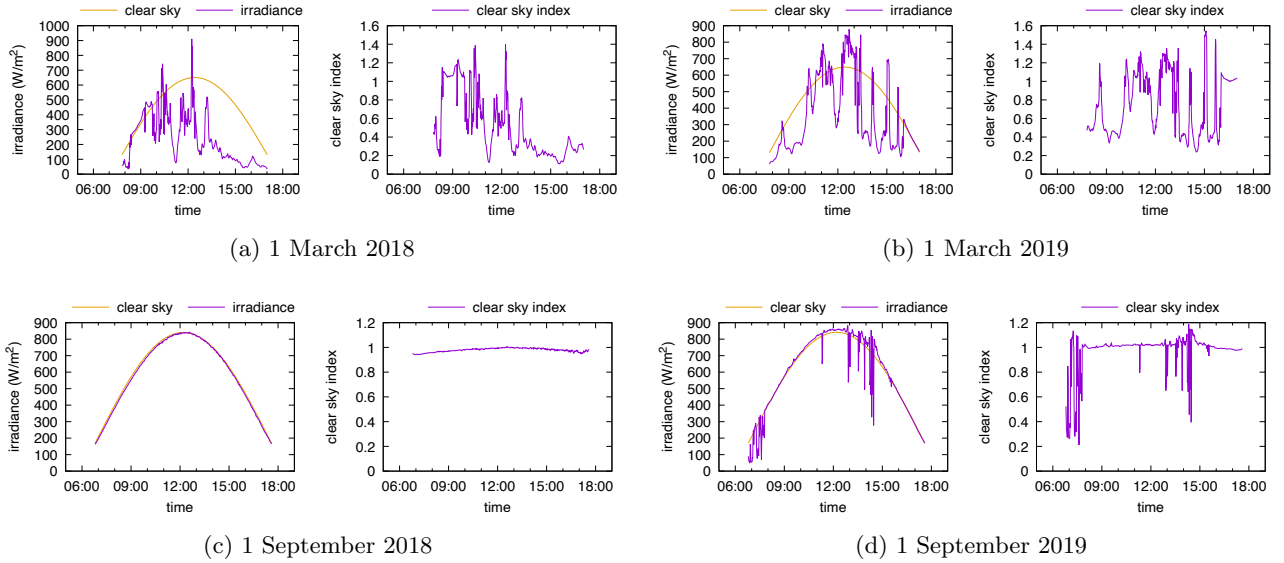


Figure 1: Global and clear sky irradiance (left) and clear sky index (right), in Portland for four selected days.

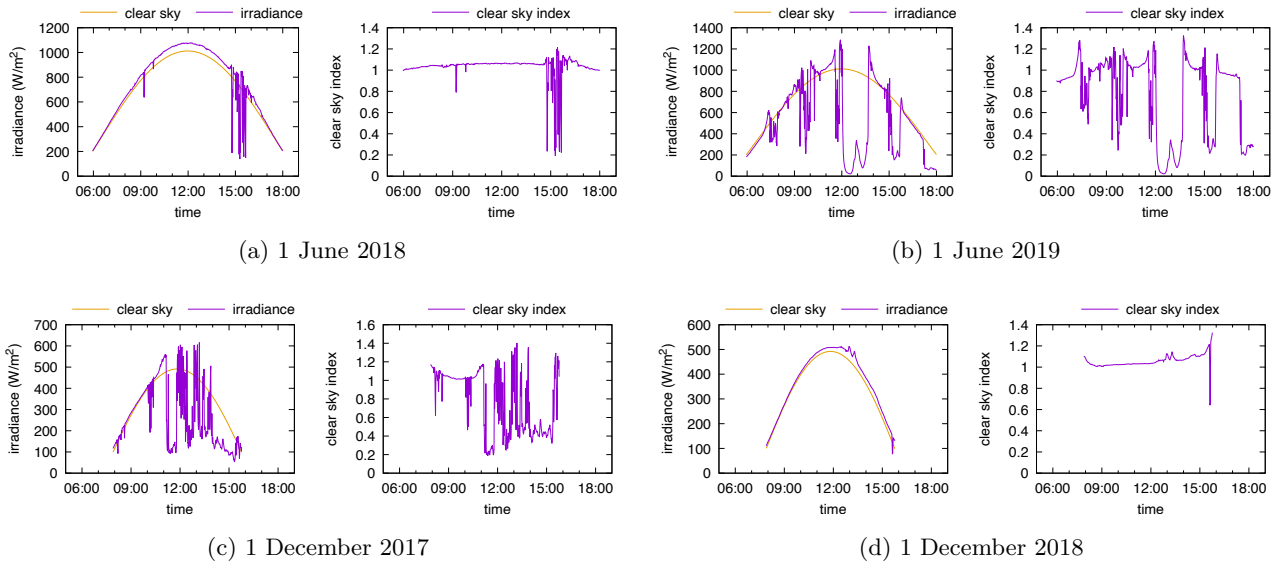


Figure 2: Global and clear sky irradiance (left) and clear sky index (right), in Golden for four selected days.

characterized by two values: the length of the day $\lambda(d)$ and the maximal irradiance I_M . Assuming in addition that the maximum is reached at noon, these values are sufficient to define a unique function with equation:

$$I_{CS}(d, m) = I_M \times \frac{\cos(2\pi/T \times (m - T/2)) - \cos(\pi\lambda(d)/T)}{1 - \cos(\pi\lambda(d)/T)} \quad (2)$$

where T is the full day length. The values of $\lambda(d)$ and I_M for a given location can be retrieved from *e.g.* [22]. If required, the model can be easily modified to account for a sunrise time t_r , a sunset time t_s and a maximum halfway between them.

2.3 Models for the direct irradiance

Miozzo et al. in [17] propose two models for the function $I_G(d, m)$. Both models assume that the process $I_G(t)$ is a semi-Markov process.

The simpler model, which we denote as the ‘Daily’ (or ‘Day/Night’) model in what follows, consists in separating day and night, and modeling the irradiance during the day as a random variate drawn from a month-dependent distribution. In our notation,

$$I_G(d, m) = \zeta_{\mu(d)}(d) \delta(d, m)$$

where $\mu(d) \in [1..12]$ denotes the month to which day d belongs, $\{\zeta_{\mu}(d); d = 1..31\}_{\mu=1..12}$ are twelve independent families of i.i.d. random variables, and $\delta(d, m)$ an on/off renewal process with values in $\{0, 1\}$.

In the second model, which we denote as the ‘Hourly’ model in what follows, the irradiance depends on the hour in addition to the month. In our notation,

$$I_G(d, m) = \zeta_{\mu(d), h(m)}(d)$$

where $h(m) = m \bmod 60$, $\{\zeta_{\mu, h}(d); d = 1..31\}_{\mu=1..12, h=0..23}$ are 12×24 independent sequences of i.i.d random variables. In this model, the semi-Markov process has 24 states, the sojourn time in every state is deterministic and equal to one hour, and transitions are deterministic as well.

In the experiments, we have replaced the (semi-Markov) on/off process of the Daily model, with a deterministic day/night alternation based on the astronomical sunrise and sunset times. Since this is a readily available data (e.g. [22]), it seems unnecessary to try to guess it statistically.

2.4 Models for the clear sky index

In [20], the function I_{CS} is constructed using an astronomical model, and the function $\alpha(\cdot)$ is assumed to be constant throughout the day, with a value that depends on a *cloudiness* value N through the equation:

$$\alpha(d, m) = 1 - \frac{3}{4} \left(\frac{N}{8} \right)^{3.4}. \quad (3)$$

The value of N is measured in oktas and lies between 0 (no clouds) and 8 (8 eighths of the sky covered). The average cloudiness for a given location is commonly measured by weather stations and averaged over one month. In our notation, we would then write $N_{\mu(d)}$ to reflect this dependence on the current month. References [28] and [27] provide some (partial) data for Portland, Oregon and Denver, Colorado, for 2001. Since this data is incomplete, we have interpolated it for the missing months. The data is displayed in Table 1.

