Title

Characterizing and Mitigating the Impact of Solar Forecast Errors on Grid Planning and Operations

Author

Wang, Guang Chao

Publication Date

2019

UNIVERSITY OF CALIFORNIA SAN DIEGO

Characterizing and Mitigating the Impact of Solar Forecast Errors on Grid Planning and Operations

A Dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Engineering Sciences (Mechanical Engineering)

by

Guang Chao Wang

Committee in charge:

Professor Jan Peter Kleissl, Chair
Professor Carlos F. M. Coimbra
Professor Mauricio De Oliveira
Professor Tara Javidi
Professor Tajana Simunic Rosing

2019

The Dissertation of Guang Chao Wang is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

DEDICATION

This dissertation is dedicated to my beloved family, for their loving support, encouragement and sacrifices.

EPIGRAPH

One, remember to look up at the stars and not down at your feet.

Two, never give up work. Work gives you meaning and purpose and life is empty without it.

Three, if you are lucky enough to find love, remember it is there and don't throw it away.

—Stephen Hawking

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

First and foremost, I am indebted to my PhD advisor Dr. Jan Kleissl. When I was wandering around at the crossroad in my life, he extended an olive branch to me with a researcher position and subsequently a PhD offer. Ever since then his continuous support, advice and guidance throughout the course of my PhD time have been critical to the success of my PhD defense and dissertation. Thank you Jan for offering me this invaluable opportunity to work and grow with you and your team for six years, which fundamentally steers my life direction to a path that I have been dreaming of.

Besides Dr. Jan Kleissl, my gratitude also goes to the rest of my committee: Profs. Carlos F. M. Coimbra, Tara Javidi, Tajana Simunic Rosing. I especially want to extend my thanks to Prof. Mauricio de Oliveira with regard to his constructive suggestions on the last part of the research in my dissertation.

I have also come to appreciate my fellow PhD colleague Benjamin Kurtz for his incredible help and support. He knows more about maths, coding, and programming tips than perhaps anyone else that I have ever met, and he has been kind enough to often impart his knowledge to me. He has been my best friend since I came to the United States twelve years ago, which is one of the biggest gains in my PhD time.

Furthermore, I have been lucky to have ideas, support, and assistance from many other friends here. Thanks particularly to Andu Nguyen, Bryan Urquhart, Changfu Li, Mengying Li, Ryan Hanna and Shuang Cui for their help and company.

Finally, and perhaps most importantly, I want to offer my gratitude and appreciation to my parents Zhengguang Wang, Chaomin Fu for their unconditional love and support, and my grandparents: Tingshu Fu, Sumian Hou, Yuelun Lu, especially Cuixiang Wang who spends countless of time in his life on everything he can think of to help me succeed. None of my achievements would be possible without his help. Thank you!

Chapter 2, in full, is a reprint of the material as it appears in “Wang, G. C., Urquhart, B. & Kleissl, J. (2019). Cloud base height estimates from sky imagery and a network of pyranometers. *Solar Energy*, 184, 594-609.” The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is a reprint of the material as it appears in “Wang, G., Kurtz, B. & Kleissl, J. (2016): Cloud base height from sky imager and cloud speed sensor, *Solar Energy*, 131, 208-221.” The dissertation author was the primary investigator and author of this paper.

Chapter 4, in full, is a reprint of the material as it is expected to appear (accepted pending major revisions) in “Wang, G. C., Kurtz, B., Bosch, J. L., de la Parra I. & Kleissl, J. (2020): Maximum expected ramp rates using cloud speed sensor measurements. *Journal of Renewable and Sustainable Energy*." The dissertation author was the primary investigator and author of this paper.

Chapter 5, in full, is a reprint of the material as it appears in "Wang, G. C., Ratnam, E., Haghi, H. V. & Kleissl, J. (2019): Corrective receding horizon EV charge scheduling using short-term solar forecasting. *Renewable energy*, 130, 1146-1158." The dissertation author was the primary investigator and author of this paper.

Chapter 6, in full, is still in an initial draft form for future publication with Dennis van der Meer and my PhD advisor Dr. Jan Kleissl. The dissertation author was the primary investigator and author of this manuscript.

VITA

2019	Ph.D. in Engineering Sciences (Mechanical Engineering), University of California San Diego, La Jolla, CA, U. S.
2012	M.S. in Finance, Simon Business School, University of Rochester, Rochester, NY, U. S.
2011	M.S. in Mechanical and Aerospace Engineering, The Ohio State University, Columbus, Ohio, U. S.
2008	B.S. in Mechatronic Engineering and Automation, Shanghai University, Shanghai, China

PUBLICATIONS

Wang, G. C., van der Meer, D. and Kleissl, J. (2020): Stochastic Optimization of EV Scheduling Incorporating Real-Time Charge Events. In Preparation.

Lappalainen, K., **Wang, G. C.** and Kleissl, J. (2020): Estimation of the largest expected PV power ramp rates. *Applied Energy*. In Preparation.

Wang, G. C., Kurtz, B., Bosch, J. L., de la Parra I. and Kleissl, J. (2020): Maximum expected ramp rates using cloud speed sensor measurements. *Journal of Renewable and Sustainable Energy*. Accepted Pending major Revisions.

Wang, G. C., Bosch, J. L., Kurtz, B., de la Parra, I., Wu, E. and Kleissl, J. (2019, June): Worst expected ramp rates using cloud speed measurements. *In 2019 IEEE 46th Photovoltaic Specialist Conference (PVSC)*. IEEE.

Wang, G. C., Urquhart, B. and Kleissl, J. (2019): Cloud base height estimates from sky imagery and a network of pyranometers. *Solar Energy*, 184, 594-609.

Wang, G. C., Ratnam, E., Haghi, H. V. and Kleissl, J. (2019): Corrective receding horizon EV charge scheduling using short-term solar forecasting. *Renewable energy*, 130, 1146-1158.

Wang, G. C. and Gross, K. C. (Dec 2018): Telemetry Parameter Synthesis System for Enhanced Tuning and Validation of Machine Learning Algorithmics. *IEEE 2018 Intn'l Symposium on Internet of Things & Internet of Everything (CSCI-ISOT)*, Las Vegas, NV.

Wang, G. C. and Gross, K. C. (Dec 2018): Real Time Empirical Synchronization of IoT Signals for Improved AI Prognostics. *IEEE 2018 Intn'l Symposium on Computational Intelligence (CSCI-ISCI)*, Las Vegas, NV.

Marcos, J., Parra, I. D., Cirs, E., **Wang, G.**, Garca, M. and Marroyo, L. (2018, September): Ramp-Rate Control in Large PV Plants: Battery vs. Short-Term Forecast. *In 2018 20th European Conference on Power Electronics and Applications (EPE'18 ECCE Europe)* (pp. P-1). IEEE.

Wang, G., Disfani, V. and Kleissl, J. (2018, May): Scenario based stochastic optimization of probabilistic EV charging scheduling. *In 2018 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia)* (pp. 552-557). IEEE.

Kuhn, P., Wirtz, M., Wilbert, S., Bosch, J. L., **Wang, G.**, Ramirez, L. and Pitz-Paal, R. (2018): Field validation and benchmarking of a cloud shadow speed sensor. *Solar Energy*, 173, 229-245.

Wang, G., Kurtz, B. and Kleissl, J. (2016): Cloud base height from sky imager and cloud speed sensor. *Solar Energy*, 131, 208-221.

PATENTS

Merged-Surface EMI Fingerprint FastScan Technique for High Sensitivity SpyChip and Counterfeit-Component Detection in Electronic Systems, K. C., Gross, G. C., Wang, M. Dayringer and A. J. Lewis, ORA190577 (Oracle Invention Disclosure, September 16, 2019).

Intelligent Data Preprocessing Algorithm for Optimal Loadshape Forecasting of Utility Distributed Energy Sources, C. T. Eyeford III, K. C., Gross and G. C., Wang, ORA190335 (Oracle Invention Disclosure, July 1, 2019).

Method and Apparatus for Optimal Multi-Mode Seasonality Characterization and Accommodation for Improved IoT Prognostics, G. C., Wang and K. C., Gross, ORA190112 (Oracle Invention Disclosure, May 22, 2019).

Insertable-Device EMI Fingerprint Characterization and Security Assurance Certification for Enterprise Servers and Storage Systems, A. J. Lewis, K. C. Gross, M. Dayringer and G. C., Wang, ORA190337 (Oracle Invention Disclosure, April 15, 2019).

Digital Twin for Advanced Prognostics and Project Ahead Forecasting of Complex Engineering Assets, K. C. Gross and G. C., Wang, ORA190241 (Oracle Invention Disclosure, April 1, 2019).

Camouflaging EMI Fingerprint Activity in Enterprise Computers for Enhanced Cyber Security, K. C. Gross, A. George and G. C., Wang, ORA190156 (Oracle Invention Disclosure, January 25, 2019).

High Fidelity Despiking/Respiking Synthesizer for Prognostic IoT Industrial Applications, G. C., Wang and K. C., Gross, ORA190103 (Oracle Invention Disclosure, December 10, 2018).

Ambient Thermal Compensation Method and Apparatus for Optimal Prognostics of Utility Critical Grid Assets, K. C. Gross, G. C., Wang and E. Wetherbee, ORA190031 (Oracle Invention Disclosure, November 9, 2018).

Automated and Optimal Time-Series Resampling Process for Big Data IoT Applications, K. C., Gross and G. C., Wang, ORA180722 (Oracle Invention Disclosure, October 23, 2018).

UnStairStepping: Method and Apparatus for Enhanced Prognostics for IoT Industries, K. C. Gross and G. C., Wang, ORA190055 (Oracle Invention Disclosure, September 11, 2018).

Automated Formulized Data Reduction (AFDR) for Big Data Time-Series Archives, K. C., Gross, G. C., Wang, S. Jeffreys, A. Wood and C. MacMillan, ORA180260-US-CIP-1 (Oracle Invention Disclosure, August 2, 2018).

Optimal Missing Value Imputation for Prognostic Analysis of IoT Time-Series Signals, G. C., Wang, K. C., Gross and D. Gawlick, ORA180645 (Oracle Invention Disclosure, June 11, 2018).

ABSTRACT OF THE DISSERTATION

Characterizing and Mitigating the Impact of Solar Forecast Errors on Grid Planning and Operations

by

Guang Chao Wang

Doctor of Philosophy in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2019

Professor Jan Peter Kleissl, Chair

In recent years, the contribution of photovoltaic (PV) power production to the electric grid has been increasing. Still, a number of challenges remain for a reliable and efficient integration of solar energy. While conventional electric power generated by gas turbines can be adjusted to follow the grid load, the stochastic nature of solar radiation makes it difficult to control the PV output, which hinders its integration in the grid. Accurate solar forecasts help grid operators integrate solar energy by enhancing power quality and reducing grid operation costs. Following the development of sky imager hardware and algorithms at UC San Diego, we present a variety of models and methodologies to reduce sky imager forecasts errors by improving the accuracy of meteorological parameters, compensating the power mismatch caused by solar forecasts errors, and mitigating the impact of solar forecast errors on real world grid planning and operations.

First, a low-cost instrument for measuring local cloud motion vectors (CMVs) was developed.

Three algorithms for estimating local cloud base height (CBH) using a single sky imager paired with either distributed ground irradiance sensors or measured CMVs were then designed and tested. Since sky imager forecasts are often used in conjunction with other instruments for measuring CBH, cloud velocity, and/or solar irradiance measurements, our approaches decrease instrumentation costs and logistical complexity. More importantly, through these algorithms, local measurements improve sky imager forecasts by adding information that is unobservable from a single sky imager.

Second, integrating battery systems into a PV plant can compensate the power imbalance caused by solar forecast errors. Battery system size can be optimized by determining the energy reserve required to offset the possible maximum power ramp. Because passing cloud shadows are the main cause of the power ramps, a simple model based on physics variables that are available globally can determine the worst power ramp rates. Local CMV measurements enable even more accurate maximum ramp rate estimates. The key merit of the method is that it is universally applicable in the absence of high frequency measurements.

Finally, issues when integrating imperfect solar forecasts in grid operations are evaluated. Both physics-based forecasts and Numerical Weather Prediction (NWP) based machine learning forecasts that are commonly utilized in the grid operations exhibit autocorrelated forecast errors. First, a deterministic valley-filling problem through EV charging is formulated to investigate how the autocorrelated forecast errors increase peak demand and cause grid net load variability. Then a corrective optimization framework is proposed to minimize the deviation of the realistic valley filling solutions from the ideal solutions. In addition, with the goal of operational deployment, stochastic programming incorporating real time updates of solar forecast and EV charge events to address real-world uncertainty is employed. The optimal valley filling problem is solved in an innovative way and executed under a predictive control scheme in the presence of autocorrelated forecast errors. The proposed corrective stochastic optimization framework successfully mitigates the impact of autocorrelated forecasts errors on grid operations.

Chapter 1

Introduction

1.1 Research Background

Photovoltaic (PV) installations have been experiencing significant growth resulting in unprecedented structural changes in the grid system [3]. However, solar generation exhibits high dependence on cloud cover variability, atmospheric aerosol levels, and other meteorological processes. The intermittent and variable nature of solar power introduces significant challenges to grid energy management regarding system stability [4], electric power balance [5], reactive power capacity [6], and frequency response and inertia [7]. As the share of electricity produced by PV increases, these effects aggravate and pose even greater challenges to the electric grid due to considerable variability of the irradiance and the reverse power flow in the distribution grid that results from PV generation.

While the needs for PV systems to be fully integrated into electricity grids arise, solar power forecasting becomes crucial, which allows grid operators and power system designers to design optimal solar PV plants as well as manage the power of demand and supply. Accurate solar forecasts facilitate the needs of grid operations and control activities, including grid regulation, power scheduling, and unit commitment in both the distribution and transmission grids [8]. If the electricity produced by solar power is perfectly known beforehand, it would be possible to dynamically adjust its output in response to real time demands from the grid. Intensive research has been done to develop forecasting models for solar resources and power output of PV plants at utility scale level in the past decades. A comprehensive review is presented by Sobri et al. [9].

Solar forecasting techniques can be primarily classified into three categories: 1) statistical methods

2) physical modeling and 3) ensemble methods [10]. Therein, physics based models the interactions between solar radiations and atmospheric components. They can be further categorized regarding the spatial and temporal resolutions, as presented by Figure 1.1. We have primarily been interested in the physics based solar forecasting through sky imager for its very high spatial and temporal resolution for time horizons less than 15-20 minutes on domains in size of a few km. While the other two methods have begun to encroach on the traditional domains where sky imagery is preferred, we target the utilities who own large scale PV plants where spatial aggregation/smoothing may not be sufficient to alleviate the need for a fine-grained forecast.

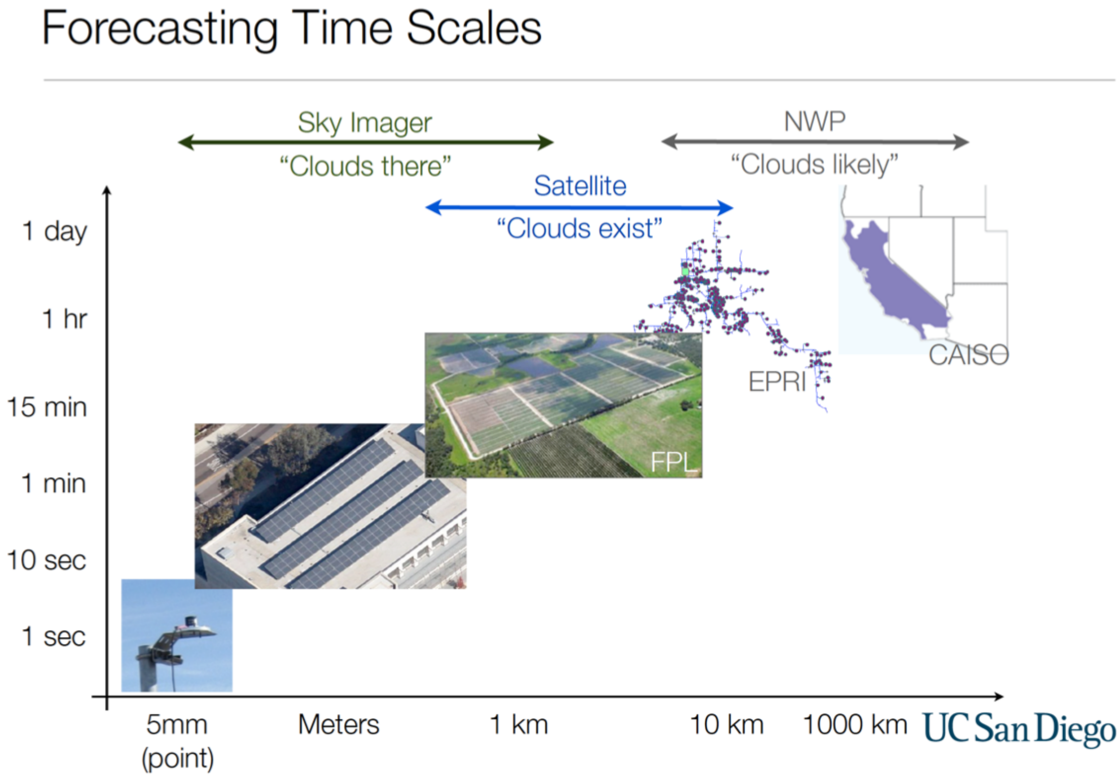

Figure 1.1: Physics based solar forecasting time scales vs. spatial resolution.