Table 1: Cloudiness measure for the two locations used in experiments

Location	Year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Portland	5.7	6.7	6.6	6.4	6.2	5.8	5.4	3.8	4.0	4.3	5.6	6.4	6.8
Denver	–	–	5.0	5.3	7.2	5.6	2.5	–	2.5	–	–	–	2.5
— completed	3.8	5.0	5.0	5.3	7.2	5.6	2.5	2.5	2.5	2.5	2.5	2.5	2.5

3 Datacenter Model

We model the datacenter (DC) with a simple flow and storage model illustrated in Figure 3. The energy supplied to the DC comes from two sources: the solar energy, via photovoltaic panels, and the grid. The solar energy feeds a battery storage. When the battery is full, the energy is lost. The battery has priority over the grid: if it is not empty, it is used by the energy consumer. If it is empty, the energy is taken from the grid.




Figure 3: Principal features of the datacenter model.

The energy consumed by the DC is due to the functioning of computing units, called “nodes” in the remainder. This consumption depends in turn on the number of nodes that are actually busy executing jobs. The number of jobs running in the DC at a given time, which we call the *workload*, is assumed to be exogenous. In other words, we do not model the complex process of job submitting, queueing, scheduling, abandoning etc. as observed in [23]. In our experiments, we will use traces of real datacenters from which we extract this number of running jobs at every instant. Those will be described in Section 4.2.

Since our focus is on modeling the solar irradiance process, we shall not develop complex models of: a) transformation of the irradiance into electric current; b) fine points of datacenter energy consumption: cooling, starting/stopping machines etc.

For a) we use a simple linear model, assuming that the electric power coming in the DC is directly proportional to solar irradiance. It is well-known that photovoltaic (PV) panels only capture a proportion of the irradiance and that this proportion depends, in particular, on temperature. Also, the current that comes out of the PV suffers losses before it can be used by the DC. Since in our assumption, the proportion is fixed, we will not mention its value explicitly and present all panels surface in *efficient* square meters: the surface necessary to produce an electric power equal to the unit irradiance, measured in W/m^2 .

For b) we use a simple affine model, assuming that nodes have a consumption P_{idle} when idle, and P_{busy} when busy. Accordingly, given that the DC has N_{nodes} computing nodes, a workload of N_{jobs} is converted into power with the following equation:

$$P = P_{idle} \times (N_{nodes} - N_{jobs}) + P_{busy} \times N_{jobs}. \quad (4)$$

Obviously, a precise modeling of these elements is mandatory for building a useful model of a complete datacenter. However, we are not interested in accurate modeling of the datacenter, but only of the irradiance. We make the working assumption that, albeit imprecise, the energy transformation and consumption models we use are sufficiently representative of reality to validate the irradiance model.

4 Methodology

In this section, we discuss the processing of real data that we have performed. In Section 4.1, we present our new solar source model and the way its parameters are learned. In Section 4.2, we describe the workload data and how we used it.

4.1 Forecast-based sources

Our driving idea for constructing a solar source will be that the curve $I_G(\cdot)$ essentially depends on the *type* of the day, referring to the sort of weather this day enjoys. Mathematically, we refine the model in (1) into:

$$I_G(d, m) = I_{CS}(d, m) \times \alpha(\tau(d), m). \quad (5)$$

The notation $\tau(d)$ refers to the *type* of the day which defines the weather conditions observed/predicted along that day. We assume $\tau(d) \in \mathcal{T}$ where \mathcal{T} is the set of day types, finite and hopefully not too large. In the experiments of Section 5 we will consider models where the number of day types ranges from 2 to 16.

For any $\tau \in \mathcal{T}$, we assume that there is a sequence of distributions $\xi^{(\tau, m)}$, for $m = 0, 1, \dots, 1439$ such that the clear sky index $\alpha(\tau, m)$ is a random variable following the distribution $\xi^{(\tau, m)}$.

For a given day d of type $\tau = \tau(d)$, it is then possible to characterize the global irradiance $I_G(d, m)$ using a model for $I_{CS}(d, m)$ and i.i.d. sequences $\{\xi^{(\tau, m)}(n), n = 0, 1, \dots, 1439\}$, drawn from distributions $\xi^{(\tau, m)}$, with:

$$I_G(d, m) = I_{CS}(d, m) \times \xi^{(\tau(d), m)}(m). \quad (6)$$

In order to classify the days in the type set \mathcal{T} , we start from per-minute irradiance data for two locations, Portland (Oregon, US; see [26]) and Golden (Colorado, US; see [2]) for a 3-year period of time. In addition to the global irradiance, the data [2, 26] reports predictions of the clear sky irradiance using Bird’s model [4]. We rely on these predictions of $I_{CS}(d, m)$ to compute the clear sky index $\alpha(\tau(d), m)$ according to (5). The subsequent step is to identify which *features* of days will be used for the classification.


The type of the day can be seen as a description of the weather forecast for that day. In its simplest expression, a day can be seen as “sunny” or “cloudy”, but much more detailed descriptions can be envisaged: hourly conditions may be available/predicted and a finer categorization of the sky condition may be made. When identifying the day type, one can rely on the average value of the clear sky index over the entire day, or on the average values obtained over a number of intervals partitioning the day duration. Partitioning the whole duration of the day in fixed intervals throughout the year is not adequate, since the quantity $\alpha(d, m)$ is measured only during daylight time. We will therefore consider that daylight time is split in a number I of equally sized intervals, seen as “phases” of the day, during which the distribution of $\alpha(d, m)$ is assumed to be constant. This amounts to refine the model in (6) into:

$$I_G(d, m) = I_{CS}(d, m) \times \xi^{(\tau(d), \varphi(d, m))}(m), \quad (7)$$


where m is a minute during daylight time, $\varphi(d, m) \in [1..I]$ is the interval number of minute m in day d , and $\{\xi^{(\tau, \varphi)}(m), m = 0..1439\}_{\tau \in \mathcal{T}, \varphi = 1..I}$ are $|\mathcal{T}| \times I$ independent sequences of i.i.d. random variables. We will limit the value of I to 12 in our experiments. Consequently the shortest days in the year in Portland with ~ 433 minutes would have eleven 36-minute long intervals and one 37-minute long interval. In Golden, with ~ 461 minutes, the shortest days in the year would have seven 38-minute long intervals and five 39-minute long ones. (Slight variations exist over the years, these figures are those seen in the training and testing data.)

To estimate day types, we use the Scikit-learn Python module [19] and more specifically KMeans, its implementation of the K-Means clustering. We split the collected data into two sets: the first one (the *training* data set) spans a period of two years, from October 2017 until September 2019, while the second one (the *testing* data set) spans a period of a full year from October 2019 until September 2020. We cluster the training data (that in the first set) using a number of features I ranging from 1 until 12 (each feature is the average clear sky index over a time interval), and we consider a number of clusters $K = |\mathcal{T}|$ ranging from 2 until 16. We therefore obtain 180 distinct clustering results, one for each configuration. For each clustering result, we predict the types of the days present in the second set, the testing data set, based on the clusters identified in the training set. To do so, we simply select the cluster which centroid is closest to the vector of features attached to the day to be classified. To test the models presented in Section 2 (and in particular (5)), we will generate synthetic sources of energy that will mimic the days present in the testing data set. The interval where the current minute lies is determined using sunrise/sunset data and the splitting of daylight time into equal intervals. Using the day type and the interval number, we generate samples from the distributions $\zeta^{(\tau, \varphi)}$, one sample per minute, and obtain the global irradiance with (7).

We illustrate in Figures 4 and 5 the types obtained for selected numbers of clusters (in other words, selected sizes of \mathcal{T}) and selected number of intervals. Each graphic in these figures reports the centroids of the clusters obtained. The numbering of the clusters changes at each run, but by looking at the centroids obtained, it is often possible to map the clusters obtained with different number of intervals. For instance, the cluster (type) numbered 3 in Figure 4a has a centroid with coordinates close to 1, which represents days where the clear sky index is close to 1 all day long. Such clear sky days are grouped in cluster 0 in Figure 4b whose centroid has coordinates close to 1. Similarly, type 0 in Figure 5a can be mapped to type 4 in Figure 5b and represents days with heavy clouds.




(a) Portland: 8 types and 4 intervals




(b) Portland: 8 types and 12 intervals

Figure 4: Types of days for Portland.



(a) Golden: 9 types and 4 intervals



(b) Golden: 9 types and 12 intervals

Figure 5: Types of days for Golden.

As for the distributions $\xi^{\tau, \varphi}$, $\tau \in \mathcal{T}$, $\varphi = 1..I$, we use the empirical distributions of the clear sky index of the days in cluster τ in the training set and view them as the *ground truth* for the day type τ . For illustration purposes, we display in Figure 6 the distributions $\zeta^{(\tau, \varphi)}$ when there are $K = 4$ types of days and the day is partitioned into $I = 3$ intervals (this implies that there are actually 3 distinct distributions for each day type).




Figure 6: Distributions of the clear sky index for Golden with 4 day types and 3 day intervals.

Once clusters have been identified, we predict the types of the days present in the testing set (which will be used in the experiments). We illustrate in Figure 7 the predictions obtained for Portland when there are three day types (clusters) and the day is split into six intervals.

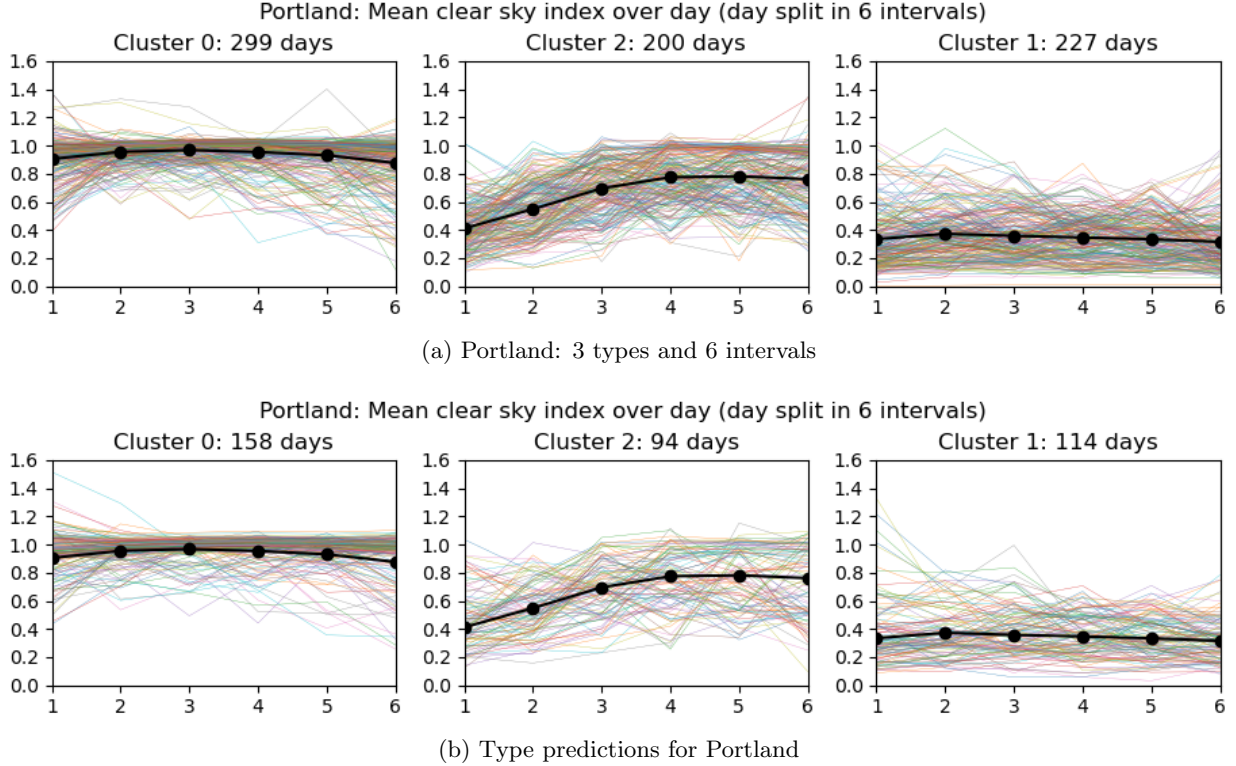


Figure 7: Days characteristics and their types for Portland: training data on top, testing data on bottom.

In the practical use of the model, the data of next day is obviously not known exactly. A proxy for it is given by weather predictions. Assuming that some hourly prediction of cloudiness is available from the weather service, the representative values for each interval can be determined, to serve as the basis for the day type prediction. If hourly predictions are not available, a coarser information can be used in conjunction with a model with fewer intervals.

4.2 Energy consumers

The solar energy collected by photovoltaic sources and stored in batteries is used to power a computing cluster. To simulate the energy consumption of a computing cluster, we rely on traces of two real clusters. The first cluster, called NEF [1], is used internally at Inria for research purposes. The workload trace that we have covers the workload experienced over a period of two years (from July 2016 until June 2018). (This trace was captured by Inria’s IT staff in July 2018 and we got exclusive rights to use it after screening by Inria’s General Data Protection Regulation officer.) The second cluster is part of a larger Google computing cluster, used mostly internally by Google employees but it runs also jobs for clients. The detailed workload description can be found in [23]. This trace covers a period of 29 days in May 2011. For each of the clusters, we first computed the number of resources/machines used over time, then integrated that number every minute (since our solar model has the minute as unit of time). The energy consumption per minute is a function of this integrated workload, see (4). Figure 8 depicts the workload of NEF, and Figure 9 depicts that of the Google cluster.

To suit our experiments, we had to pre-process these traces in order to give them the same order of magnitude. To begin with, the Google trace clearly fluctuates between around 11900 and around 12500 resources. This fluctuation being small relatively to the total, the datacenter might be seen as having an almost constant consumption. In order to give more variability to the trace, we have subtracted the baseline value 11900. As a result, the workload fluctuates between 0 and 600 approximately. In comparison, the NEF trace fluctuates between 0 and 2000 most of the time.




Figure 8: Workload (number of used resources) of NEF over time.




Figure 9: Workload (number of used machines) of the Google cluster over the days (May 2011).

The global characteristics of these traces are summarized in Table 2. Observe that the unit used is relevant for the minute-based model.

Table 2: Characteristics of the workloads used in simulation

Workload	duration	average number of resource \times second in a minute	average load
NEF	730 days	68406	1140
Google	29 days	25765	429

5 Experiments

We now present the experimental setup and the result of these experiments.

5.1 Synthetic solar traces

In our experiments, we have used the models of solar sources described in Section 2.

The sources we call “Sinusoidal” are defined by the clear-sky irradiance model (2) with the clear-sky irradiance given by (2). We have used astronomical data and recorded weather data to define their parameters. These sources generate a deterministic trajectory.

The models described in Section 2.3 are denoted “Daily” and “Hourly”. To find their parameters, we have used the training data to form the empirical distributions of monthly daylight irradiance, or monthly hourly irradiance. The generation of one trajectory consists in drawing one sample from these distributions every day (for the Daily sources) or every hour (for the Hourly sources).

For our new forecast-based source defined in Section 4.1, which we call “Synthetic”, we have used the day type predictions of the testing data and the distributions $\xi^{(\tau, \varphi)}$ for $\tau \in \mathcal{T}$ and $I = 1..I$, for generating trajectories for given days of the year.