On the other hand, due to the chaotic nature of atmospheric dynamics, precise solar power forecasting can be extremely difficult. The erroneous solar power forecasts hinder solar power integration since forecast errors can cause sub-optimal solutions in grid planning or even worsen grid operation. If solar forecasts are ever to be operationally used in the grid system, simultaneous actions need to be taken to alleviate the concerns and address the issues of imperfect solar forecasts in the following three aspects: 1) reduce solar forecast errors 2) compensate the power mismatch resulted from solar forecast errors 3) reduce the impact of solar forecast errors in real world grid operations. The combined actions and resulting improvements are preferable to ensure a secure and economic integration of PVs into the grid system. Thus, each of the above is covered in this thesis and further introduced in the following chapters.

1.2 Problem of Interest

With the research background discussed in the prior section, the particular grid integration problem of interest to this thesis is the classic grid oversupply problem as a result of higher penetration of PV installations. The oversupply scenario happens when all anticipated renewable energy generation dominated by PV production significantly exceeds the real-time power demand during solar peak time. That yields a net load curve appearing a "belly" shape in the mid-afternoon that quickly ramps up to produce an "arch" afterward, which is typically described as the net load energy valley or duck curve [11]. Figure 1.2 presents this power mismatch between significant PV output and grid load demand during the solar peak hours, whereas the potential for this energy valley is anticipated to increase because of the continually increasing level of PV penetrations.

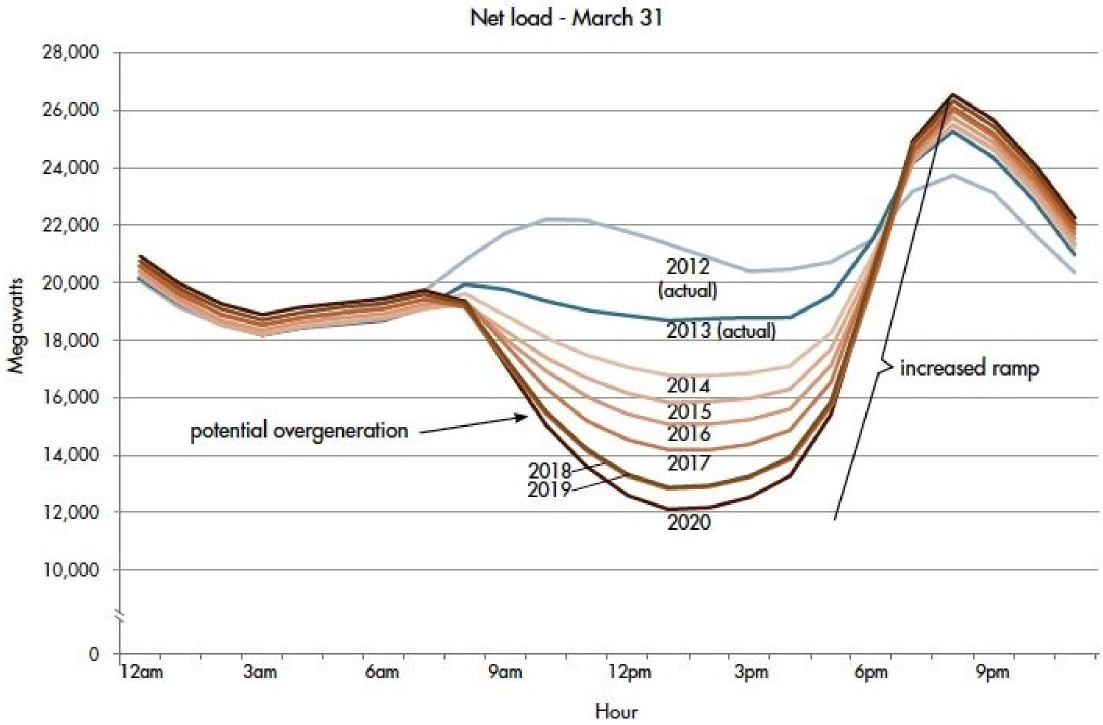

Figure 1.2: The duck curve shows steep ramping needs and overgeneration risk caused by increasing PV installations. Photo Credit: California ISO [1].

The energy valley is a concern because it can cause bidirectional power flow and wide voltage swing, and that steep up-ramp after solar production time additionally exerts a large power peak onto the grid network. One solution is to reinforce the power grid. However, this involves a huge amount of investments in upgrading grid infrastructures, which is costly. The alternative solution is to strategically increase power demand through energy storage systems (ESS) during the solar time to absorb the excessive solar energy, which is known as a valley filling strategy. In recent years, the adoption of electric vehicles (EV) as a load demand is becoming more popular since it naturally carries both cost and environmental advantages. However, uncontrolled EV charging behavior in large numbers can also negatively influence the grid network [12] by creating load peaks in the aggregated load curve, degrading power quality [13] and increasing operational cost [14].

Thus, it fills two needs with one deed if EVs can serve as ESSs that absorb the surplus solar energy or supplement solar generation during the partly cloudy time under smart charging. Therefore, in this thesis, the proposed grid integration study to implement and validate all the developed techniques considers a real world valley filling problem through scheduling EV charging.

1.3 Motivations and Objectives

1.3.1 Reduce PV Forecast Errors

Cloud base height and cloud motions are important parameters for physics based high resolution solar forecasting. Incorrect CBH leads to incorrect projections of cloud shadows, and incorrect cloud velocity results in inaccurate shadow positions and timing of shadowing events, causing solar power forecast errors. Typically, the measurements for the two variables can be obtained from in-situ or remote sensing instruments, satellite data, and other indirect measuring methods. However, these approaches suffer from either prohibitive equipment or operating costs, high compute cost, inaccurate measurements because of the spatial variability in cloud cover, or insufficient temporal resolution, which adds the logistical complexity to the forecasting methodology or degrades the accuracy of intra-hour solar forecasts. It will be preferable if high resolution CBH and cloud speed measurements can be obtained at the forecast site in a low-cost instrumentation setup. Thus, we develop an instrument to measure the local cloud motion vectors (CMVs), and propose a total of three methods to estimate local CBHs utilizing a sky imager and the measured CMVs or distributed irradiance sensors that are present in nearly every solar plant nowadays. The objective of our methodologies is that the existing infrastructures for sky imagery forecasts are conveniently utilized to estimate cloud height as an ancillary product, which in return provides more accurate physics variable inputs to the sky imagery solar forecasting to reduce the forecast errors.

1.3.2 Compensate PV Forecast Errors

The mismatch between the predicted solar power and actual solar production generally are offset by ESS deployed near to the solar plants, otherwise more spinning and supplemental energy reserve are required by grid operators to maintain the stability of the PV-connected grid. However, the high cost of batteries is the major drawback of their commercial applications. Intensive battery sizing research has been done to determine an appropriate battery capacity for a given grid problem to bring it in a cost competitive position for balancing solar forecast errors. However, the studies typically require high frequency (O(1s)) solar irradiance or power measurements, and 1-min data resolution was found to be barely acceptable without inducing significant modeling errors [15]. While high frequency solar irradiance data are rarely available, only the worst power ramp rate is required to determine the required power and energy capacity rating for ESS to buffer all unexpected power ramps. Thus, in the second part of this thesis, we formulate an analytical relation between cloud velocity, PV plant dimension, and recent PV power measurements / cloud optical depth to estimate the maximum power ramp for optimal energy storage design. The objective of this work is to provide a convenient yet economical approach to sizing the battery capacity which only requires simple data input at coarse temporal resolutions that are publicly accessible.

1.3.3 Characterize the Impact of Solar Forecast Errors

With all the effort of improving forecast accuracy, the error magnitude, typically quantified by root mean square error (RMSE), is often about 10% of the installed solar power capacity [9]. However, that number is not simply carried through to the grid resource planning problem. Besides the forecast error magnitude, realistic forecast error also exhibits a structure of autocorrelation, which means the timeseries forecast errors are correlated in some ways between the present time and a later time. In the literature, grid integration studies typically assume an independent structure for timeseries forecast errors, even though this hypothesis is not explicitly stated. When solar forecast is used for a single time step ahead energy dispatch, the error independence structure is less concerned because the forecasting technique can be updated with new information at each step. When solar forecast is used for multiple time steps ahead energy dispatch, the error dependence structure presents great challenges to the problem. For example, for a fixed magnitude of forecast errors, the positive autocorrelation can severely increase the need for ESS capacity by an order of magnitude, which is an effect of “error independence pitfall” [16]. Thus, the autocorrelation property of timeseries forecast errors should be taken into account in the grid integration studies.

Note that the popular machine learning based PV forecasts naturally have smaller autocorrelation of forecast errors than physics-based PV forecasts. But since ML models usually rely on data sources, such as Numerical Weather Prediction (NWP) that also inherits autocorrelated forecast errors from its time derivative process, in practice ML forecasts can still show a certain level of autocorrelated errors. We believe both magnitude and structure of forecast errors are important determinants in the grid integration studies, while the latter unfortunately received fewer attentions by being modeled through simple statistical methods. Thus the objective of this work is to thoroughly characterize and quantify the impact of autocorrelated solar forecasts errors in the grid planning and operations, and to develop a predictive control based corrective framework to mitigate that impact with more accurate short-term forecasts.

1.3.4 Mitigate the Impact of Solar Forecast Errors on Grid Optimization

While autocorrelated solar forecast errors can be corrected in a deterministic optimization framework, real world problems always include unknown events that were not anticipated in the planning stage. Stochastic programming is a common approach to accommodate the uncertainties by generating a large pool of samples of uncertainties based on their empirical distributions. Each sample is called a scenario, which is one possible realization of the uncertain variables modeled in the grid integration studies.

In the classical approach, the expected value of the stochastic objective function is replaced with a statistical sample average (known as Sample Average Approximation, SAA) of all scenarios. Thus the problem can be directly solved using a deterministic optimization algorithm. The resulting solutions converge to the solution of the true problem if the scenario size is large. If the computational cost is not a concern, then an alternative approach is to solve a multitude of scenarios and assume the average of the resulting solutions as the optimal solution for all scenarios.

However, in grid resource planning problems, the aforementioned two approaches do not guarantee the satisfaction of real energy dispatch needs despite the consideration of scenarios since these scenarios were generated based on the historical data at the time of scenario generation. Moreover, a simple average over all independently solved solutions does not represent the true optimal solution because the average will not satisfy the objective function in every scenario. Our objectives are as follows:

1) eliminate the ambiguity introduced by averaging by solving the optimal solution out of the objective function using all scenarios at once, yielding an optimal solution with greater confidence.
2) incorporate real-time data updates into the proposed stochastic optimization framework, demonstrating the ability to leverage improved short-term solar forecast accuracy and power system data to solve real world grid integration problems involving uncertain events.

1.4 Overview of the Dissertation

The chapters that follow describe the scope of this dissertation with more details of my work in the field of solar grid integration and operations.

To improve the solar forecast accuracy, Chapter 2 presents two methods to estimate local CBH measurements using a single sky imager and distributed ground solar irradiance measurements. The first method is based upon the correlation between ground-observed global horizontal irradiance (GHI) time series and a modeled GHI time series generated from a sequence of sky images geo-rectified to a candidate set of CBH. The estimated CBH is taken as the candidate that produces the highest correlation coefficient. The second method integrates a numerical ramp detection method for ground-observed GHI time series with solar and cloud geometry applied to cloud edges in a sky image. The derived CBHs from both methods are benchmarked against a collocated ceilometer and stereographically estimated CBHs from two sky imagers, yielding promising results.

Further, we improve the CMV capturing algorithm of our in-house CSS for more accurate local cloud speed measurements and faster sampling rate, as presented in Chapter 3. Specifically we utilize all on-board phototransistors to fit a cosine function of cross-correlation time delay versus sensor pair direction, inspired by the geometry of linear cloud edge passing over the array of sensors. Subsequently we integrate the enhanced CSS with angular cloud speed from a sky imager to derive local CBHs at the forecast site. Two months at the UC San Diego campus were used for validation against measurements from nearby Meteorological Aerodrome Reports (METAR) and an on-site ceilometer. Given the low cost of CSS, the fact that the proposed method utilizes the existing sky image measurements to reduce the sky imager forecast errors makes it conveniently and preferably applicable.

Solar forecast errors cause unexpected power ramps. To facilitate the stability of grid operation, an ESS nearby the PV plant is preferred to smooth the power output and maintain the power ramp rates within permissible limits. Chapter 4 presents an analytical relationship between the worst expected ramp rate, cloud speed measurements (from our CSS), and the geometrical layout of the PV plant. This method provides a convenient yet economical approach to enable optimal PV plant and ESS design, which is friendly to simple data input with coarse temporal resolution. The ability of the proposed method to bracket actual ramp rates was assessed over 10 months under different meteorological conditions, and the largest observed ramp was contained with the worst-case estimate.

In Chapter 5, the impact of autocorrelated forecast errors is studied in a net load valley filling problem through EV charging. We use day-ahead persistence forecast error and short-term forecasts from a sky imager system to preserve the autocorrelation structure of forecast errors without statistical approximations. The discrepancies in valley filling solutions between using solar forecasts with and without autocorrelated errors are fully investigated. Throughout one month of simulations, the ability to flatten net load is thoroughly assessed under practical forecast accuracy levels achievable from persistence, sky imager, and perfect forecasts. Last, we demonstrate the feasibility and benefits of incorporating short-term operational forecast to a valley filling problem to mitigate the negative effects of autocorrelated forecast errors existing in the day-ahead persistence forecasts.

Finally, because real world problems always experience unanticipated events in the planning stage, in Chapter 6 we advance the prior deterministic optimization framework to solve the same valley filling problem but consider the uncertainties in PV forecasts and EV charge events that are realized by scenarios, which are subsequently corrected by the real EV and PV measurements in an novel stochastic optimization framework. The common good solution out of all scenarios is solved and executed in each time step of the problem. The new corrective approach is validated against two common approaches adopted in the grid integration studies: 1) dumb charging approach 2) day-ahead averaging approach. Our proposed method exhibits the best valley filling performance while guarantees the real EV charge needs throughout one month of simulations, demonstrating the ability to mitigate the influence of autocorrelated forecast errors in a real word valley filling problem and the operational use.

1.5 Nomenclature

The lists of symbols, abbreviations, subscripts and superscripts throughout this work are given in Tables 1.1, 1.2, and 1.3 respectively.

1.5.1 Symbols

Table 1.1: Symbols used in this work.

Symbol	Description
α	cloud direction in reference to North-South direction
β	EV charge power
$\overline{\beta}$	EV charge power limit
$\boldsymbol{\beta}$	aggregated EV charge power
δ	degree of ramp rate overestimation
Δ	temporal resolution, 15 min
ϵ	noncompliance rate
ζ	mean of timeseries GHI measurements
θ	Zenith coordinates of a pixel in the sky image
θ_{m}	sky imager field of view in degrees from the vertical
θ_{sensor}	angle offset between sensors on the CSS
λ	distance along a ray from observation point of sky imager
μ	cloud velocity scaling factor
ξ	compliance rate of power ramp estimate
σ	ramp rate compliance indicator
ϕ	azimuth coordinates of a pixel in the sky image
ϕ_{ij}	Angle between the line connecting sensors i and j and line (a-c) of CSS
ϕ_s	solar azimuth angle
$\phi_{\perp}$	angle of the CMV perpendicular to the detected cloud edge

Continued on next page.

Table 1.1 – Symbols used in this work, continued from previous page.

Symbol	Description
χ_{ED}	cumulative EV energy demand
χ_0	initial state of charge
$\overline{C}$	Maximum allowed SOC
$\underline{C}$	Minimum allowed SOC
e	notation of nRMSD
f	preferred grid power profile
G	solar generation
ΔH	cloud base height error
h	sky imager elevation or weighting matrix
H	cloud base height
H^{ceil}	CBH measured by ceilometer
H_j	CBH candidate
H^{model}	CBH estimate from the proposed methods
i	index of variable in the context
$\boldsymbol{I}$	identity matrix
j	evaluation window index
k	index of time steps
K	number of samples in 20 minutes at 30 second intervals
k_f	time index of EV departure
k_i	time index of EV arrival
kt	clear sky index
$kt_{\mathrm{in}}/kt_{\mathrm{out}}$	average cloud optical thickness for $\Delta S_{\mathrm{in}}/\Delta S_{\mathrm{out}}$ of the irradiance field
$kt_{\mathrm{min}}/kt_{\mathrm{max}}$	smallest and largest kt from recent history
L	length of PV system
L	substation load demand
m	length of the evaluation window in minutes

Continued on next page.

Table 1.1 – Symbols used in this work, continued from previous page.

Symbol	Description
M	number of modeled CBH values
n_{p}	number of cloud map pixels in one dimension
N	total number of data points/variables in the context
O	sky imager position
p^*	solar power that would be produced by the net area ΔS under clear sky condition
$\hat{p}_{\mathrm{cs}}$	area-normalized clear sky power production
p_{cs}	clear sky power production
p_k	CSS sensor pair number
r	radius of the CSS sensor circle
r	grid net load target in scalar
$\boldsymbol{r}$	grid net load target in vector
R	length of cloud map in one dimension; correlation coefficient
R^2	the goodness of curve fit
s	total time steps in a day = 96
$\boldsymbol{s}$	an intermediate variable = $(\tan\theta_s \sin\phi_s, \tan\theta_s \cos\phi_s, 1)^{\mathrm{T}}$
$\Delta S_{\,\mathrm{in}}/\Delta S_{\,\mathrm{out}}$	outgoing and incoming portion of irradiance field
S	size of PV plant
t	time
Δt_f	forecast time step
t_i	initial timestamp used to compute Δt
Δt_{ij}	time shift of cloud arrival time between CSS sensors i and j
t_0	current time
Δt	cloud travel time
t_{end}	time of last EV departure
t_{start}	time of first EV arrival
u_{x}	x-axis component of $\mathbf{u}$

Continued on next page.

Table 1.1 – Symbols used in this work, continued from previous page.

Symbol	Description
u_{y}	y-axis component of $\mathbf{u}$
$\hat{\mathbf{u}}$	cloud pixel speed [pixel s^{-1}]
$\mathbf{u}$	Cloud pixel speed [m s^{-1}]
$\mathrm{U}_{\mathrm{CSS}}$	CSS cloud speed
$\mathrm{U}_{\mathrm{pixel}}$	USI derived cloud pixel speed
$\mathrm{U}_{\mathrm{USI}}$	USI derived cloud speed
v	cloud speed in scalar
$\boldsymbol{v}$	cloud speed in vector
v_{real}	true cloud velocity vector
$v_{\perp}$	velocity of the CMV perpendicular to the detected cloud edge
w	weighting factor for deriving an actual cloud velocity
W	width of PV system
x	decision variable of the optimization problem
x_{g}	x-axis component of $\mathbf{x}_{\mathrm{g}}$
$\Delta\mathbf{x}$	cloud shadow horizontal shift [m]
$\Delta\hat{\mathbf{x}}$	cloud displacement in the sky image [pixels]
$\Delta\mathbf{x}_{\mathrm{c}}$	cloud projection error [m]
$\mathbf{x}_{\mathrm{c}}$	intersection of cloud motion line and cloud boundary [m]
$\hat{\mathbf{x}}_{\mathrm{c}}$	intersection of cloud motion line and cloud boundary [pixels]
$\mathbf{x}_{\mathrm{g}}$	vector describing ground station location [m]
$\hat{\mathbf{x}}_{\mathrm{g}}$	Vector describing ground station location [pixels]
$\mathbf{x}_{\mathrm{s}}$	intersection of solar beam and cloud map
y_{g}	y-axis component of $\mathbf{x}_{\mathrm{g}}$
z_{g}	z-axis component of $\mathbf{x}_{\mathrm{g}}$

1.5.2 Abbreviations

Table 1.2: Abbreviations used in this work.

Abbreviations	Description
AGL	Above Ground Level
AMSL	Height Above Mean Sea Level
arr	arrival
CBH	Cloud Base Height
CMV	Cloud Motion Vector
cpl	compliance
CSS	Cloud Shadow Speed Sensor
dep	departure
DERiM	Distributed Energy Resource Interconnection Map
ED	Energy Demand
ESS	Energy Storage Systems
EV	Electric Vehicle
GCSE	Geometric Cloud Shadow Edge
GHI	Global Horizontal Irradiance
HDR	High Dynamic Range
LCE	Linear Cloud edge
LCE-CFM	Linear Cloud edge Curve Fitting Method
MBE	Mean Bias Error
MCP	Most Correlated Pair method
METAR	Meteorological Aerodrome Reports
NAM	North American Mesoscale
NL	Grid Netload
nMBE	Normalized Mean Bias Error
nRMSD	Normalized Root Mean Square Difference

Continued on next page.

Table 1.2 – Abbreviations used in this work, continued from previous page.

Abbreviations	Description
NWP	Numerical Weather Prediction
opt	optimization
PDT	Pacific Daylight Time
PST	Pacific Standard Time
PV	Photovoltaic
QP	Quadratic Programming
QRF	Quantile Regression Forests
RHC	Receding Horizon Control
RMSD	Root Mean Square Difference
RR	Ramp Rate
SAA	Sample Average Approximation
SCE	Southern California Edison
SIND	Solar Integration National Dataset
SOC	State of Charge
SPVP	Solar Photovoltaic Program
TSC	Time Series Correlation
UCSD	University of California San Diego
USI	UC San Diego Sky Imager
UTC	Coordinated Universal Time
V1G	Unidirectional EV charging
V2G	Vehicle-to-Grid
WCS-RR	Worst Case Scenario Ramp Rate

1.5.3 Subscripts and Superscripts

Table 1.3: Subscripts and superscripts used in this work.

Subscripts and Superscripts	Description
ceilo	ceilometer
csk	clear sky model
CSS+USI	using CSS cloud speed and USI cloud pixel speed
corrective	using the proposed corrective optimization framework
DA	using day-ahead averaging approach
dumb	using uncontrolled charging approach
model	derived from the proposed model
obs	observations
p	derived from using persistence only
p+perfect	derived from using persistence and perfect solar forecasts
p+USI	derived from using persistence and USI solar forecasts
perfect	derived from using perfect forecasts only
T	transpose

Chapter 2

Cloud Base Height Estimates From Sky Imagery and a Network of Pyranometers

2.1 Introduction

2.1.1 Impact of CBH on Intra-hour Solar Power Forecasting with a Sky Imager

CBH plays a vital role in intra-hour solar power forecasting. For typical mid-latitude solar zenith angles of 45°, a difference of 100 m in CBH causes a 100 m translation of the cloud shadow on the ground. Though CBH induced errors vanish when the solar irradiance is averaged in time and/or in space, considering short time scales, local power output forecast errors can be in excess of 60% of clear sky production levels if a CBH error causes the wrong sky condition (clear or cloudy) to be forecast. Thus, accurate CBH estimation is critical for predicting local power ramps over short time scales. For sky imager solar forecasts that are based on the geometry between the sun, clouds, and ground, CBH is required for mapping the cloud field from sky images to the atmosphere and then projecting to the ground. Specifically, the mapping process consists of three geometry steps: 1) projection of the clouds in the sky image into a plane in the sky (termed "cloud map", see Section 2.2.3) at the CBH; 2) forward motion of the cloud map in time; 3) projection of cloud map onto the ground. Thus, an erroneous CBH leads to three different scaling errors listed below (see the nomenclature for variable definitions and Section 2.3.2 for derivations):

1) The cloud projection error is:

$$\Delta \mathbf{x}_c = \Delta H \cdot (\tan\theta \sin\phi, \tan\theta \cos\phi, 1)^{\mathrm{T}}, \tag{2.1}$$

where $\mathbf{x}_c$ is a 3D-vector describing position error for a given CBH error ΔH, and (θ, ϕ) are respectively the zenith and azimuth pointing angles corresponding to a pixel obtained using pixel coordinates and the camera geometric calibration (e.g. [17]). ΔH linearly scales cloud horizontal position in the radial direction and stretches or shrinks the cloud about a center point at the sky imager, and the scaling error is more sensitive to ΔH at farther spatial distance (outer pixels) caused by the nonlinear effect of $\tan\theta$.

2) Physical cloud velocity error. Because the cloud velocity derived from sky image is in units of pixels, a conversion to actual cloud velocity in units of m/s requires scaling the pixel velocity with cloud base height, resulting in a linear scaling error by ΔH.

3) Cloud shadow projection error. When the cloud map is advected and projected onto the ground, the vertical shift ΔH causes a uniform horizontal shift $|\Delta \mathbf{x}|$ in shadow position following the expression:

$$|\Delta \mathbf{x}| = \Delta H \tan\theta_s, \tag{2.2}$$

which is exaggerated at larger solar zenith angles θ_s. Thus, CBH errors also cause shadows or sunlight to be predicted at locations that are shifted further as the distance from the sky imager increases.

2.1.2 CBH Measurement Techniques

Cloud base height (CBH) can be measured directly using in-situ and remote sensing instruments such as radiosondes [18], ceilometers [19, 20], and satellites [21]. A radiosonde is a battery-powered telemetry instrument package that vertically profiles the atmosphere as the balloon ascends, yielding CBH estimates. Although the CBH measurements from a radiosonde are accurate, the observations are usually taken at most twice daily and at discrete and sparse locations, making them unsuitable for use in intra-hour solar energy forecasting. Ceilometers are the most common CBH observational tool and are regularly installed at airports. It emits a pulsed near-infrared vertical laser beam and measures a vertical profile of atmospheric backscatter from which CBH is derived. Since ceilometer is expensive, it has limited application outside of airports in most countries except in the UK, where ceilometer is a standard component of weather stations.

Indirect CBH measurements using ground based thermal infrared cameras [22, 23] and derived data from remote-sensing techniques such as spectroradiometers [21] are also feasible. The assumption that clouds are blackbodies usually leads to an overestimation of CBH derived by infrared cloud imagers [23]. Satellite-measured cloud top near-infrared radiance [24] or measured cloud top temperature with an atmospheric temperature profile [25] can be used to obtain cloud top height with wide spatial coverage, but CBH is difficult to detect from satellites and time delays in data dissemination limit its application in short-term solar power forecasting. Numerical weather prediction offers another alternative to obtain CBH [26].

CBH can also be obtained from sky imagery. The application of stereogrammetric techniques using two sky imagers was investigated by Allmen and Kegelmeyer [27] and Kassianov et al. [28]. Nguyen and Kleissl ([29], referred to as NK14) further generalized and improved accuracy and computational efficiency of the approach introduced by Kassianov et al. [28] for (binocular) stereographic CBH estimation: a 2D georeferenced projection is used to overlay images from each camera. The CBH is the cloud height associated with the minimum normalized matching error, which implicitly assumes a single cloud layer. More sophisticated stereo-vision techniques can offer a 3D view of the cloud base using the standard technique of matching image patches along epipolar curves [27, 29, 30]. These methods are computationally intensive and provide high spatial resolution CBH within a pair of images. The stereographic method requires at least two sky imagers and accurate geometric calibration of the imaging system (e.g. [17]). Wang et al. [31] and Kuhn et al. [32, 33] demonstrated that CBH can be obtained from a single sky imager and an independent measurement of cloud speed. Because angular cloud speed determined from sky images is proportional to cloud speed and CBH, CBH can also be derived from a collocated cloud speed sensor [34] and sky imager. In [31] and for the same location as in this paper, typical daily root mean square differences were 126 m or 17% of the observed CBH. But the raw (instantaneous) CBH measurements need to be filtered to derive a robust CBH, which makes CBH outputs infrequent (one CBH output every 50 sec for 27 partly cloudy days and every 250 sec for 21 overcast days, on average).

2.1.3 Objectives and Structure of the Paper

CBH is a required input for some sky imager-based short term solar power forecasting variants [35, 36]. The variety of methods presented in Section 2.1.2 can produce accurate CBH information at different temporal and spatial scales, however either equipment or operating costs are prohibitive, or computational requirements are high, or the temporal resolution is insufficient for intra-hour solar power forecasting.

Cameras are ubiquitous and low cost, and nearly every solar power installation has meters for solar irradiance or power. Therefore, existing and low cost infrastructure provides an opportunity to estimate cloud height as an ancillary product if the irradiance distribution on the ground is measured in space and time. Thus, the objective of this work is to provide a low-cost alternative to estimate CBH using such irradiance measurements and a single sky-pointing camera. CBH is estimated using two related methods requiring a single sky imager and irradiance sensors distributed within the footprint of the sky imager, i.e. within the cameras field of view. Both methods are new and have not been been presented before. In the first method, CBH is estimated by correlating ground-observed global horizontal irradiance (GHI) measured using a set of pyranometers with GHI modeled using a sky imager irradiance forecast [35]. Modeled GHI time series are generated from a sequence of sky images geo-rectified to a candidate set of CBH. The second method estimates CBH by matching ramp event timings from pyranometer-measured GHI to cloud shadow arrival times derived from cloud geometry and sun triangularization adapted to sky imagery. The presentation of the latter method provides a new mathematical description of the forecast approach used in [35].

This paper is organized as follows. The measurement equipment, including the sky imaging system and forecasting procedure, is briefly described in Section 2.2. Section 2.3 introduces the CBH estimation methods. Section 2.4 presents the overall performance in a set of 30 days, and then validates CBH from both methods against ceilometer data and the NK14 stereographic method in a case study. Section 2.5 provides detailed discussion regarding the performance and limitation of the proposed methods. Finally Section 2.6 provides conclusions and future work.

2.2 Experimental Data and Sky Imager Forecast Procedure

2.2.1 Ground Measurements

The University of California, San Diego (UCSD) designed and developed a sky imager system specifically for short-term solar power forecasting applications (Fig. 2.1, [37]). The UCSD Sky Imager (USI) features a high-quality image sensor and lens contained in a thermally controlled, compact environmental housing, and capture software employing a high dynamic range (HDR) imaging technique. The USI uses an Allied Vision GE-2040C camera which has a 15.15 × 15.15 mm ON Semiconductor KAI-04022 CCD sensor (originally developed by Kodak). The Sigma 4.5 mm focal length fisheye lens provides a 180 degree field of view with 1748 × 1748 pixels covering the sky hemisphere. Thermal stability of the camera is achieved using two thermoelectric coolers for the entire enclosure, a copper heat sink, and a fan attached to the camera to keep it at the ambient enclosure temperature. The dome on the USI is a 1.6 mm thick, neutral density (ND2) acrylic hemisphere with a UV protective coating. Additional information can be found in [38]. The USI used in this analysis is installed next to one of the six pyranometers shown in Figure 2.2 and Table 2.1.

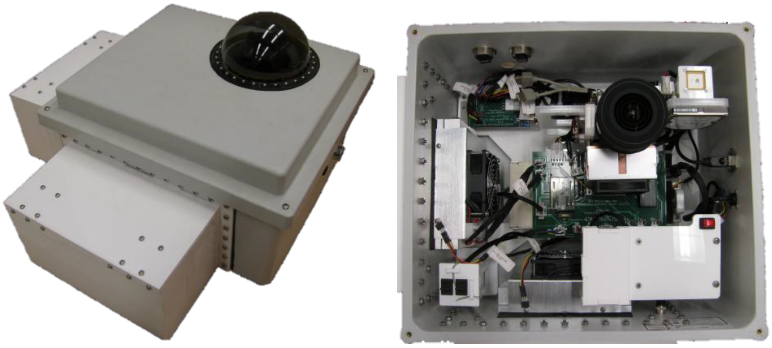
(a) (b)

Figure 2.1: The University of California, San Diego Sky Imager (USI). (a) Outer view showing the enclosure with dome and white radiation shields for the coolers; (b) a top view of the open system showing the components inside the enclosure.

GHI data sampled at 1 Hz is obtained from six weather stations with Li-COR 200SZ pyranometers installed at the locations shown in Figure 2.2 and Table 2.1. In addition, a Vaisala CT25K ceilometer located on EBU2 computes CBH every 20 seconds from backscatter returns. Due to the small sampling area (< 0.1° cone above the ceilometer), the heterogeneity of cloud field, as well as cloud formation and movement, the 20-second ceilometer output is not always representative of the CBH in the field of view of the sky imager. Therefore, consistent with NK14, a 15-minute median filter is applied to ceilometer measurements prior to comparison with the proposed methods.

Table 2.1: Locations of USI and pyranometers used for CBH estimation and their respective distances to the USI (*re-tabulated with permission from* [2]).