5.1.1 Trajectories

One first step with validating the model is to generate synthetic traces and compare them with reality. Figures 10 and 11 depict the daily global irradiance of several days together with synthetic data for these days. We have not represented trajectories of Daily and Hourly sources since their piecewise-constant shape makes them look very different from real traces. Likewise, the Sinusoidal curves may look similar to the real one on a clear day, but the maximal irradiance usually differs because of the averaging of the cloudiness over one month.

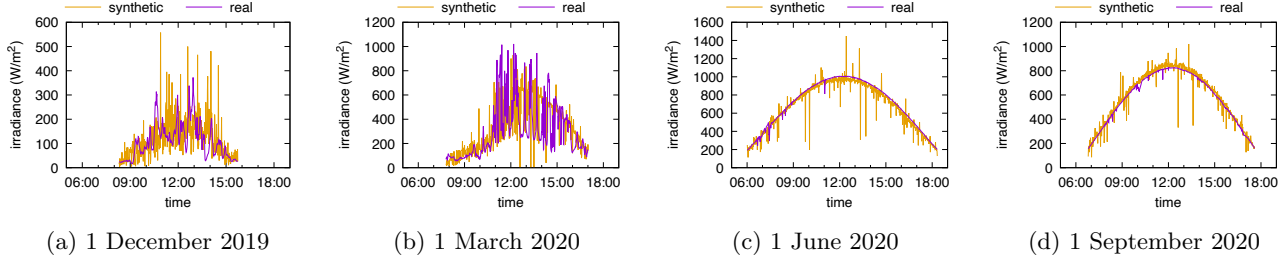


Figure 10: Global irradiance in Portland: measured and synthetic (16 types and 12 intervals).

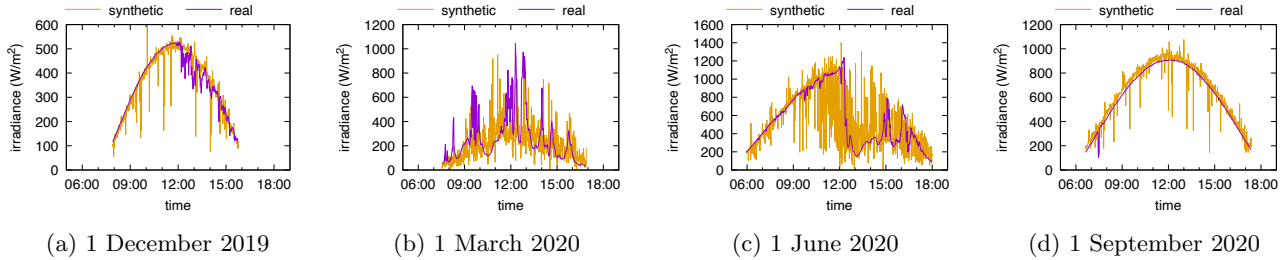


Figure 11: Global irradiance in Golden: measured and synthetic (16 types and 12 intervals).

From Figures 10 and 11 we conclude that the clustering apparently identifies correctly some trends in the weather: perfect days on 1 June 2020 and 1 September 2020 in Portland and 1 September 2020 in Golden; perfect mornings then a degradation on 1 December 2019 and 1 June 2020 in Golden. On the other hand, we also see the effect of the irradiance distributions attached to clusters, coupled with the re-sampling of irradiances at each minute: even on perfect days, the synthetic irradiance is irregular. These irregularities may be smoothed out by the buffering in the battery, though.

5.1.2 Global characteristics

In a second step, we examine the global, longer-term characteristics of the sources. Consistent with our premises that stationarity is difficult to exhibit in the irradiance phenomenon, we do not favor the spectral analysis of the signal obtained since this is a feature of wide-sense stationary signals. Instead, we settle for averages of the power over given periods. The average yearly power generated by the sources we use, are given in Table 3.

Table 3: Power of the different solar sources used in experiments (W/m^2)

Source	Portland, Oregon	Golden, Colorado
Benchmark	162.87	191.39
Pure Sinusoidal	252.51	269.18
Damped Sinusoidal	157.27	215.43
Monthly Night/Day	168.05	193.43
Monthly and Hourly	168.19	193.40
Synthetic	161.104 ± 0.027	188.380 ± 0.052

For the benchmark situation, this average is computed over the testing year. The Synthetic source selected is the one with $K = 12$ clusters and $I = 4$ intervals. The value displayed is the empirical average over 36 samples,

together with the half-width of the 95% confidence interval. The source called 'Pure Sinusoidal' is the one given by (2). The one called 'Damped Sinusoidal' is considering the clear-sky index of (3) in addition.

Tables 4 and 5 display the average power of different sources computed over each month. For sinusoidal sources, the cloudiness values have been taken from Table 1 (line 'Denver, completed'). Discrepancies between averaged yearly powers may be explained by small variations in the number of days in the year (leap days) or days missing in the records due to malfunction of the recording equipment.

Table 4: Yearly and monthly average powers for Portland, Oregon (W/m²)

source	year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
real, training (day)	327	140	174	296	322	396	488	485	434	355	268	156	130
real, testing (day)	329	131	237	312	385	384	389	492	482	302	283	179	116
real, training	162	54	76	147	180	243	313	304	251	184	119	62	48
real, testing	163	51	103	155	216	236	249	308	279	156	129	72	43
Daily, training	168	69	76	148	181	244	314	305	253	185	123	63	49
Daily, testing	170	71	104	156	217	237	251	310	280	158	130	72	43
Hourly, training	168	74	76	148	181	243	314	305	252	185	123	63	49
Hourly, testing	169	76	104	156	216	236	250	308	279	157	130	72	43
Sinusoidal w/ weather	157	67	100	155	219	285	326	368	316	240	144	81	55

Table 5: Yearly and monthly average powers for Golden, Colorado (W/m²)

source	year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
real, training (day)	379	262	324	414	414	382	444	475	437	436	317	267	243
real, testing (day)	387	268	313	387	442	475	466	445	433	412	348	270	241
real, training	187	106	144	205	228	228	275	289	248	225	146	110	94
real, testing	191	108	139	192	244	284	288	270	245	212	161	112	94
Daily, training	193	111	145	207	230	230	277	290	249	226	147	111	95
Daily, testing	198	116	140	194	245	285	290	272	247	213	161	113	95
Hourly, training	193	117	145	206	229	228	276	289	249	225	147	111	95
Hourly, testing	198	121	139	193	245	284	289	270	246	212	161	112	95
Sinusoidal w/ weather	215	118	159	213	160	300	401	389	346	279	205	147	122

The purpose of these tables is to illustrate several features of the data. First, the monthly variation is obvious, but it is worth stressing the observation that the ratio of the most sunny month over the less sunny one may be larger than 4. Second, if the variations of one year (testing) to the average of the two previous ones (training) are sometimes large in a month, although usually smoothed out for the whole year. The four first lines contain the real data, split over training/testing and with the presentation of statistics restricted to daylight hours (tagged with 'day'). This allows to assess that the irradiance does vary over the months, not just as a consequence of varying durations of daylight time over the year.

5.2 Workload

In simulations, we use two real traces that have been described in Section 4.2.

5.3 Datacenter

The characteristics of the datacenter are as follows. There are $N_{nodes} = 2300$ computing nodes, which is a number larger than the maximum number of resources used in the real traces we use (see Figure 8 and Figure 9). The power consumption of these computing nodes (see Equation 4) are taken as $P_{idle} = 52.5$ W and $P_{busy} = 112.5$ W. A rationale for this choice is that values $P_{idle} = 105$ W and $P_{busy} = 165$ W have been measured on Grid5K platform nodes in 2015 [14]). The values we choose account for an improvement in hardware consumption equivalent to roughly running twice as many nodes with the same energy.