Station Name	Latitude	Longitude	Altitude (m)	Distance to USI (m)
USI	32.8722	-117.2410	140	-
BMSB	32.8758	-117.2362	111	603
CMRR	32.8806	-117.2353	111	1074
EBU2	32.8813	-117.2330	101	1257
HUBB	32.8672	-117.2534	24	1288
MOCC	32.8784	-117.2225	103	1857
POSL	32.8807	-117.2350	110	1103

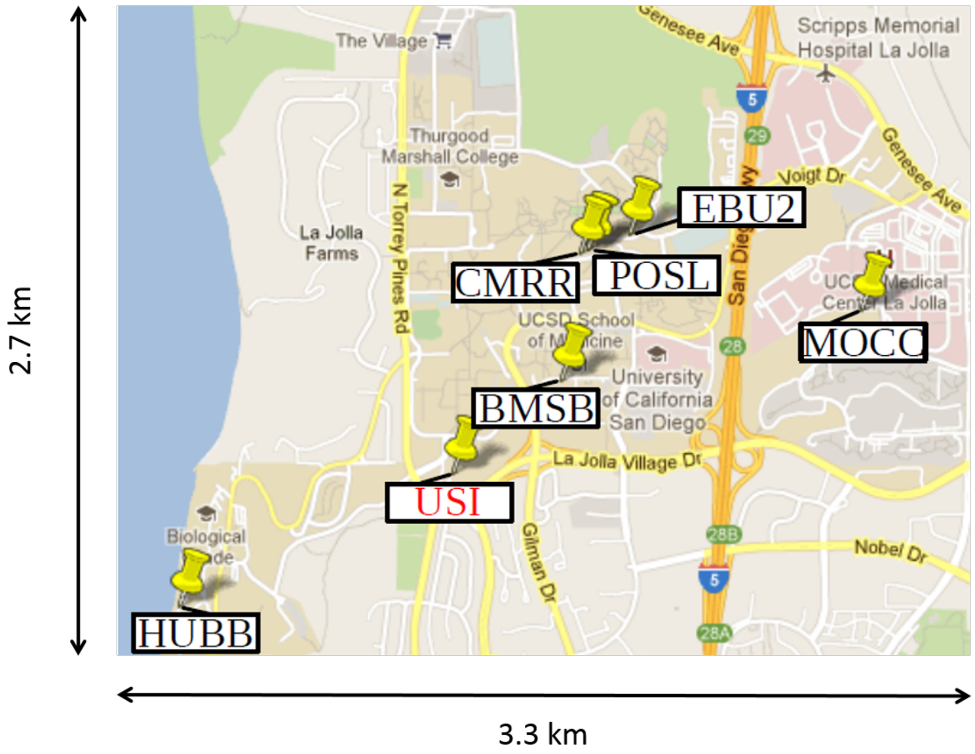

Figure 2.2: Locations of the six pyranometers and the USI on the UCSD campus. The ceilometer is located on EBU2. *Reprinted with permission from* [2]. ©Google Maps.

2.2.2 Evaluation Dataset

The CBH estimation methods are evaluated using two different sets of CBH measurements: (1) an on-site ceilometer on 33 days and (2) the NK14 2D stereography method on 3 days. Thirty-three cloudy days from 2012 to 2016 were selected based on the following criteria:

1) Data availability from sky imager, ceilometer and pyranometers.
2) Cloudy conditions: clear and rainy days were excluded.
3) Cloud type: opaque clouds such as stratocumulus, cumulus, and stratus, since they are most relevant to solar forecasting.
4) Cloud height predominantly less than 1000 m. Four days were chosen with cloud heights greater than 1000 m.
5) Lack of rain: less than 2 hours of rain.

Finally, only time periods with solar zenith angles less than 75° are considered. Moreover, during an intensive operating period in 2012, two sky imagers were installed, which allowed 2D stereography to be applied to four days, as reported in NK14. December 14, 2012 was characterized by broken stratocumulus clouds above a few cumulus clouds. On December 26, a single layer of low scattered cumulus clouds was observed. December 29 was overcast with stratus clouds. Jan 1, 2013 analyzed in NK14, was not included in this paper because several station outages limited GHI measurements to only two stations.

2.2.3 Sky Imager Forecast Procedure

The USI can be used to geolocate clouds, to measure cloud angular velocity, and to track cloud motion [35, 39]. These measurements are then used to forecast future cloud locations up to 15 minutes ahead. The forecast procedure is outlined in the flow chart of Fig. 2.3. A brief overview of the USI forecast procedure is given in the remainder of this section. For more information, the reader is referred to [2, 35, 37, 40]. Similar sky imager systems and forecast procedures can be found in [36, 41, 42].

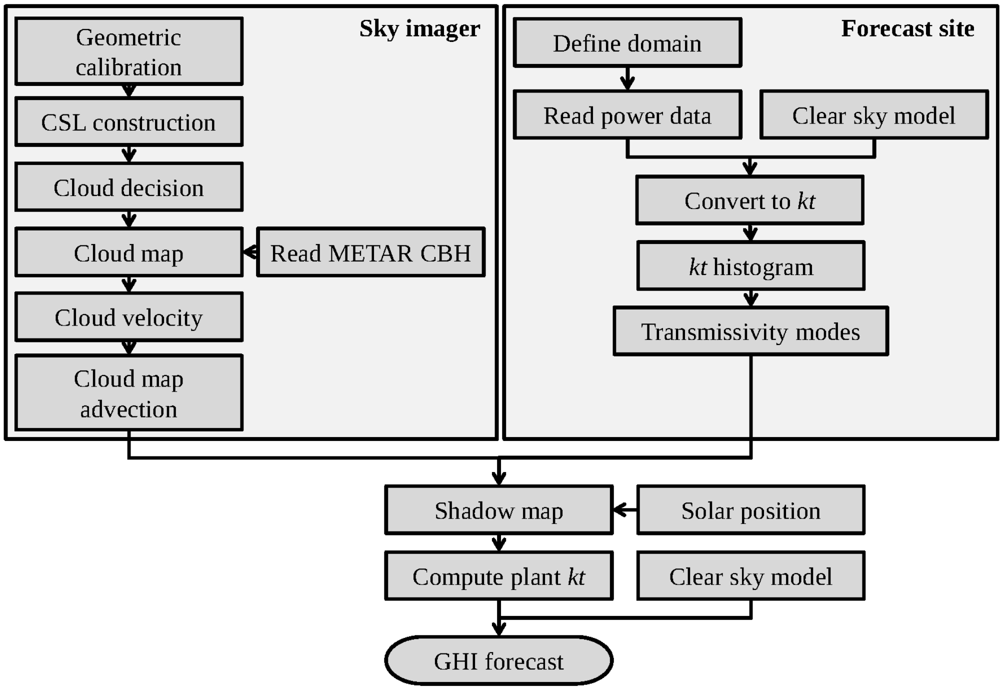

Figure 2.3: Flowchart of USI forecast procedure. Sky image processing (left) is combined with the clear sky index from local ground observations (right) to produce spatial irradiance forecasts (*reprinted with permission from* [2]).

Based on images taken every 30 seconds, cloudy pixels are detected and using lens-camera geometry, images are transformed to a rectified planar grid [27]. CBH is then used to register each pixel to a latitude, longitude, and altitude (geo-rectification [35]). The resulting geo-referenced map of clouds is termed the "cloud map", which is a planar mapping of cloud position at a specified altitude above the forecast site. The cloud map at the current time $t = t_0$ yields the "nowcast", while future cloud positions ($t > t_0$) are determined through cloud advection at discrete time steps delivering the forecast. The ability to resolve the horizontal cloud structure near the horizon is limited due to perspective effects (look vectors are nearly parallel to horizontal cloud base) and due to the longer distance to the clouds, causing a single pixel to subtend a much larger projected area. Both of these factors introduce errors when using the perimeter of the cloud map (more discussion in Section 2.5.3).

Cloud pixel velocity is obtained by applying a cross-correlation method to the red-blue ratio of two consecutive sky images. The cloud speed $\mathbf{u}$ [m s^{-1}] is then calculated from cloud pixel velocity $\hat{\mathbf{u}}$ [pixel s^{-1}] using a scaling factor μ, which is a function of CBH as:

$$\mathbf{u} = \frac{1}{\mu(H)}\hat{\mathbf{u}} = \frac{1}{\mu}\frac{\Delta\hat{\mathbf{x}}}{\Delta t_f}, \tag{2.3}$$

where $\Delta\hat{\mathbf{x}}$ is the cloud displacement in the image, Δt_f is the image capture interval (here also equal to the forecast time step), and the ˆ indicates units of pixels. Equation 2.9 in Section 2.3.2 gives the expression for $\mu(H)$. The cloud velocity is then used to advect the planar cloud map to generate cloud position forecasts for each forecast horizon. Since the distance from the sun to the Earth is much larger than the distance from the clouds to the Earth (i.e. the direct solar beam for locations on Earth is essentially parallel), cloud shadow speed is essentially identical to cloud speed.

The forecast procedure used in this work is developed for a single sky imager. The default CBH source for a single sky imager is meteorological aerodrome reports (METAR). METAR stations, which use a ceilometer, report high quality CBH data but are limited in temporal resolution (typically hourly reports) and are spatially sparse. Therefore, spatial variability in cloud cover causes differences between CBH at the sky imager location and the nearest METAR station. These limitations are the main motivation for this work.

2.3 Methods for CBH Estimation

A Time Series Correlation (TSC) method and a Geometric Cloud Shadow Edge (GCSE) method will be introduced in this section. Both methods only require a single sky imager and time-synchronized measurements of GHI or solar power output at surrounding stations. For TSC, at each ground station GHI is simulated for a set of CBHs and cross-correlated with GHI measurements at the corresponding ground sites. For GCSE, cloud arrival and departure times are determined from the GHI time series using ramp detection. CBH is then derived by matching these detected cloud arrival times with cloud arrival times simulated using USI cloud imagery and cloud position forecasts.

2.3.1 TSC Method

Most of the large-magnitude variability in GHI time series is introduced by cloud shadows approaching or departing a location. In fact, as described in [31], cloud shading events implicitly contain CBH information: the duration of the shading event is proportional to the length of cloud (and cloud shadow) in the direction of cloud motion (cloud velocity assumed to be constant). Using an independent cloud speed measurement (e.g. [43, 44] along with cloud pixel speed estimated in the USI forecast procedure (Section 2.2.3), CBH can be derived based on Eqn. (2.1).

TSC estimates CBH using a grid search performed over a set of candidate CBH values H_j. For each ground measurement station (indexed by $i = 1 \ldots N$), GHI is modeled over the last 20 minutes for each H_j ($j = 1 \ldots M$) using USI nowcasts from a 20 min sequence of geo-rectified sky images captured at sampling rate of 30 sec (i.e. a total of $K = 41$ image samples). For each station, the correlation coefficient R_{ij} is computed between each modeled GHI time series $\mathrm{GHI}_i\left(t; H_j\right)$ and the observed GHI time series $\mathrm{GHI}_i^{\mathrm{obs}}(t)$:

$$R_{ij} = \frac{1}{K\sigma_i^{\mathrm{obs}}\sigma_{ij}} \sum_{k=1}^{K} \left[\mathrm{GHI}_i^{\mathrm{obs}}\left(t_o + k\Delta t_f\right) - \zeta_i^{\mathrm{obs}}\right]\left[\mathrm{GHI}_i\left(t_o + k\Delta t_f; H_j\right) - \zeta_{ij}\right], \tag{2.4}$$

where ζ_i^{obs} and ζ_{ij} are the means of $\mathrm{GHI}_i^{\mathrm{obs}}(t)$ and $\mathrm{GHI}_i\left(t; H_j\right)$ over the K samples, respectively, and σ_i^{obs} and σ_{ij} are the corresponding standard deviations. For each of the N stations, this yields M correlation coefficients. The coefficients are then averaged across stations for each value of H_j to generate a correlation score for each CBH candidate:

$$R_j = \frac{1}{N} \sum_{i=1}^{N} R_{ij}. \tag{2.5}$$

Initially, a weighting scheme using the inverse sky imager to weather station distance was applied, however performance was similar, and thus only a simple average is used here. After R_j has been computed for all CBH candidates H_j, the CBH candidate corresponding to the largest correlation score R_j is selected as the CBH estimate.

Theoretically, TSC can yield a CBH every 30 seconds (i.e. sampling rate of sky images) because a correlation can always be established. However, since CBH in clear or rainy conditions is irrelevant to solar forecasting, TSC results with correlation coefficients below 0.5 are excluded. Moreover, the performance of TSC degrades in homogenous cloud cover or clear conditions because the variations in the time series are small and correlation between modeled and measured GHI is expected to be similar for all CBH candidates. As further discussed in Section 2.5, under these conditions, using the maximum correlation is not a reliable way to estimate CBH. Fortunately, for solar power forecasting applications, in cases of uniform sky cover, the impact of CBH error is mitigated.

2.3.2 GCSE Method

2.3.2.1 Cloud Shadow Geometry

The coordinate system origin is the sky imager position. The coordinate axes are aligned such that *x* is positive east, *y* is positive north, and *z* is positive up, and earth curvature effects are ignored. The location of a ground station in this coordinate system is then $\mathbf{x}_\mathrm{g} = \left(\mathrm{x_g}, \mathrm{y_g}, \mathrm{z_g}\right)^\mathrm{T}$, where T indicates transpose. The ray pointing to the sun from point $\mathrm{x_g}$ can be parameterized as:

$$\mathbf{x}_s(\lambda) = \mathbf{x}_\mathrm{g} + \lambda \begin{bmatrix} \sin\theta_s \sin\phi_s \\ \sin\theta_s \cos\phi_s \\ \cos\theta_s \end{bmatrix}, \tag{2.6}$$

where θ_s is the solar zenith angle, ϕ_s is the solar azimuth angle, and λ is the distance from $\mathbf{x}_\mathrm{g}$ towards $\mathbf{x}_s$ in meters. Assuming a planar layer of clouds, we can compute the intersection of $\mathbf{x}_s(\lambda)$ with the clouds by setting the z-coordinate to the CBH above the sky imager: $\mathrm{x}_{s,z}(\lambda) = \mathrm{x}_{g,z} + \lambda \cos\theta_s = H - h$, where H is the cloud base height and h is the height of the sky imager (both heights referenced above ground level [AGL]). This gives $\lambda = \left(H - h - \mathrm{x}_{g,z}\right)\sec\theta_s$; the point $\mathbf{x}_s$ in the cloud layer is then:

$$\mathbf{x}_s = \mathbf{x}_\mathrm{g} + \left(H - h - \mathrm{x}_{g,z}\right) \begin{bmatrix} \tan\theta_s \sin\phi_s \\ \tan\theta_s \cos\phi_s \\ 1 \end{bmatrix} = \mathbf{x}_\mathrm{g} + \left(H - h - \mathrm{x}_{g,z}\right)\mathbf{s}, \tag{2.7}$$

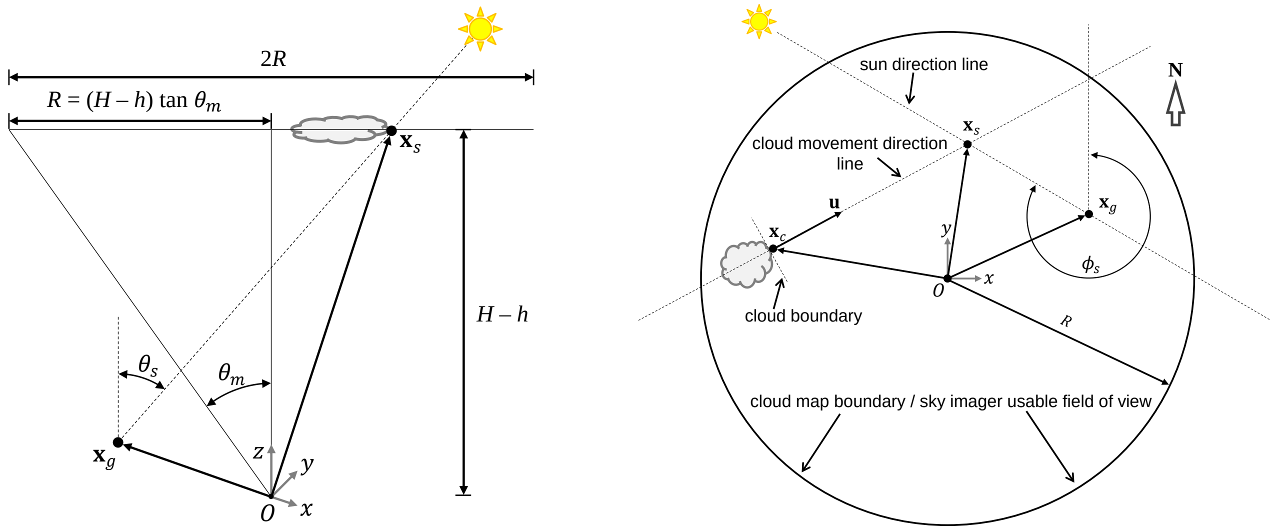

(a) (b)

Figure 2.4: (a) Cross-section and (b) plan view of the geometric relationship between sky imager position O, a ground station $\mathbf{x}_g$ and the cloud intersection point $\mathbf{x}_s$ (Note: to improve the illustration clarity, $\mathbf{x}_g$ and $\mathbf{x}_s$ are shown in different relative locations in each subfigure).

where $\mathbf{s} = (\tan\theta_s \sin\phi_s, \tan\theta_s \cos\phi_s, 1)^{\mathrm{T}}$. Figure 2.4a illustrates this geometric relation when sky imager position O, $\mathbf{x}_{\mathrm{g}}$, and $\mathbf{x}_s$ are coplanar in azimuth (although in general they are not coplanar) and Fig. 2.4b shows a top-down-view of the geometric configuration.

Depending on the spatial configuration of the cloud field, at any given time clouds may or may not be present at point $\mathbf{x}_s$, which is the point at which clouds must be present to shade the station located at $\mathbf{x}_g$ (To shade a sensor at $\mathbf{x}_g$, clouds may actually be anywhere along $\mathbf{x}_s(\lambda)$, but again we are assuming a planar cloud field at height H). Assuming a constant cloud velocity $\mathbf{u} = \left(\mathrm{u}_x, \mathrm{u}_y, 0\right)^{\mathrm{T}}$, we estimate the current position of a cloud $\mathbf{x}_c$ that will move to point $\mathbf{x}_s$ in Δt seconds:

$$\mathbf{x}_c(H, \Delta t) = \mathbf{x}_{\mathrm{s}}(H) - \Delta t \mathbf{u}, \tag{2.8}$$

where the meaning of the input argument to $\mathbf{x}_s$ has been changed from slant distance λ to CBH H following Eqn. (2.5). Hereinafter, Δt is referred to as the cloud travel time (Eqn. 2.9).

To search the image for clouds that could potentially cause shadowing of the sensor, we search the surface $\mathbf{x}_c(H, \Delta t)$ parameterized by H and Δt. This requires the following conversion from space coordinates to image coordinates. The usable field of view of the sky imager for cloud imaging is $2\theta_m$, and the corresponding width of the cloud map is $2R = 2(H-h)\tan\theta_m$. The number of pixels spanning the cloud map diameter is set to n_p (The cloud map is an "undistorted" plane-projected version of the original distorted image, taking into account the camera calibration). The projection requires interpolation of the image and n_p can be set to a suitable value based on the footprint of the sky image. In this paper, $n_\mathrm{p} = 1251$ is the default value used in our sky imager forecast algorithm. The conversion from units of meters to pixels is then:

$$\mu(H) = \frac{n_\mathrm{p}}{2(H-h)\tan\theta_m}\left[\frac{\text{pixels}}{\text{meter}}\right]. \tag{2.9}$$

Combining Eqns. (2.7) and (2.8) and multiplying by μ gives:

$$\hat{\mathbf{x}}_c(H,\Delta t) = \hat{\mathbf{x}}_\mathrm{g} + \frac{n_\mathrm{p}}{2}\frac{\left(H-h-\mathrm{x}_{\mathrm{g,z}}\right)}{(H-h)\tan\theta_m}\mathbf{s} - \Delta t\hat{\mathbf{u}}, \tag{2.10}$$

where the ˆ indicates coordinates have been converted to units of pixels. When $|\hat{\mathbf{x}}_c(H,\Delta t)| > \frac{n_\mathrm{p}}{2}$, the cloud point is outside of the cloud map and the cloud state cannot be retrieved (i.e. it is outside of the sky imagers usable field of view). Additionally, we only consider cases where $|\hat{\mathbf{x}}_\mathrm{g}| \leq \frac{n_\mathrm{p}}{2}$, as $|\hat{\mathbf{x}}_\mathrm{g}| > \frac{n_\mathrm{p}}{2}$ occurs if the station is outside the cloud map because H is too low, shrinking the cloud map (i.e. R is small). When the latter criterion is not met, it is possible that the shadow projection of the cloud map may still encompass the station, however for $|\hat{\mathbf{x}}_\mathrm{g}| > \frac{n_\mathrm{p}}{2}$ the station is "far" and the reliability of the results is questionable. Interestingly, setting $\hat{\mathbf{x}}_\mathrm{g}$ to the sky imager location $(0,0,0)^\mathrm{T}$ shows that forecasts at the sky imager location do not depend on CBH. Using Eqn. (2.10), we can solve for the cloud travel time:

$$\Delta t = \frac{1}{|\hat{\mathbf{u}}|}\left|\hat{\mathbf{x}}_\mathrm{g} - \hat{\mathbf{x}}_\mathrm{c}(H,t) + \frac{n_\mathrm{p}}{2}\frac{\left(H-h-\mathrm{x}_{\mathrm{g,z}}\right)}{(H-h)\tan\theta_m}\mathbf{s}\right|. \tag{2.11}$$

2.3.2.2 Ramp Detection

A ramp detection procedure is used to determine the start of down ramps in the ground-observed GHI data. Down ramps are associated with cloud edge arrival times, and thus locating down ramps allows timing the expected passage of a cloud edge over the station. Precise ramp timings require a high sampling rate, and in this analysis a 1 Hz dataset is used.

Figure 2.5 presents a case study where the detection process described below is applied to a GHI time series. Ground-observed GHI, sampled at 1 Hz at each station, is converted to clear sky index using the Kasten clear sky model (improved and described by [45]). A Gaussian filter is then applied to smooth the data (top subplot). The size of the filtering window is selected as 10 min, an empirical tradeoff value between effective noise reduction and signal shape preservation. Consistent with convention, the filter width is set to 3 standard deviations which comes out to 100 seconds. At each time step in the smoothed series, we compute the maximum difference in clear sky index between the current data point and any subsequent point within 90 seconds yielding a time series of maximum ramp magnitudes (blue and green curve in the bottom subplot). All such ramp points with a clear sky index change in magnitude of greater than 0.3 (30% clear sky index ramp) are collected (red). Then local extrema are located with the MATLAB implementation of Findpeaks[1], which usually gives a single time instant corresponding to the start time of each large ramp (black lines). When more than one ramp extremum is found per ramp, the point with greater associated ramp magnitude is selected. Finally, because sometime large ramps exhibit non-monotonic characteristics, causing the detected start time to deviate, the ramp event start time is corrected if there is a local maximum in the clear sky index within 5 seconds from the detected time instant can be found (refer to Section 2.5.2 for more details).

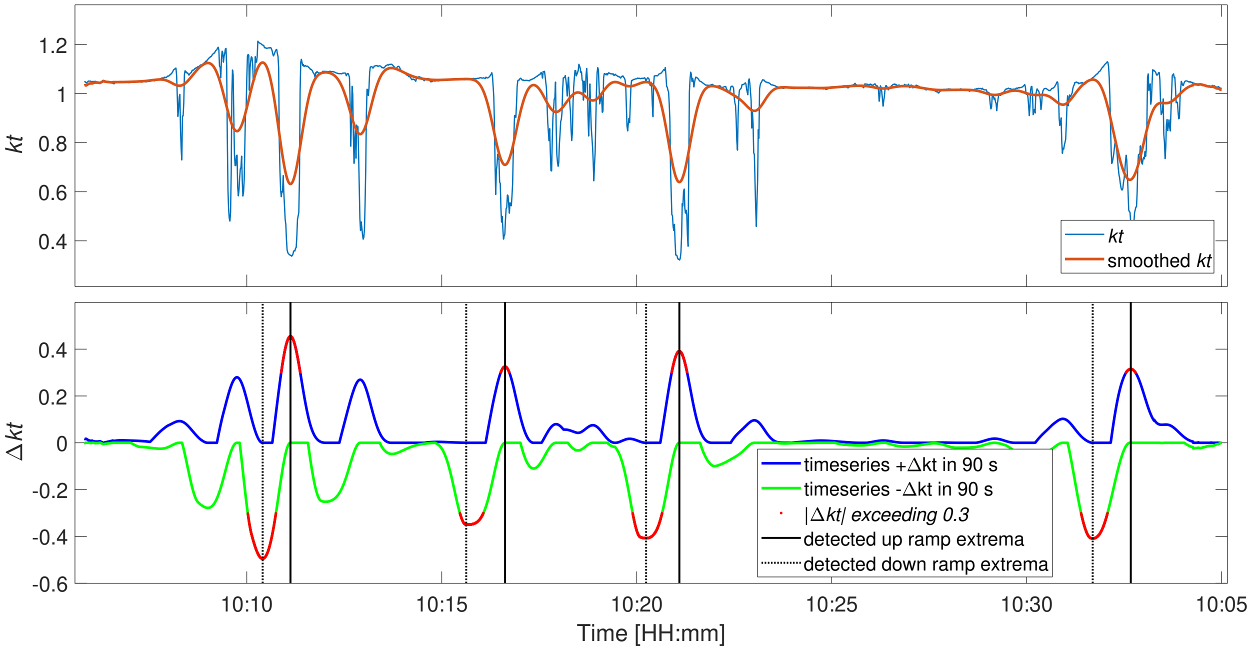

Figure 2.5: Illustration of the procedure to detect ramps in the normalized time series GHI (*kt*). Top: The time series *kt* is smoothed by a Gaussian filter with filter width of 10 min and standard deviation of 100 sec. Bottom: The maximum difference in *kt* between within time window of 90 sec is computed, resulting in time series ramp points of (blue and green). The points with an associated ramp magnitude of less than 0.3 are excluded and the remaining points are kept (red). The local extrema are located by MATLAB implementation of Findpeaks (black and dashed black line).

Figure 2.6 illustrates the outcome of a real execution of the procedure in Figure 2.5 to both BMSB

---

[1] https://www.mathworks.com/help/signal/ref/findpeaks.html, accessible as of March, 2019

and EBU2 stations (refer to Figure 2.2 for station name and location). At the current time t_0 = 13:06:00 LST, the prior 10 minute GHI data is collected, and $t_0 - 10$ min is defined as the initial time t_i = 12:56:00 LST. Since more than one large down ramp occurred in the ten minute window, the down ramps closest in time to t_i are selected and the ramp start time instants are determined. The detected down ramp start times are t = 12:59:46 LST for BMSB and t = 13:01:19 LST for EBU2 (red dots) yielding cloud travel times defined in Eqn. (2.11) of $\Delta t = |t - t_i|$ = 226 s for BMSB and 319 s for EBU2, respectively.

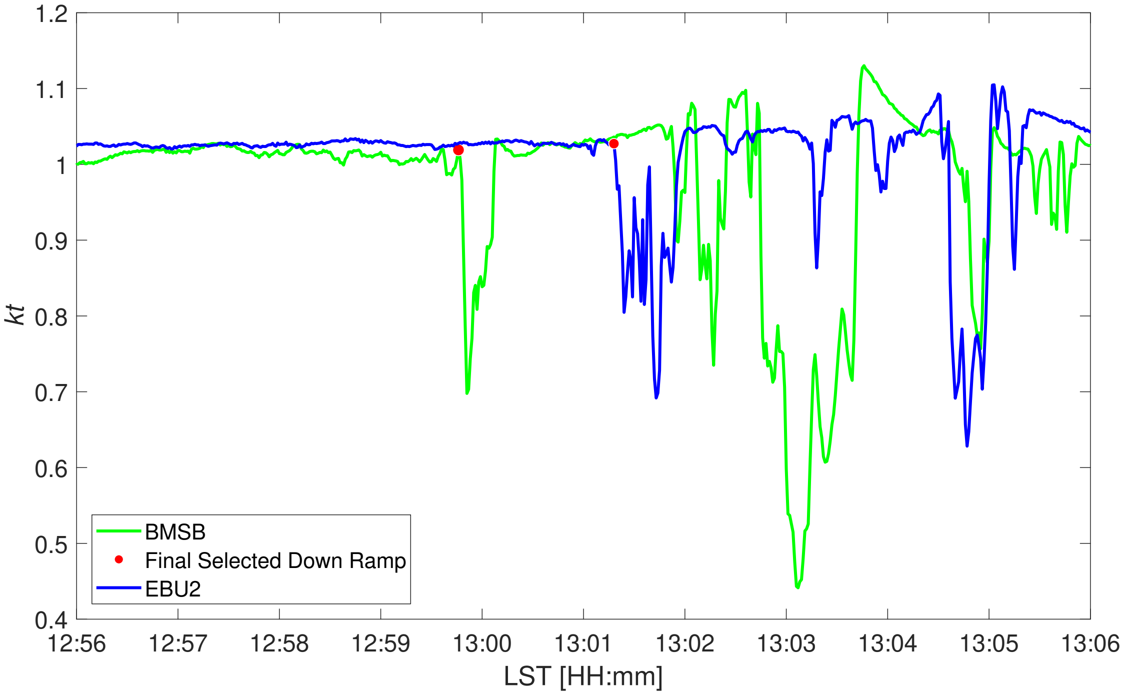

Figure 2.6: Illustration of the proposed ramp detection procedure to determine the start time of the down ramp for the BMSB and EBU2 station at initial time t_i = 12:56:00 LST on May 19, 2014. The start time of the final selected down ramp closest to t_i is marked as a red dot.

2.3.2.3 Using the H-Δt Map to Estimate CBH

Equation 2.10 provides an expression for cloud map pixel location as a function of cloud base height H, cloud travel time Δt, and cloud velocity $\mathbf{u}$. The cloud state of the cloud map at location $\hat{\mathbf{x}}_c(H, \Delta t)$ can be clear sky, thin cloud, or thick cloud. The range of H considered in our analysis is 300 m to 2500 m in 50 m increments based on the common CBH range for coastal Southern California. Δt is varied from 0 to 10 min in 5 sec increments. Using the grid of H and Δt (velocity is assumed constant during Δt), the pixel position in the cloud map is computed (Eqn. 2.10), and the cloud state is extracted. This results in a transformation of the cloud map which we call H-Δt map.

Figure 2.7 visualizes the H-Δt map for the time window and GHI data corresponding to Figure 2.6. For illustration purposes, the CBH range is set to 500 m to 1800 m in 10 m increments and Δt varies from 0 to 350 sec in 5 sec increments. The vertical yellow lines are placed at Δt_{BMSB} and Δt_{EBU2}, indicating the respective station cloud travel times (as determined in Fig. 2.6). CBH candidates are obtained from the H-Δt map by searching for cloud condition transitions around lines of constant Δt_i. The most commonly occurring CBH candidate across all stations is selected as the CBH estimate. If two or more CBH candidates are equally common then they are averaged. If none of the stations returns a CBH candidate, no CBH estimate is generated. Red crosses in Fig. 2.7 indicate the CBH candidates are 620 m for BMSB, and 540 m and 660 m for EBU2. Thus, the CBH candidates from the two stations are averaged to be 606 m. The concurrent ceilometer reading at 12:59:00 LST indicates a single cloud layer at 610 m.

Note that the Δt axis scales linearly with the cloud velocity, so uncertainty in $\mathbf{u}$ contributes directly to uncertainty in Δt. This and other potential errors in down ramp timing estimates (see discussion in Section 2.5.2) justify extending the cloud condition transition search to a search window of 60 seconds around lines of constant Δt_i. While this process induces more CBH candidates, it reduces the times when no CBH estimate is output.

An alternate, more intuitive presentation of the H-Δt map is the cloud shadow distribution plan view in Fig. 2.8, generated using the t_i cloud map advected at cloud velocity $\mathbf{u}$ to time $t_i + \Delta t$. In each subplot of Fig. 2.8, a shadow is just about to pass over the stations (note the shadow adjacent to each red dot and the direction of cloud movement). The plan view gives the cloud shadow distribution for all stations at a single $(H, \Delta t)$ pair, whereas the H-Δt map gives the possible cloud condition for a single station at a range of $(H, \Delta t)$.