Battery capacity in experiments will range from 1 GJ to 10 GJ, that is, less than what is necessary to store 1 day of energy input or consumption. Indeed, for a solar source of mean power 150 W/m² over a day, and a

PV panels surface of 1200 (efficient) m², the daily energy input is 15.51 GJ. The panels surface will range from 1200 m² to 12000 m².

These values are chosen so that the proportions of time the battery is empty or full are not trivially close to 0 or 1: in those situations, the complexity of the solar input model is irrelevant to predict the correct value.

5.4 Results

5.4.1 First experiments

For each city and for each energy source, we performed a first series of simulations considering either of the two consumers and various battery capacities and surfaces of the photovoltaic panels. The six settings presented in Table 6 were used.

Table 6: Settings used in the first series of experiments

Consumer	Google	Google	NEF	NEF	NEF	NEF
Battery capacity	5000 MJ	10000 MJ	5000 MJ	10000 MJ	6000 MJ	6000 MJ
Panels surface	1200 m ²	1200 m ²	1200 m ²	1200 m ²	1200 m ²	12000 m ²
Label	G5000	G10000	N5000	N10000	N1200	N12000

For each single simulation we record the empirical distribution of the battery level. We only report in the following on the values of two performance metrics, namely the percentage of time that the battery is full and the percentage of time that the battery is empty. We do not name these metrics “probabilities” because they do not result from some stationary phenomenon. When the source is random in nature, the percentage is obtained by averaging 36 independent replications. Occasionally, we display the corresponding 95% confidence interval.

Figures 12 and 13 display the performance metrics obtained for the cities of Portland and Golden, respectively, when the energy sources described in Sections 2.3, 2.4 and 4.1 are used, and for a single setting of battery capacity and panels surface. Each point in the figure represents the performance of a single energy source in terms of the full/empty battery proportion. The naming of the source is that of Section 5.1. Our forecast-based model described in Section 4.1 has 180 configurations according to the considered numbers of day types and day intervals. All configurations with the same number of day type (i.e. number of clusters) use the same point shape/color and are denoted by their number of clusters; thus there are 12 similar points with the same label, one for each possible number of day intervals.

For a given energy source model, the closer its point to that of the benchmark the better it is, as this suggests the energy model captures the essential features needed for the use case.

It is clear from the graphs displayed in Figures 12 and 13 that the “Daily”, “Hourly” and “Sinusoidal” models fail to predict the performance in terms of the percentage of time when the battery is full/empty. Therefore, they should not be used for the purpose of managing a datacenter fed by photovoltaic panels. Instead, our energy model with any of its 180 configurations yields a performance very similar to that achieved by the benchmark, making it a good candidate for this use case scenario.


The performance metrics are further discussed in Appendix A (for the percentage of time when the battery is *empty*) and Appendix B (for the percentage of time when the battery is *full*). A close look at the individual values reported in each of the graphs displayed in Figures 18–21 does not allow to identify the configuration for which our forecast-based energy source model achieves the closest performance to the benchmark’s regardless of the experimental setting. In the next section we will present our approach to select the best achieving configuration.

5.4.2 Selection of most promising configurations

To assess which, out of all 180 configurations considered in our energy model, best approaches the benchmark, we consider four selection criteria, as explained next. Each configuration has been tested in six experimental settings yielding two performance metrics. In total, we have 12 pairs of values (benchmark, model) for each configuration. We compute


- the maximum relative distance among the pairs of values (benchmark, model);
- the average relative distance over all pairs;
- the root mean square distance over all pairs;

Google experiment for Portland, battery capacity 5000, by cluster number




(a) battery capacity 5000 MJ, panels surface 1200 m²

Google experiment for Portland, battery capacity 10000, by cluster number




(b) battery capacity 10000 MJ, panels surface 1200 m²

NEF experiment for Portland, battery capacity 5000, by cluster number




(c) battery capacity 5000 MJ, panels surface 1200 m²

NEF experiment for Portland, battery capacity 10000, by cluster number




(d) battery capacity 10000 MJ, panels surface 1200 m²

NEF experiment for Portland, panel size 1200, by cluster number



(e) battery capacity 6000 MJ, panels surface 1200 m²

NEF experiment for Portland, panel size 12000, by cluster number




(f) battery capacity 6000 MJ, panels surface 12000 m²

Figure 12: Portland experiments with Google consumer (a-b) and NEF consumer (c-f).


- a weighted sum of the previous three criteria.

This last criterion deserves some words of explanation. Whereas the maximum relative distance and the average relative distance are of the same order of magnitude as the relative distance between benchmark and model, the root mean square distance has the same order of magnitude as the standard deviation of the benchmark values. The weighted sum is such that all three terms have the same order of magnitude.


The selection criterion is to choose the configuration having the smallest value.




(a) battery capacity 5000 MJ, panels surface 1200 m²




(b) battery capacity 10000 MJ, panels surface 1200 m²




(c) battery capacity 5000 MJ, panels surface 1200 m²



(d) battery capacity 10000 MJ, panels surface 1200 m²



(e) battery capacity 6000 MJ, panels surface 1200 m²



(f) battery capacity 6000 MJ, panels surface 12000 m²

Figure 13: Golden experiments with Google consumer (a-b) and NEF consumer (c-f).

In Figures 14-15, we report graphically the values of the first three criteria obtained for all 180 configurations for Portland and Golden.

As we want the configuration in terms of the number of day intervals and the number of day types that achieves the smallest distances, we need to look for those configurations with the brightest colors in each graphic. Our observations are as follows. First, it came as a surprise to us to see huge differences between neighboring configurations. For instance, from Figure 14a we observe that for Portland the maximum relative distance

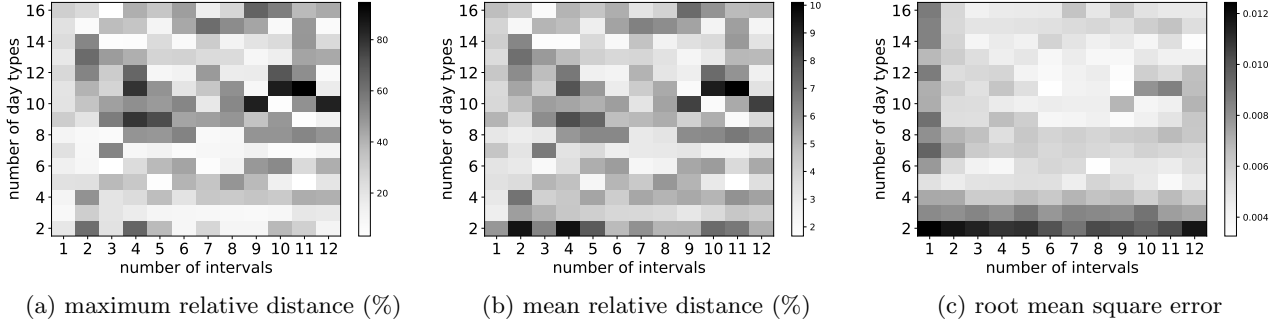


Figure 14: Performance of our forecast-based energy source model for Portland.

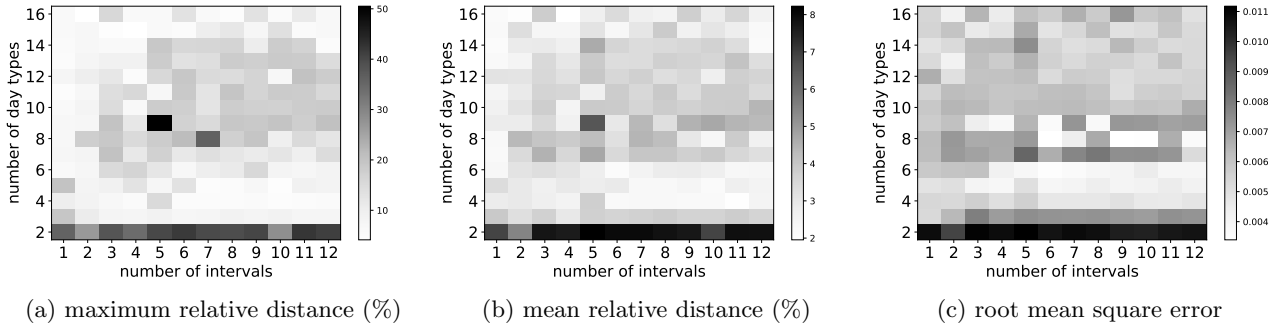


Figure 15: Performance of our forecast-based energy source model for Golden.

drastically changes when the number of day intervals changes from 9 to 10 for a number of day types equal to 10 (or 11, same observation). The value goes from 84.79 to 4.85 (and from 9.13 to 85.40 when there are 11 day types). Second, we observe an interesting trend in Figure 15a where configurations in the center of the graphics achieve worse maximal relative distances than configurations at the periphery. This is observed also in Figure 15b although to a lesser extent. There are no clear trend in any of the graphics of Figure 14. Third, many configurations are candidates for being the best achiever as they have bright spots in all three graphics for each city. To identify these with certainty, we resort to a ranking analysis.

We first discuss the case of the city of Portland. Table 7 lists the top-10 configurations according to the sum of all three criteria. Configuration (5, 11) is ranked first according to both the maximum/mean relative distance and the sum of three criteria, and ranked 8th according to the root mean square distance. The fact that this configuration yields small errors according to the three criteria is confirmed visually in Figure 14, although the neighboring configurations are not as good. This configuration with 5 day types and 11 day intervals is selected as the most promising one (for Portland) and will be used for further experiments.

Table 7: Top model configurations for Portland: ranking by sum of three criteria

day types	day intervals	maxRelDistance (%)		meanRelDistance (%)		RMSE×1000		Sum	
		value	rank	value	rank	value	rank	value	rank
5	11	3.17	1	1.67	1	3.75	8	8.58	1
6	3	3.83	2	2.11	7	4.30	20	10.24	2
16	3	4.23	3	2.06	6	4.37	25	10.66	3
14	3	4.69	4	1.91	3	4.07	16	10.67	4
10	10	4.85	6	2.05	5	4.02	14	10.92	5
14	4	5.41	8	1.89	2	4.69	43	11.98	6
9	11	4.89	7	2.12	8	5.02	56	12.04	7
5	5	4.70	5	2.33	12	5.26	69	12.29	8
6	7	5.80	9	2.30	11	5.50	87	13.60	9
7	11	6.10	10	2.52	15	5.63	93	14.25	10

We now follow the same approach for the city of Golden. Table 8 displays the top-10 configurations with respect to the sum of all criteria, together with the individual ranking for each of them. When adding criteria about a small maximum and small mean relative distance, only three remain: configurations (16, 2), (5, 10) and (4, 10). None of these configurations are in the top 10 according to the lowest root mean square error achieved in the experiments. Among the three finalists, we select the configuration with 16 day types and 2 day intervals as the most promising one (for Golden) since it is ranked first according to the maximum relative distance criterion (and the sum criterion), and third according to the mean relative distance criterion. This configuration is ranked 23rd according to the smallest root mean square error criterion; note that the first-ranked configuration has value 0.00339, compared to our choice’s value 0.00426. We have used this configuration for further experiments.

Table 8: Top model configurations for Golden: ranking by sum of three criteria

day types	day intervals	maxRelDistance (%)		meanRelDistance (%)		RMSE×1000		Sum	
		value	rank	value	rank	value	rank	value	rank
16	2	4.14	1	2.09	3	4.26	23	10.49	1
5	10	4.92	4	2.22	8	4.23	20	11.37	2
5	7	5.34	6	2.30	18	4.13	19	11.77	3
15	4	4.36	2	2.46	26	4.96	41	11.77	4
15	3	4.89	3	2.46	27	4.62	32	11.97	5
4	10	5.43	7	2.25	9	4.66	33	12.34	6
5	8	5.61	9	2.34	20	4.48	27	12.43	7
4	11	5.82	15	2.03	2	4.62	31	12.47	8
6	6	7.04	37	2.17	4	3.57	4	12.78	9
4	6	5.75	11	2.28	14	4.77	35	12.80	10

By selecting for each city the configuration that achieved the smallest maximum relative distance in the first series of experiment, we are confident of the robustness of our energy model with this configuration.

For completeness, we report in Appendix C the relative distance between benchmark and model for each performance metric (empty/full battery proportion) achieved by each city’s selected configuration over the first series of experiments.

5.4.3 Further experiments

To show how the energy models perform on a wider range of parameters, we performed a second series of simulations. For each city and for each of the energy sources, we first varied the battery capacity and fixed the surface of the photovoltaic panels, then did the opposite (we fixed the battery capacity and varied the surface of the photovoltaic panels). For each single simulation we record the empirical distribution of the battery level.

We depict in Figures 16-17 the battery full/empty percentages as a function of the battery capacity or alternatively as a function of the panels surface. There are five curves in each figure, one for each of the energy sources considered, namely: Sinusoidal, Daily and Hourly (described in Sections 2.3 and 2.4), Benchmark (which corresponds to the global irradiance measured during a one-year period [2, 26]), and Synthetic (which corresponds to our forecast-based energy model described in Section 4.1 with 5 day types and 11 day intervals in the case of Portland, and with 16 day types and 2 day intervals in the case of Golden). It is obvious that the performance achieved when using our forecast-based sources matches perfectly the benchmark performance. This strengthens the previous observations that we made with experiments and confirms the supremacy of our energy models over the other models in the use case at hand.

6 Conclusions

We have considered in this paper the problem of modeling the solar energy at a scale fine enough to be useful for the analysis of solar-powered computer systems. We have proposed a novel way of generating samples of solar irradiance for future days, based on the knowledge of the weather outlook. This method involves clustering a recorded trace of irradiance with easily computed features. These features are based on the *clear sky index* instead of the clear sky irradiance. We demonstrate, through simulations, that our new solar source performs better in predicting battery shortage and overflows than previously proposed models, when feeding a datacenter submitted to a workload obtained from real traces.