2.4 Results

2.4.1 Median Filtering and Error Metrics

The non-uniform ceilometer measurements are first resampled to the TSC and GCSE time steps through nearest neighbor interpolation. A sliding 15 minute median filter is then applied to the raw output of TSC, GCSE, and the resampled ceilometer measurements. To quantify the differences between the proposed methods and the ceilometer output, the mean bias difference (MBD) and the root mean square difference (RMSD) were used:

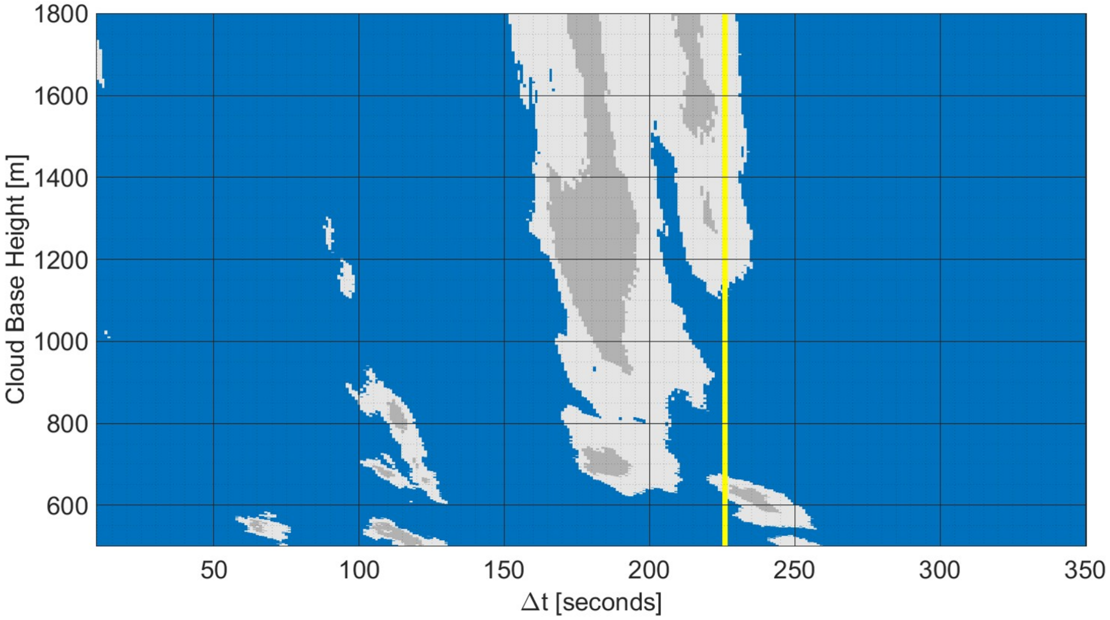

(a)

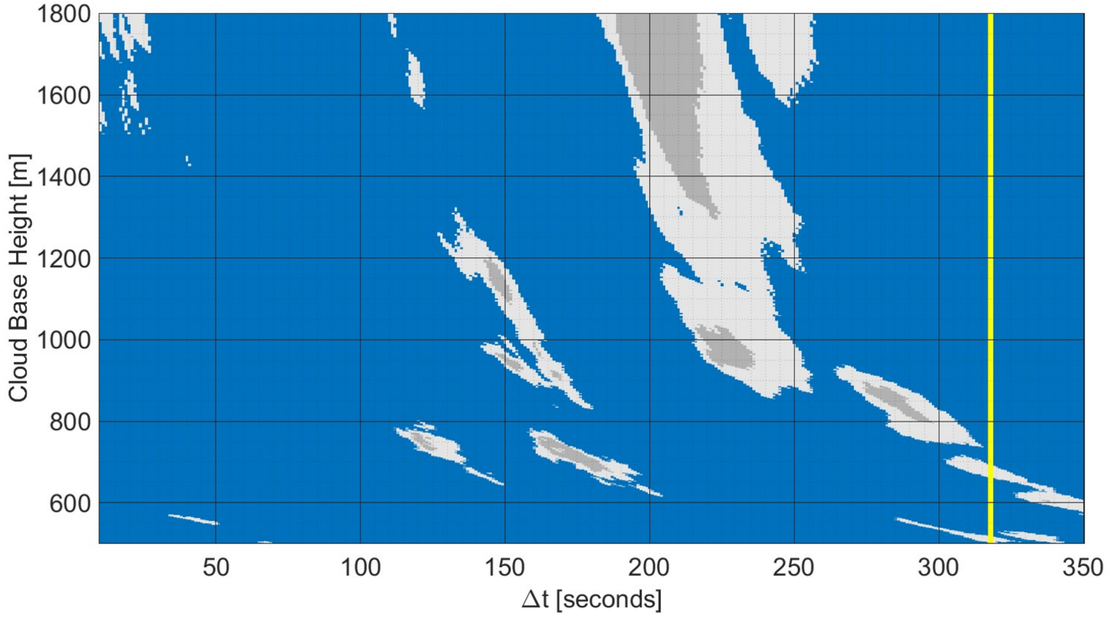

(b)

Figure 2.7: CBH versus arrival time (or H-Δt map) for the initial time $t_i = 12:56:00$ LST on May 19, 2014. Ramp events, indicated by the vertical yellow lines, were detected at $\Delta t = 226$ s and $\Delta t = 319$ s for (a) BMSB and (b) EBU2 ground stations, respectively (see Figure 2.2 for locations). Left-to-right transitions from clear to cloudy (i.e. down ramps) along the yellow line indicate CBH candidates (red cross). Blue, white, and grey colors represent clear sky, thin clouds, and thick clouds, respectively.

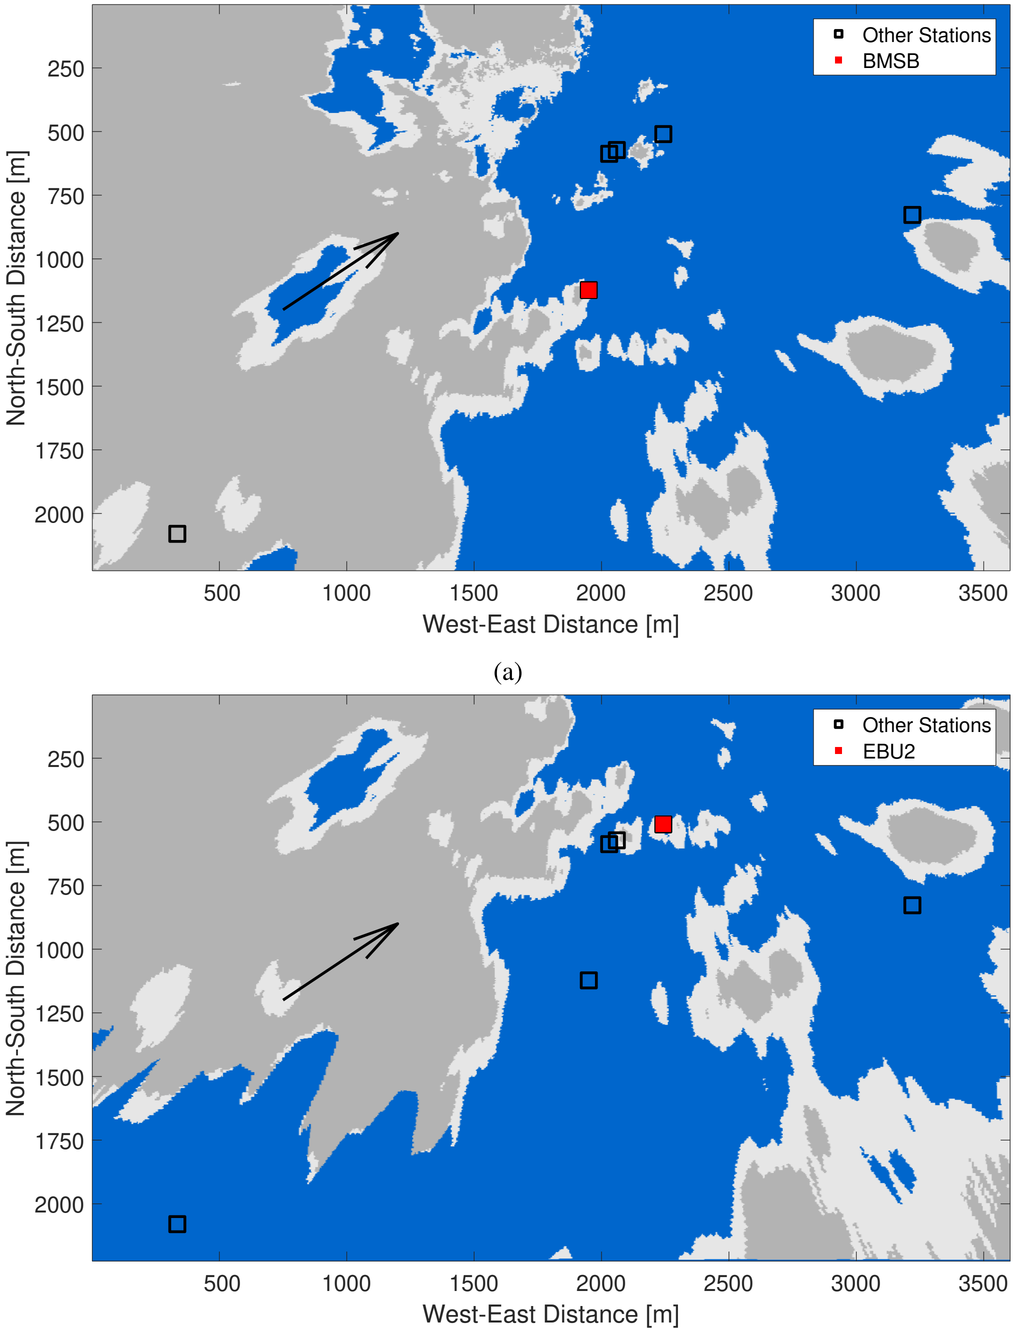

(a)

(b)

Figure 2.8: Advected cloud shadow map generated from the sky image taken at $t_i = 12:56:00$ LST on May 19, 2014 using arrival time of the down ramp of (a) $\Delta t_{\mathrm{BMSB}} = 226$ s and (b) $\Delta t_{\mathrm{EBU2}} = 319$ s and a CBH of 606 m determined in Fig. 2.7. BMSB (a) and EBU2 (b) ground stations are shown as red filled squares. Empty squares represent the five other ground stations. The arrow indicates the cloud motion vector, showing the cloud shadows moving towards northeast. The arrow magnitude indicates the distance traveled by a cloud in 30 s. Blue, white, and grey colors represent clear sky, thin cloud, and thick cloud, respectively.

$$\mathrm{MBD} = \frac{1}{N}\sum_{n=1}^{N}\left(H_n^{\mathrm{model}} - H_n^{\mathrm{ceil}}\right), \tag{2.12}$$

$$\mathrm{RMSD} = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\left(H_n^{\mathrm{model}} - H_n^{\mathrm{ceil}}\right)^2}, \tag{2.13}$$

where N is the total number of data points, H_n^{model} is the CBH from the TSC and GCSE methods, and H_n^{ceil} is the corresponding ceilometer measurement at time index n. MBD and RMSD are divided by the daily average CBH to obtain a normalized MBD (nMBD) and normalized RMSD (nRMSD). Normalization provides a better comparison across days RMSD is expected to be proportional to the true cloud height), whereas the un-normalized metrics give a better characterization of CBH accuracy for solar power forecasting. Periods with rain, either falling or droplets remaining on the imager, were excluded from the evaluation since neither ceilometer nor sky imager methods perform reliably under those conditions. Rainy periods are shaded in yellow.

2.4.2 Evaluation over 30 Cloudy Days

The performance for 30 days, spanning all seasons and multiple cloud types is summarized in Table 2.2 and presented in Fig. 2.12. Stratocumulus and cumulus clouds were most common on the selected days. Only four of the 30 days had CBHs exceeding 1000 m, so the evaluation provided is predominantly for low cloud conditions consistent with the dominant climatology of coastal Southern California. Overall, TSC outperformed GCSE for this extended data set, with TSC achieving an RMSD of 133 m versus 163 m for GCSE. The standard deviation of daily RMSD for TSC was 72.3 m versus 92.9 m for GCSE, indicating the performance of TSC is more consistent across days. TSC had a small positive bias, versus a small negative bias for GCSE.

The number of CBH values reported per day varies markedly between TSC and GCSE. GCSE yields no result if there are no clouds detected that will shade the station. This will occur during periods with sufficiently homogeneous cloud conditions and specifically periods with clear or overcast conditions along the cloud motion vector such that no CBH candidates are generated for the available ground stations, i.e. there are no clouds within $\Delta t_i \pm \sigma_t$ for each station i in the H-Δt map. Additionally, GCSE cannot generate CBH if no down ramps are located. These limitations cause GCSE to issue 34% less CBH than TSC averaged over 30 days.

2.4.3 Comparison to NK14 on Select Days

Table 2.3 and Fig. 2.9 present further validation against NK14 on three days. While it produces scattered raw results, TSC captures the major CBH transition on all three days. In contrast, the CBH estimates from GCSE are not as scattered likely because of the internal quality control that requires CBH output consensus between stations. RMSD errors for TSC and NK14 are less than 300 m RMSD and 20% (nRMSD) averaged over the three days. GCSE, however, has RMSD and nRMSD of over 400 m and 27%, respectively, performing consistently worse than the other two methods. The MBD and nMBD show that the bias of GCSE is almost twice that of TSC for these three days. Note that nRMSD seems higher for Dec 26 on both methods; however, the absolute error on Dec 26 is not unusual and the large error can be attributed to the normalization by a smaller CBH. NK14 beats both TSC and GCSE on all three days, though the performance of TSC is close to that of NK14.

2.5 Discussion

2.5.1 TSC Performance

TSC computes the average correlation coefficient between 20 minutes of measured and modeled GHI across several ground stations. Correlation coefficients are computed for a range of CBH values, and the CBH corresponding to the maximum correlation is output. While the bias of the method over 30 days is small at 1.3% nMBD, the random error is significant at 18.9% nRMSD. Although this may seem high, it is within 3 percentage points of the stereographic method. The following subsections highlight different factors affecting the performance of TSC method.

GHI Sampling and Correlation

Forty-one (41) samples are used to compute the correlation coefficients which are subsequently averaged across stations. The 20 minute sample duration may be insufficient to yield a reliable CBH estimate, but is chosen empirically to allow the method to be sufficiently dynamic to track intra-hour changes in CBH. Increasing the time window may reduce the estimate variance at the expense of being unable to react to rapid CBH changes. An alternative to increasing the sample duration is to decrease the sampling (image

Table 2.2: The monthly average (“Avg.”) and standard deviation (Std.) of the daily error metrics weighted according to the number of data points. Ceilometer daily averages are reported as “Mean CBH”. Refer to Table 2.4 for error metrics by day.

Date	Mean CBH	TSC method					GCSE method				
		MBD	nMBD	RMSD	nRMSD	No. Points	MBD	nMBD	RMSD	nRMSD	No. Points
	[m]	[m]	[%]	[m]	[%]	[-]	[m]	[%]	[m]	[%]	[-]
Avg.	715	2.2	1.3	132.8	18.9	581	-33.9	-3.1	162.6	20.8	382
Std.	217	72.0	-	72.3	-	160	101	-	92.9	-	176

Table 2.3: Comparison of cloud base height estimates. Ceilometer daily averages are reported as “Mean CBH”. The average for each column (“Avg.”) is weighted according to the number of data points in each day. Rainy periods are excluded.

Date	Mean CBH	TSC method				GCSE method				NK14	
		MBD	nMBD	RMSD	nRMSD	MBD	nMBD	RMSD	nRMSD	RMSD	nRMSD
	[m]	[m]	[%]	[m]	[%]	[m]	[%]	[m]	[%]	[m]	[%]
14-Dec	1814	88	4.9	291	16	-223	-12.3	394	21.7	262	14.3
26-Dec	1164	-128	-11	288	24.8	-323	-27.8	399	34.3	206	17.7
29-Dec	1625	-103	-6.3	299	18.4	169	10.4	440	27.1	272	16.8
Avg.	1534	-47.8	-4	293	19.50%	-104.3	-8.3	413	27.4	246	16.3

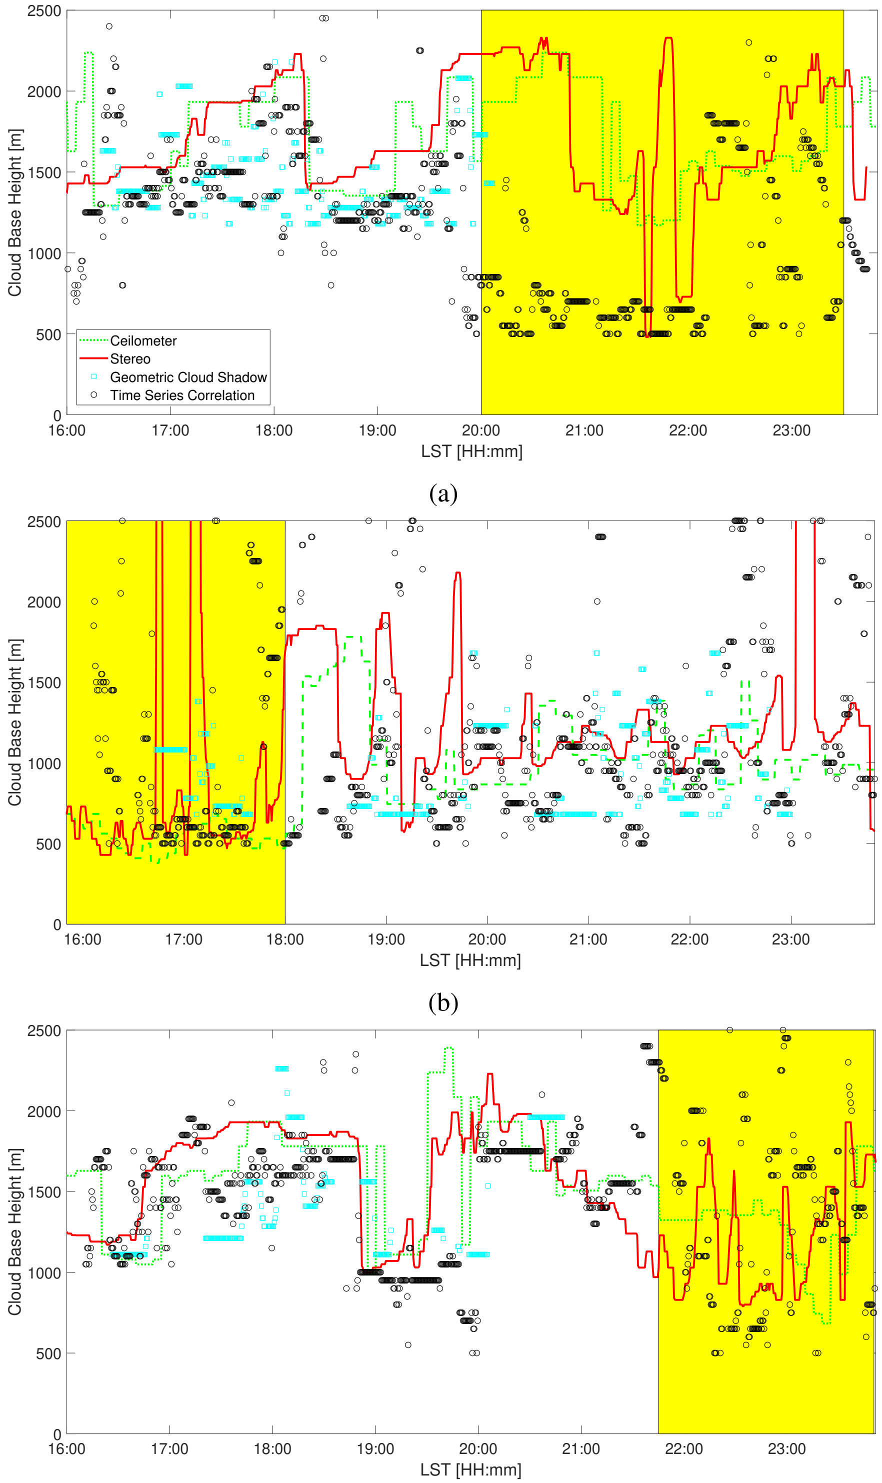

(a)

(b)

(c)

Figure 2.9: Cloud base height comparison between the TSC, GCSE, and 2D stereographic method (NK14), and ceilometer measurements for (a) Dec 14 (b) Dec 26 and (c) Dec 29, 2012. Yellow highlights show periods of rain that are ignored for the quantitative comparison.

capture) time step (i.e. 30 seconds) in this work. Beyond increasing the number of samples, treating the sensor network as an array and applying array signal processing methods may provide a lower variance CBH estimate.

Sensitivity of CBH to Correlation Coefficient

During certain periods, the variation of the correlation coefficients R_j over all CBH candidates was found to be small. For example, Figure 2.10 gives the mean correlation R_j (Eqn. 2.3) at different H_j for a selected period. The maximum R_j is very similar to the minimum with R_j ranging from 0.9 to 1. While the changes in R_j are small relative to its range, the relative changes in H_j are considerable at 1050 m to 1700 m. In this case, due to the small difference between the minimum and maximum correlation coefficient, the selected H_j may be determined by small and somewhat random fluctuations in R_j which is not desirable behavior for an accurate and robust CBH estimation algorithm. Small variations in the correlations are caused by homogeneous cloud cover (e.g. overcast condition) or a cloud projection that is insensitive to CBH changes (e.g. collocated sky imagery and pyranometer).

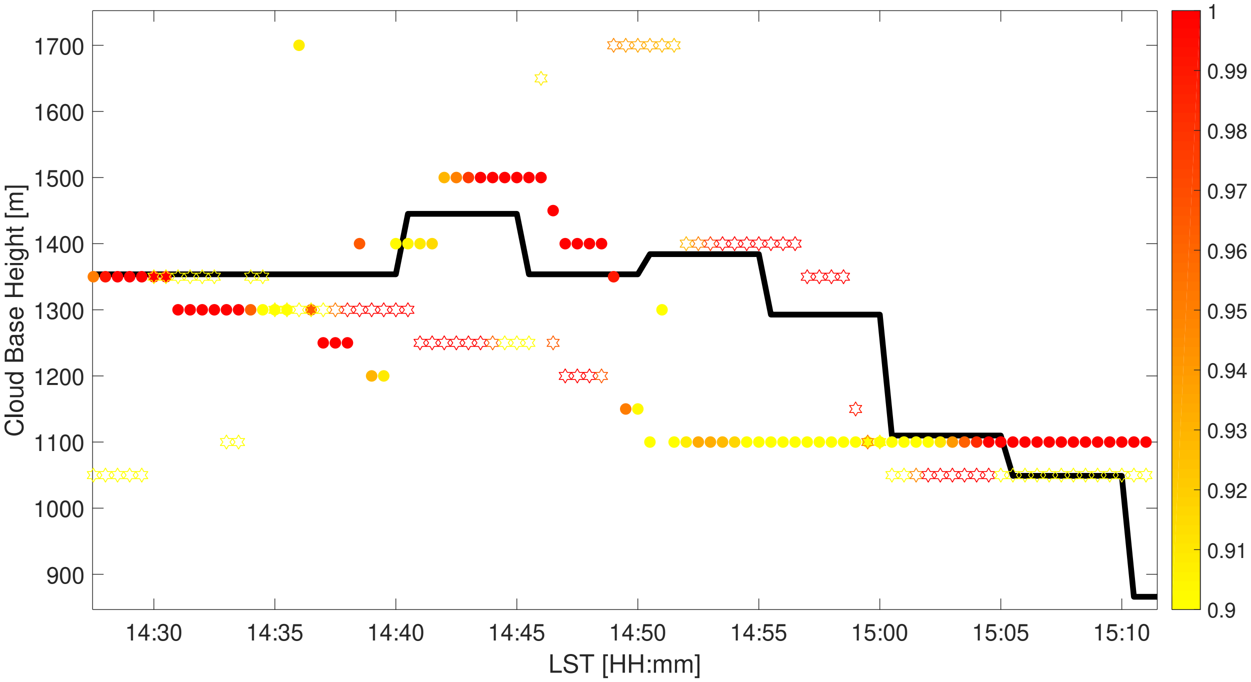

Figure 2.10: Example of CBH estimates for the TSC method versus the ceilometer for a 45 minute period on Dec 29, 2012. The color of each symbol indicates the average correlation coefficient R_j (Eqns. 2.2 and 2.3) between the observed and simulated nowcast GHI from the set of stations. For each time step the CBH (y-axis) and its associated maximum R_j (filled circles) and minimum R_j (open hexagrams) are shown.

Moreover, larger imager to station distance can promote errors in the GHI time series from the sky imager. As indicated in Eqn. (2.9), a larger pixel zenith angle (more distant station) results in cloud projection being more sensitive to CBH changes because the cloud projection error scales with $\tan\theta$. In addition, the lower pixel resolution for the outer part of the sky image at larger pixel zenith angle can cause larger random errors in shadow projection at the ground station.

2.5.2 GCSE Performance

GCSE combines ramp detection with an analytic-geometric component derived from the sky imager forecast. Down ramp events are detected for each ground station and associated cloud edges are matched in each station's H-Δt map. Since the construction of H-Δt map is a matrix indexing operation for an image, it takes less than a second to construct H-Δt map on a typical Intel I5-powered workstation, making operational use feasible.

In terms of nRMSD, GCSE performed over 10 percentage points worse than the NK14 method over the three days studied. In the more extensive 30 day comparison, GCSE improved substantially with an nRMSD of 20.8%. For all error metrics, the GCSE performed worse than TSC. This is in large part due to the modeling complexity and assumptions involved (see Section 2.5.3). An accurate down ramp start time from GHI observations is required for the GCSE to work correctly. The method described in Section 2.3.2.2 is a reasonable approach if the ramp is monotonically up or down. But in some cases ramps exhibit local extrema, causing the proposed approach to misidentify the start time.

Figure 2.11 provides an example on Dec 26 with scattered cumulus clouds. The left plot shows a large down ramp with a complex *kt* time series: two local extrema are identified in the time series maximum ramp points difference (Figure 2.5) at 12:23:21 LST (black dashed line) and 12:23:52 LST (green dashed line). The associated ramp event start times are determined at 12:23:26 LST (black dot) and 12:23:47 LST (green dot), respectively, by searching within 5 seconds for a local maximum in *kt*. While visual inspection suggests that the black dot is a reasonable ramp start time, the *kt* variation around the two original local extrema is small, making identifying the start time somewhat random. These small "pre-ramp" events are likely caused by the multiscale nature of clouds and associated deformations around the cloud boundary.

The impact of this ambiguity in the local extremum is illustrated in the right plot in Figure 2.11. The black dot corresponds to black line $(H, \Delta t)$ = (775 m, 86 s) and the green dot corresponds to green line $(H, \Delta t)$ = (625 m, 107 s) with $\Delta t = 0$ at 12:22:00 LST. The two local extrema that are spaced by 31 s cause a 150 m difference in CBH. In this case, the local extremum # 1 is slightly greater than # 2, so it is selected per the procedure in Section 2.3.2 and the associated local maximum (i.e. black dot) is used to determine the CBH candidate. While in this case the final selected CBH is closer to the ceilometer measurement of 866 m than the alternate, similar ambiguities in local extremum and subsequent CBH variation were common in the analysis.

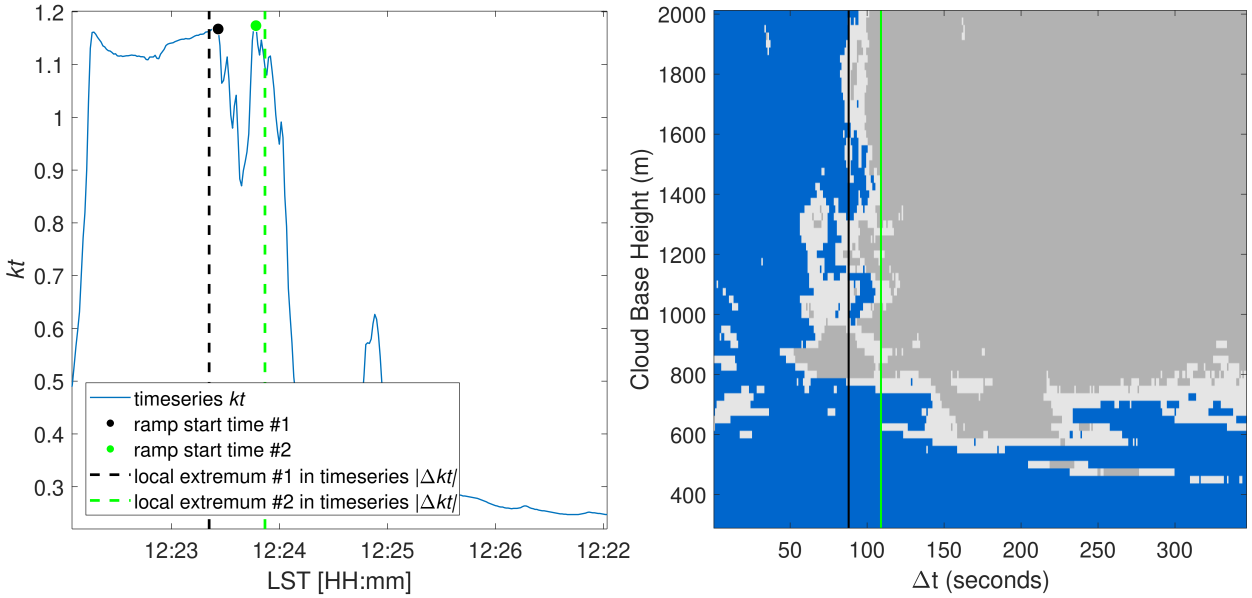

Figure 2.11: Sensitivity of CBH to ramp start time and ambiguity in ramp start time estimation. Left: Two local extrema (dashed lines) are identified due to a non-monotic time series of *kt* (the dots show how following ramp detection each ramp start time is adjusted to the local maximum in kt within 5 sec). Right: The corresponding H-Δt map with $\Delta t = 0$ sec at 12:22:00 LST. The vertical lines in the right plot correspond to the colored dots in the left plot. The actual CBH measurement from ceilometer is 866 m.

2.5.3 Other Modeling Errors Affecting CBH Estimation

Both TSC and GCSE rely on derived products generated in the USI forecast procedure that apply simplifying assumptions and inject additional uncertainty into CBH estimation. Naturally, since sky images are the key input to both methods, TSC and GCSE are not operational at night. Cloud edges derived from sky imagery rely on the cloud decision process determining where clouds "begin". The methods to detect cloud presence are generally accurate, but there is some inherent uncertainty in a binary pixel classification as being-cloudy or cloud-free [40] particularly near cloud edges which may have a diffuse and blurred transition. This affects both TSC and GCSE.

Another issue is that extensive cloud evaporation and formation can cause GCSE to fail because the "frozen" cloud advection assumption is violated. Consider the case where a cloud forms between time t_o when a sky image is taken, and the time when that clouds edge causes a down ramp at time t_1. Another stable cloud that was present in the sky image at t_o causes a down ramp at time t_2 where $t_2 > t_1$. Although the down ramp occurring at t_1 is detectable in GHI data, the cloud map generated from data at t_o only has the information of the cloud which passes at t_2. For TSC, this increases the separation between the measured and modeled GHI time series, affecting correlation coefficients across the CBH grid search. The GCSE ramp detection algorithm will return a ramp occurrence time of $\Delta t_1 = t_1 - t_o$ which does not have a matching cloud edge in the H-Δt map. The H-Δt map search process will yield the best available clear-cloudy transition at Δt_1 which is likely to be incorrect.

Besides, both methods are affected by overcast conditions with homogeneous cloud cover. The TSC method identifies concurrent cloud edge events using the correlation coefficient. If the 20 minute sample window does not contain any significant cloud-edge induced fluctuations, the correlation coefficients are small and likely no CBH will be output. The GCSE method does not provide CBH in overcast conditions either; while ramp detection may still be feasible due to variability of cloud optical depth in overcast conditions, the cloud travel time cannot be estimated from the binary H-Δt map. Fortunately, in overcast or clear conditions the solar irradiance can be predicted accurately without CBH because all stations are likely covered by the same sky condition and receive similar irradiance.

Additionally, the pixel resolution for the outer part of the sky image at larger pixel zenith angle is degraded, making the estimated cloud cover more uniform over the 20-minute comparison interval. Any station whose shadow projection comes from these perimeter image sections will lack detailed cloud structure. This less detailed cloud structure yields lower correlation for the TSC, and larger errors in identifying the timing of sky condition changes for the GCSE. Interestingly, for the GCSE, the forecast at the sky imager position is unaffected by CBH and thus the forecast GHI does not suffer from CBH errors.

The temporal resolutions for TSC and GCSE differ: the TSC output rate is one sample per 30 seconds as set by the image capture frequency, but data availability may be less frequent due to low correlation coefficients. The GCSEs output rate depends on the existence of sufficient variability in cloud cover, and the ability to find a consensus CBH candidate. For the dataset presented here, GCSE outputs 34% less CBH samples than TSC, for an average of one GCSE sample every 75 seconds. While this lower output rate is sufficient for short-term solar power forecasting, it may be a limiting factor for other applications or in other sky conditions with less cloud cover. The valid time of the CBH estimates also differs between methods: Since TSC correlates the last 20 minutes of GHI data, the estimated CBH applies to those 20 minutes. While GCSE utilizes only a single dominant down ramp in the GHI time series the CBH strictly applies to that time instant only.

The cloud velocity estimation of the sky imager is actually an apparent cloud edge velocity, which is a combination of cloud speeds due to advection along with cloud formation or evaporation occurring from image to image. These cloud dynamics introduce real or apparent fluctuations in cloud speed which negatively affects the performance of GCSE because construction of the H-Δt map assumes that the cloud velocity remains constant over the CBH estimation interval (typically 10 minutes). TSC is insensitive to cloud speed variability as it does not employ a cloud advection scheme.

Last, multiple cloud layers and cloud three-dimensionality [46] can degrade the performance because both methods operate under assumption of single-layered planar cloud cover.

2.5.4 Number of Stations and Spatial Diversity

Geographic variations at the individual sites may affect both TSC and GCSE. For TSC, averaging correlation coefficients at each CBH blurs potential station-to-station differences in correlation coefficient due to real differences in CBH. For GCSE, station-to-station ramp timing errors may cause inconsistent CBH candidates, preventing an accurate CBH estimate. The current limitation of our setup was the availability of only six stations, four of which were located within 600 m of each other resulting in more correlated GHI data and little diversity in perspectives. A logical extension to this work is to examine the impact of adding additional ground stations. At large solar installations, weather stations, reference cells, and individually metered inverters can all be used to improve spatial distribution of stations.

2.6 Conclusions and Future Work

The objective of this paper was to propose two methods for CBH estimation requiring a single sky imager together with spatially distributed irradiance or power output measurements, providing an alternative CBH estimation technique to direct, in-situ, or multi-camera approaches. These new methods can serve as a low-cost alternative to ceilometers for sky imager based short term solar power forecasting in which the cloud height information is required [35, 36].