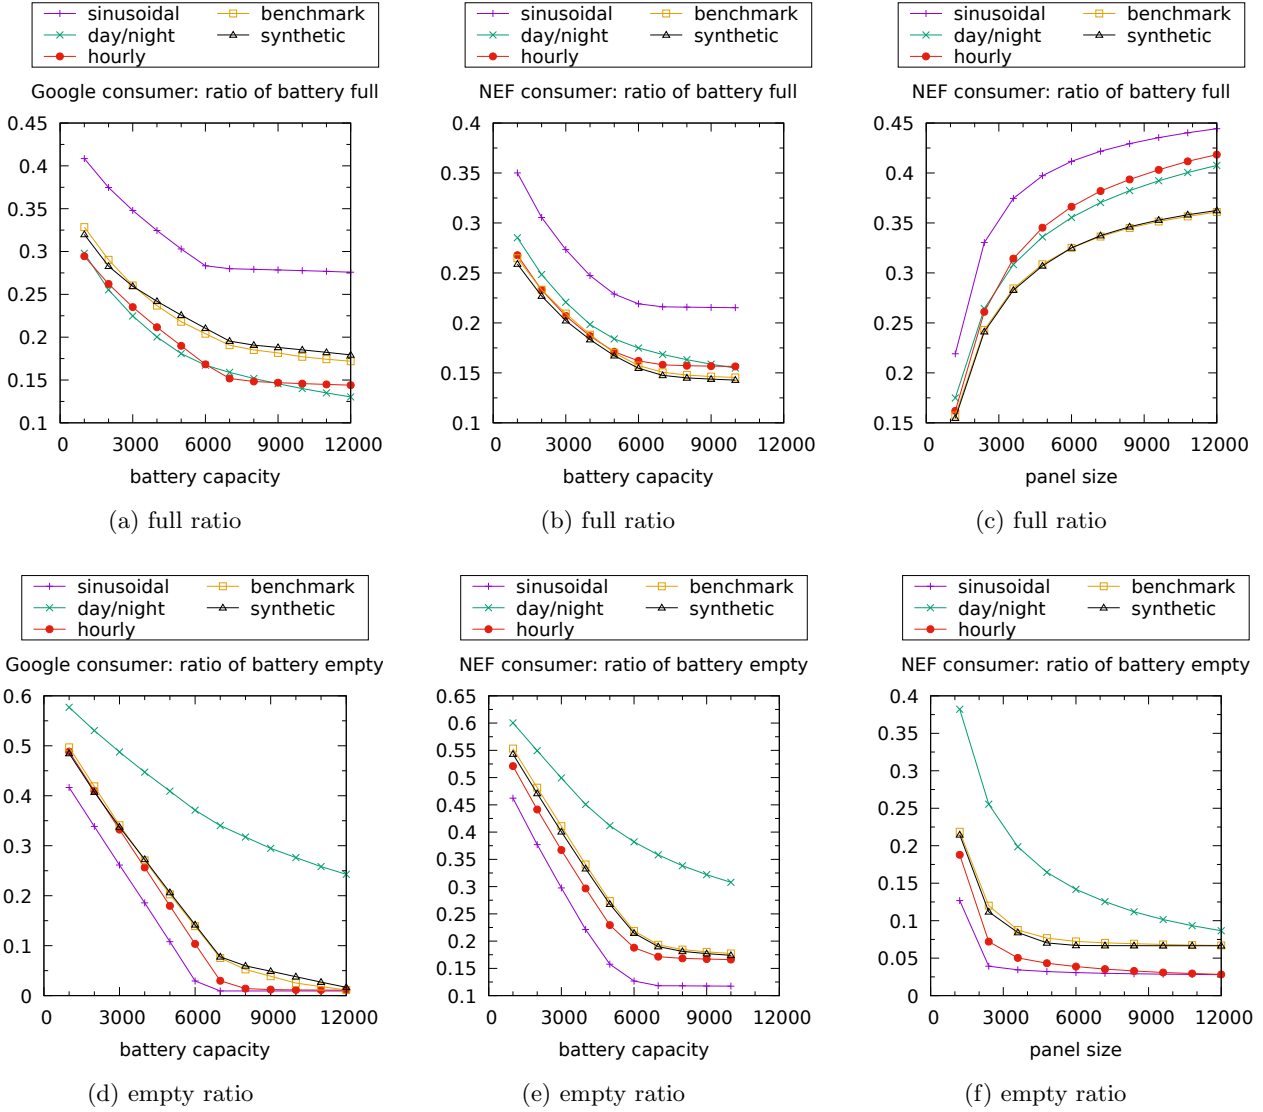


Figure 16: Experiments for Portland: full ratio (a-c) and empty ratio (d-f) as the battery capacity varies (a-b, d-e; panels surface 1200 m²) and as the panels surface varies (c, f; battery capacity 6000 MJ).





Figure 17: Experiments for Golden: full ratio (a-c) and empty ratio (d-f) as the battery capacity varies (a-b, d-e; panels surface 1200 m²) and as the panels surface varies (c, f; battery capacity 6000 MJ).

Next steps in the analysis will aim at confirming the robustness of our approach by confronting the synthetic source to more energy consumption models (i.e. datacenter workload models), and fitting the model to more solar data. With more experimental results, we will investigate further the question of selecting the proper configuration of clusters and intervals for our model. We also plan to develop a queuing-theoretic analysis as well as a control-theoretic analysis of datacenters, based on our new solar source model. For this purpose, a (semi)Markovian representation of the clear sky index process may prove handy.

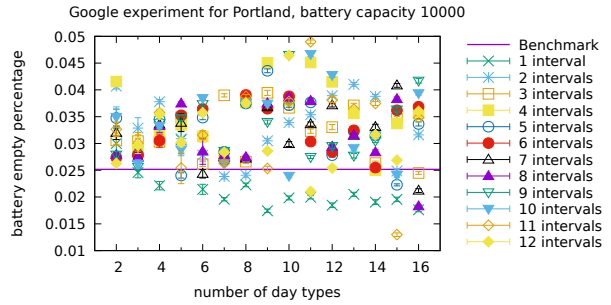
The data and the programs used in this study will, after review, be made available through `gitlab`.

A Empty ratio: Detailed performance

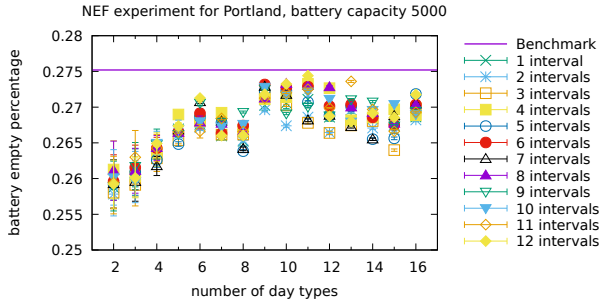
The percentage of time when the battery is *empty* is depicted with its 95% confidence intervals in Figures 18 (for Portland) and 19 (for Golden) with NEF/Google consumers and for several settings of the battery capacity and panels surface. We observe that some configurations of our energy models perform better than others in a single setting, but this changes from one setting to another. When the configuration has a small number of types, equal to 2 or 3 (or even 4), the confidence intervals are larger and the performance metrics farther from the benchmark than the other configurations in most cases. We conclude from these figures that no configuration in terms of number of day types and number of day intervals is always the best and other selection criteria should be determined to select the most promising configurations. Nevertheless, we believe that one has the possibility to use any of the configurations (except for 2-4 day types) and still be confident in the accuracy of the results.




(a) battery capacity 5000 MJ, panels surface 1200 m²



(b) battery capacity 10000 MJ, panels surface 1200 m²




(c) battery capacity 5000 MJ, panels surface 1200 m²




(d) battery capacity 10000 MJ, panels surface 1200 m²

Figure 18: *Cont.*




(e) battery capacity 6000 MJ, panels surface 1200 m²




(f) battery capacity 6000 MJ, panels surface 12000 m²


Figure 18: Empty ratio, Portland: Detailed performance of our models for several settings of the battery capacity and panels surface with Google consumer (a-b) and NEF consumer (c-f).




(a) battery capacity 5000 MJ, panels surface 1200 m²




(b) battery capacity 10000 MJ, panels surface 1200 m²




(c) battery capacity 5000 MJ, panels surface 1200 m²



(d) battery capacity 10000 MJ, panels surface 1200 m²



(e) battery capacity 6000 MJ, panels surface 1200 m²




(f) battery capacity 6000 MJ, panels surface 12000 m²


Figure 19: Empty ratio, Golden: Detailed performance of our models for several settings of the battery capacity and panels surface with Google consumer (a-b) and NEF consumer (c-f).

B Full ratio: Detailed performance


Similarly to what was done in Appendix A, the percentage of time when the battery is *full* is depicted with its 95% confidence intervals in Figures 20 (for Portland) and 21 (for Golden) with NEF/Google consumers and for several settings of the battery capacity and panels surface. We are again unable to visually identify with certainty the best model configuration. The observations made in Appendix B are valid here.




(a) battery capacity 5000 MJ, panels surface 1200 m²




(b) battery capacity 10000 MJ, panels surface 1200 m²




(c) battery capacity 5000 MJ, panels surface 1200 m²



(d) battery capacity 10000 MJ, panels surface 1200 m²




(e) battery capacity 6000 MJ, panels surface 1200 m²




(f) battery capacity 6000 MJ, panels surface 12000 m²


Figure 20: Full ratio, Portland: Detailed performance of our models for several settings of the battery capacity and panels surface with Google consumer (a-b) and NEF consumer (c-f).




(a) battery capacity 5000 MJ, panels surface 1200 m²




(b) battery capacity 10000 MJ, panels surface 1200 m²




(c) battery capacity 5000 MJ, panels surface 1200 m²



(d) battery capacity 10000 MJ, panels surface 1200 m²



(e) battery capacity 6000 MJ, panels surface 1200 m²



(f) battery capacity 6000 MJ, panels surface 12000 m²

Figure 21: Full ratio, Golden: Detailed performance of our models for several settings of the battery capacity and panels surface with Google consumer (a-b) and NEF consumer (c-f).

C Relative distance achieved by the selected configurations

We use the labels introduced in the last row of Table 6 to refer to the settings used in the first series of experiments. The relative distance between benchmark and model is reported in Table 9 for the selected configurations for each city. In addition we report the mean obtained over all 6 values regarding the same performance metric.

Table 9: Relative distance (in percentage) in the first series of experiments

City	Types	Intervals	Metric	G5000	G10000	N5000	N10000	N1200	N12000	Mean
Portland	5	11	empty	3.17	1.07	2.11	1.33	1.97	0.59	1.70
			full	2.36	2.95	1.30	1.51	1.45	0.23	1.63
Golden	16	2	empty	3.82	3.01	4.14	3.92	3.54	0.51	3.15
			full	0.60	3.37	0.75	0.21	0.16	1.06	1.03

References

- [1] Clusters home, visited on Nov, 17, 2020. https://wiki.inria.fr/ClustersSophia/Clusters_Home.
- [2] Afshin Andreas and T. Stoffel. NREL Solar Radiation Research Laboratory (SRRL): Baseline Measurement System (BMS); Golden, Colorado (Data). Report DA-5500-56488, NREL, 1981.
- [3] R. E. Bird and R. L. Hulstrom. Review, Evaluation, and Improvement of Direct Irradiance Models. *Journal of Solar Energy Engineering*, 103(3):182–192, 08 1981.
- [4] Richard E. Bird and Roland L. Hulstrom. A simplified clear sky model for direct and diffuse insolation on horizontal surfaces. Technical Report SERI/TR-642-761, Solar Energy Research Institute, February 1981.
- [5] 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 *Proceedings of SustainIT: Sustainable Internet and ICT for Sustainability*, pages 1–8, Madrid, Spain, April 2015.
- [6] J. V. Dave, P. Halpern, and H. J. Myers. Computation of incident solar energy. *IBM Journal of Research and Development*, 19(6):539–549, November 1975.
- [7] Ioannis Dimitriou, Sara Alouf, and Alain Jean-Marie. A markovian queueing system for modeling a smart green base station. In *Proceedings of EPEW: European Performance Evaluation Workshop*, volume 9272 of *LNCS*, pages 3–18, Madrid, Spain, August 2015. http://dx.doi.org/10.1007/978-3-319-23267-6_1.
- [8] J.M. Fourneau. Modeling green data-centers and jobs balancing with energy packet networks and interrupted poisson energy arrivals. *SN COMPUT. SCI.*, pages 1–28, 2020.
- [9] Yashar Ghiassi-Farrokhfal, Srinivasan Keshav, Catherine Rosenberg, and Florin Ciucu. Solar power shaping: An analytical approach. *IEEE Transactions on Sustainable Energy*, 6(1):162–170, January 2015. <http://dx.doi.org/10.1109/TSTE.2014.2359795>.
- [10] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, and R. Bianchini. Designing and managing data centers powered by renewable energy. *IEEE Micro*, 34(3):8–16, 2014.
- [11] Lianhong Gu, Jose D. Fuentes, Michael Garstang, Julio Tota da Silva, Ryan Heitz, Jeff Sigler, and Herman H. Shugart. Cloud modulation of surface solar irradiance at a pasture site in southern Brazil. *Agricultural and Forest Meteorology*, 106(2):117–129, January 2001.
- [12] Christian A. Gueymard. A review of validation methodologies and statistical performance indicators for modeled solar radiation data: Towards a better bankability of solar projects. *Renewable and Sustainable Energy Reviews*, 39:1024 – 1034, 2014.
- [13] M. Jurado, J.M. Caridad, and V. Ruiz. Statistical distribution of the clearness index with radiation data integrated over five minute intervals. *Solar Energy*, 55(6):469–473, December 1995.
- [14] Wafa Khelif. How sustainable datacenters can be ? Master thesis, University Nice Sophia Antipolis, France, August 2015.
- [15] Fanxin Kong and Xue Liu. A survey on green-energy-aware power management for datacenters. *ACM Comput. Surv.*, 47(2), November 2014.
- [16] Zhenhua Liu, Yuan Chen, Cullen Bash, Adam Wierman, Daniel Gmach, Zhikui Wang, Manish Marwah, and Chris Hyser. Renewable and cooling aware workload management for sustainable data centers. *SIGMETRICS Perform. Eval. Rev.*, 40(1):175—186, June 2012.
- [17] Marco Miozzo, Davide Zordan, Paolo Dini, and Michele Rossi. SolarStat: Modeling photovoltaic sources through stochastic Markov processes. In *Proceedings of 2014 IEEE International Energy Conference*, pages 688–695, Dubrovnik, Croatia, May 2014. <http://dx.doi.org/10.1109/ENERGYCON.2014.6850501>.
- [18] Giovanni Neglia, Matteo Sereno, and Giuseppe Bianchi. Geographical load balancing across green datacenters: A mean field analysis. *ACM SIGMETRICS Performance Evaluation Review*, 44(2), 9 2016.

- [19] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [20] Christian Piedallu and Jean-Claude Gégout. Multiscale computation of solar radiation for predictive vegetation modelling. *Annals of Forest Science*, 64(8):899–909, January 2007.
- [21] Dimitra Politaki and Sara Alouf. Stochastic models for solar power. In *Proceedings of EPEW: European Performance Evaluation Workshop*, volume 10497 of *LNCS*, pages 282–297, Berlin, Germany, September 2017. http://dx.doi.org/10.1007/978-3-319-66583-2_18.
- [22] ptaff.ca. Sunrise, sunset daylight in a graph. <https://ptaff.ca/soleil/>.
- [23] Charles Reiss, John Wilkes, and Joseph L. Hellerstein. Google cluster-usage traces: format + schema. Technical report, Google Inc., Mountain View, CA, USA, November 2011. Revised 2014-11-17 for version 2.1. Posted at <https://github.com/google/cluster-data>.
- [24] H. Suehrcke and P.G. McCormick. The frequency distribution of instantaneous insolation values. *Solar Energy*, 40(5):413 – 422, 1988.
- [25] Sun Sun, Fiodar Kazhamiaka, Srinivasan Keshav, and Catherine Rosenberg. Using synthetic traces for robust energy system sizing. In *Proc. e-Energy'19*, Phoenix, USA, June 2019.
- [26] F. Vignola and Afshin Andreas. University of Oregon: GPS-based Precipitable Water Vapor (Data). Report DA-5500-64452, NREL, 2013. <http://dx.doi.org/10.7799/1183467>.
- [27] weatherexplained.com. 2001 denver, colorado (den) - average, low, world, daily, high, snowfall, days, normals, means, and extremes, reference notes:, retrieved on Oct, 28, 2020. <http://www.weatherexplained.com/Vol-2/2001-Denver-Colorado-DEN.html>.
- [28] weatherexplained.com. 2001 portland, oregon (pdx) - average, low, world, daily, high, snowfall, days, normals, means, and extremes, reference notes:, retrieved on Oct, 28, 2020. <http://www.weatherexplained.com/Vol-5/2001-Portland-Oregon-PDX.html>.