The TSC method, is comparatively simple and the more reliable of the two proposed methods. The GCSE method, relies on a complex stack of models: cloud detection, cloud velocity estimation, cloud shadow forecasting, and down ramp detection. The construction of the H-Δt map is a novel feature of this work, and its utility is demonstrated for the purposes of cloud edge matching and CBH estimation. Overall, the GCSE method performed slightly worse (1 percentage point larger nRMSD) than the TSC method. For both methods, the nRMSD remained below 21% for all 30 days. On the other hand, the CBH estimate derived from a sky imager coupled with a cloud speed sensor in our previous work [31] yielded better accuracy (17% nRMSD) on a different set of 30 days, owing partially to the strict filtering of the raw cloud speed measurement.

Future efforts will involve improving both sky imager cloud detection and cloud velocity estimation, which will also benefit solar power forecasting with a sky imager. Chow et al. [39] proposed optical flow to enable detection of multiple cloud layers as well as their respective cloud pixel speeds, which is an improvement to the cross-correlation velocity estimation method used in this work. Adding more and more distributed ground stations will also help improve the robustness of the methods. Finally, validation under different meteorological conditions more relevant to continental climates would further substantiate the general applicability of the methods.

Table 2.4: Comparison of cloud base height estimates between TSC and GCSE methods for 30 days. Ceilometer daily averages are reported as “Mean CBH”.

Date	Mean CBH	TSC method					GCSE method				
	[m]	MBD [m]	nMBD [%]	RMSD [m]	nRMSD [%]	No. Points	MBD [m]	nMBD [%]	RMSD [m]	nRMSD [%]	No. Points
09/20/2013	613	-90.5	-14.8	123.6	20.2	616	-104.7	-17.1	129.9	21.2	157
09/21/2013	695	40.6	5.8	109.9	15.8	752	45.3	6.5	124.4	17.9	362
09/22/2013	618	137	22.2	144.5	23.4	750	129	20.9	138.9	22.5	217
10/02/2013	598	-14.5	-2.4	66.7	11.2	337	-109.9	-18.4	161.1	26.9	408
02/04/2014	1110	-41.5	-3.7	294.4	26.5	410	-161.6	-14.6	296	26.7	712
02/09/2014	747	32.7	4.4	118.9	15.9	606	-51.1	-6.8	156.8	21	473
02/10/2014	648	88.3	13.6	135	20.8	822	76.5	11.8	118.2	18.2	291
03/22/2014	1116	-241.4	-21.6	365.9	32.8	424	-389.6	-34.9	444.7	39.9	416
03/23/2014	657	-27.7	-4.2	122.3	18.6	696	-40.8	-6.2	118.6	18.1	219
03/25/2014	651	-40	-6.1	282.6	43.4	750	-33.7	-5.2	258.7	39.7	141
03/27/2014	971	43.5	4.5	268	27.6	504	94.9	9.8	315.7	32.5	646
05/20/2014	783	33.1	4.2	162.7	20.8	460	-98.3	-12.5	197.7	25.2	679
06/10/2014	633	-18.6	-3	101.7	16.6	552	-81.6	-12.9	120.8	19.1	283
04/20/2015	662	-65.9	-10	107.2	16.2	531	-85.1	-12.9	110.2	16.6	209
05/21/2015	1263	-65.6	-5.2	288.8	22.9	285	-73.8	-5.8	324	25.6	213
04/12/2016	460	58.6	12.7	89.4	19.4	471	62.1	13.5	108.9	23.7	389
04/22/2016	485	42.7	8.8	89.4	18.4	563	-0.9	-0.2	87.6	18.1	404
04/25/2016	998	-2.8	-0.3	147.4	14.8	760	-11.1	-1.1	145.7	14.6	847
05/21/2016	1004	51.9	5.2	97.1	9.7	525	-23.5	-2.3	148.4	14.8	507
05/23/2016	916	61.9	6.8	124.4	13.6	361	69.7	7.6	128.4	14	411
05/26/2016	771	99.2	12.9	124.3	16.1	309	31.3	4.1	212.4	27.5	342
07/17/2016	409	73.5	18	97.1	23.7	672	70.9	17.3	93.3	22.8	266
08/11/2016	555	17.4	3.1	93.4	16.8	374	-15.7	-2.8	82.9	14.9	610
08/26/2016	725	-49	-6.8	82.5	11.4	773	-60	-8.3	85.3	11.8	323
08/28/2016	376	52.8	14	64.5	17.2	744	38.4	10.2	56	14.9	520
09/02/2016	514	-46.6	-9.1	62.9	12.2	614	-59.3	-11.5	68.4	13.3	280
09/08/2016	603	-42.5	-7	80.6	13.4	605	-53.5	-8.9	83.4	13.8	253
09/09/2016	469	18.1	3.9	41.2	8.8	654	19.5	4.2	40	8.5	278
09/12/2016	714	27.3	3.8	102.7	14.4	699	26.3	3.7	93.2	13.1	237
10/12/2016	689	-118.9	-17.3	156.4	22.7	805	-134.4	-19.5	171.5	24.9	377

2.7 Appendix

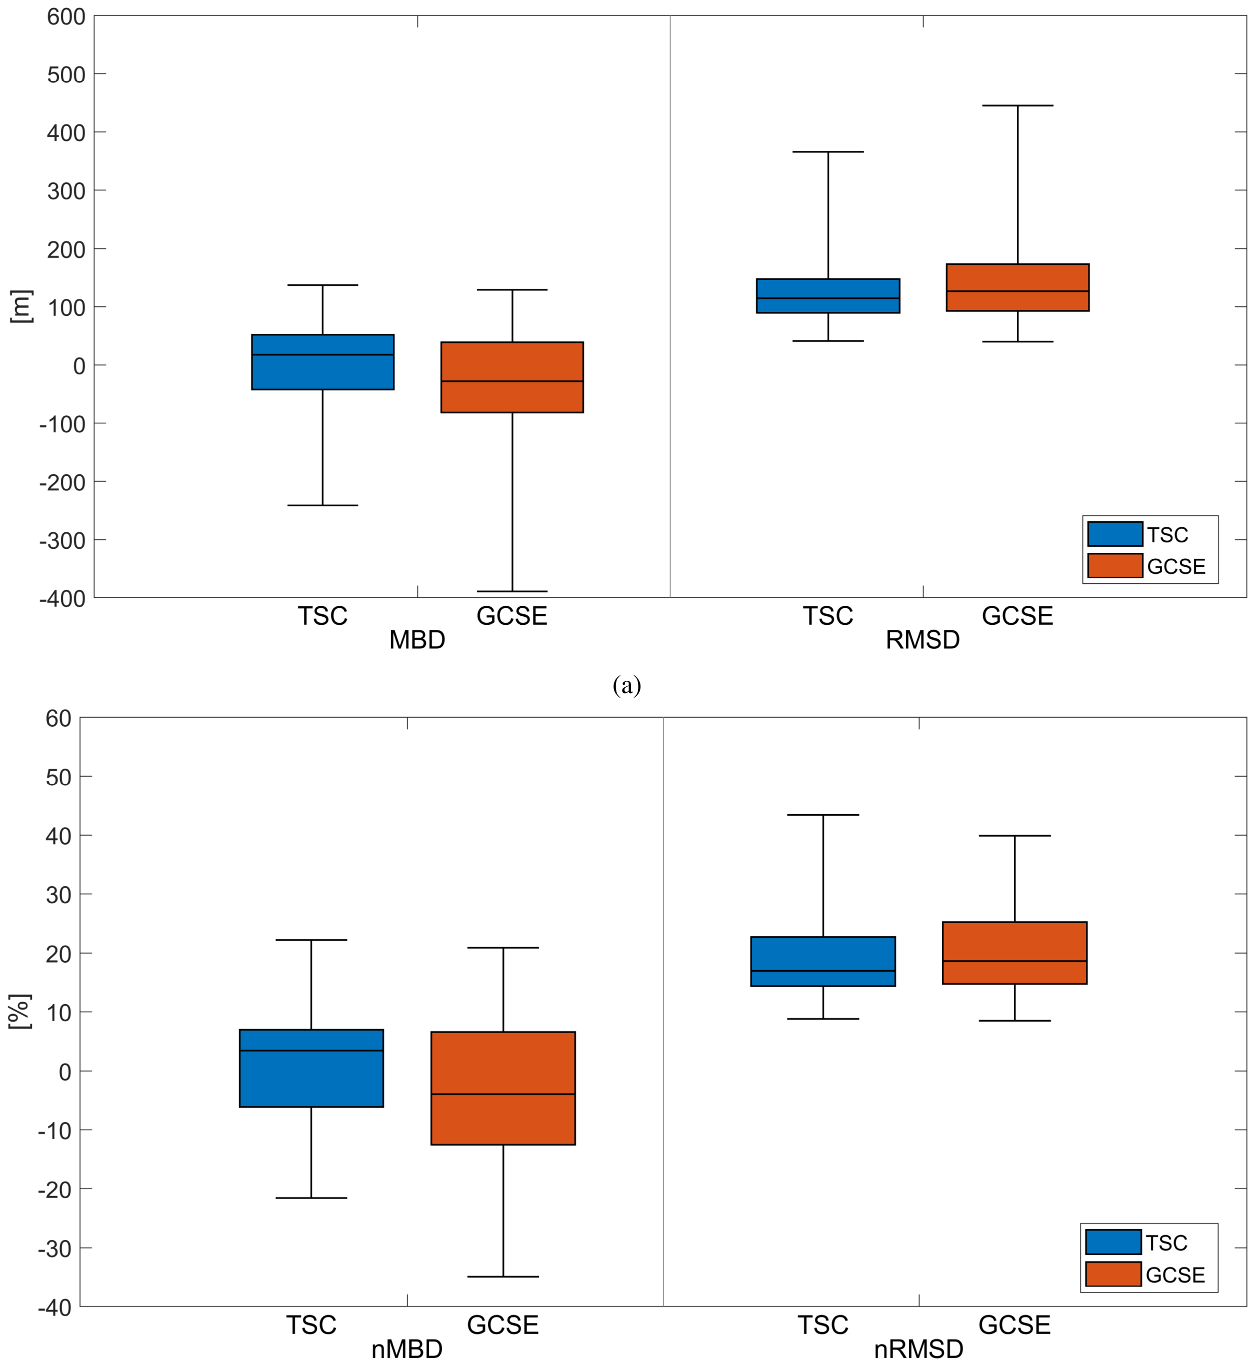

(a)

(b)

Figure 2.12: Box plot of CBH errors (a) and normalized CBH errors (b) for TSC and GCSE methods over 30 days. This figure is a visual representation of the performance metrics in Table 2.4.

2.8 Acknowledgments

This Chapter, in full, is a reprint of the material as it appears in “Wang, G. C., Urquhart, B. & Kleissl, J. (2019). Cloud base height estimates from sky imagery and a network of pyranometers. *Solar Energy*, 184, 594-609”. The dissertation author was the primary investigator and author of this paper.

Chapter 3

Cloud Base Height from Sky Imager and Cloud Speed Sensor

3.1 Introduction

Cloud base height (CBH) plays an important role in many solar energy applications. For example, [47] incorporate CBH to generating synthetic irradiance signals. While an accurate source of CBH become less critical in larger-scale forecasting such as satellite image based forecasting, it does matter in short-term solar forecasting which is becoming vital in the solar industry as solar penetration increases ([2]). As the cloud is observed by the sky imager, variations in CBH change the distance between the latitude and longitude of the center of the cloud and its shadow on the ground. In addition the physical cloud size (and its shadow size) scales linearly with CBH. Hence, incorrect CBHs lead to offsets between modeled and actual cloud shadow. Finally, inaccurate cloud speed associated with CBH errors causes errors in the estimated arrival time of cloud shadows, which leads to offsets in ramp forecasting.

The most common CBH measurement techniques include radiosondes [18] and ceilometers [19, 20]. A radiosonde is a battery-powered instrument package that vertically profiles the atmosphere. Although the measurement is accurate as it is taken in-situ, the observations are usually taken only twice daily at major airports. This frequency is not sufficient for intra-hour forecasting. Ceilometers, as the most common CBH observational tool, emit a high intensity near-infrared laser beam vertically. A vertical profile of atmospheric backscatter is then obtained and CBH can be computed multiple times per minute. Ceilometer

CBH measurements are usually reported in meteorological aerodrome reports (METAR). While METAR stations report high quality CBH data, limited temporal resolution (hourly reports) and spatial heterogeneity in cloud cover and CBH, can cause differences between METAR and local CBH. Since the cost of ceilometers is relatively high, their application outside of airports is limited in most countries, although ceilometers are standard at weather observation stations in the UK.

A few indirect methods of CBH estimation have emerged during the past decade. [26] estimate CBH based on the output of a Numerical Weather Prediction model. CBH estimates with ground based infrared measurements [22, 23] were developed based on the monotonic relationship between CBH and downwelling thermal infrared radiance. The assumption that clouds are blackbodies leads to an over-estimation of the CBHs derived by infrared cloud imagers [23]. Satellite images [24, 25]; estimate cloud height with great spatial coverage and resolution, but the fact that satellite radiance is primarily a function of cloud top height limits its application in short-term solar forecasting. The stereographic method using two or more sky imagers was initially proposed by [27] and refined by [28]. Nguyen and Kleissl [29] further improved stereographic CBH detection and determined CBH using a 2D method for single homogeneous cloud layers and an enhanced 3D method to provide CBH with high spatial resolution. However, the stereographic method requires two sky imagers spaced 1.23 km apart and accurate geometric calibration of the imaging systems is critical [38].

The cloud shadow speed sensor (CSS) [34] or cloud speed measurements from spatially distributed irradiance or power sensors within a power plant [43] offer an alternative to direct CBH measurements when combined with a sky imager. Since the cloud pixel speed (or angular cloud speed) determined by the sky imager can be expressed as the ratio of cloud speed [m s^{-1}] and CBH, CBH can be computed from collocated sky images and cloud motion vectors (CMVs). Hence, accurate CMV estimation is critical to CBH computation. While existing CMV methodology was proposed by [44], we present an enhanced CMV methodology that is more suitable for CBH computation. Some limitations of the approach and validation should be disclosed upfront. The CMV as derived from the CSS applies to the cloud edge approaching the sun, but cloud pixel speed is determined in the entire field of view of the sky imager, resulting in inconsistency in CBH computation. Furthermore, the ceilometer measurement used for validation presents temporally averaged CBH at zenith (versus at solar zenith for the CSS). Therefore, random differences between computed CBH and ceilometer CBH are expected for validation, but biases should be small.

The principal objective of this paper is to propose a method that offers an accurate local CBH for sky imager solar forecasting. This method incorporates a cloud speed sensor with an enhanced algorithm to a sky imager, and the package provides an affordable and convenient approach to estimate CBH compared to a ceilometer. This paper is organized in five sections. The UCSD CSS and data availability will be described in Section 3.2. A new algorithm to derive cloud speed from CSS raw data is described in Sections 3.3.1 and 3.3.2. Sections 3.3.3 and 3.3.4 introduce a method that combines CSS and UCSD sky imager (USI) results to determine CBH. Section 3.4 provides CBH validation against an on-site ceilometer. Section 3.5 provides conclusions on the method, applications, and future work.

3.2 Hardware

3.2.1 Instrumentation and Setup

The CSS [34] is a compact system that measures cloud shadow motion vectors (CMVs). The system offers an affordable technique to measure CMVs with material costs of less than US$ 400. It consists of an array of eight satellite phototransistors (TEPT4400, Vishay Intertechnology Inc., USA) positioned around an identical phototransistor located at the center of a half circle of radius 0.297 m, covering 0-105° in 15° increments (Fig. 3.1). The sensors have a spectral response ranging from approximately 350 to 1000 nm with peak sensitivity at 570 nm. Sensor response time was determined experimentally in a laboratory controlled environment and was found to be 21 μ s rise time (10 - 90% response). High-frequency irradiance data are taken from all sensors and fed to a microcontroller (chipKIT Max32, Digilent Inc., USA). The on-board static memory allows fast storage of 6,000 10-bit data points per cycle. Due to the high sampling frequency, the measurements are not continuous. With the sampling rate of 667 samples s^{-1}, 6,000 data points fill up the on-board memory in approximately 9 sec. These 9 sec of data are then processed to determine one CMV as described in Section 3.3. During this process, the raw data collection has to be temporarily suspended for about 9 sec resulting in a temporal resolution of CMVs of about 18 sec. The CMVs used in this analysis were taken from a CSS located at 32.8810° N, -117.2328° W, and 106 m height above mean sea level (AMSL) (marked as CSS in Fig. 3.2).

Sky images were taken every 30 sec by a USI located at 32.8722° N, -117.2409° W, 129 m AMSL (marked as USI1-2 in Fig. 3.2). The USI is designed and developed for short-term solar forecasting applications [38]. It features a high quality imaging sensor and lens contained in a thermally controlled and compact environmental housing. The capture software is employed with a high dynamic range (HDR)

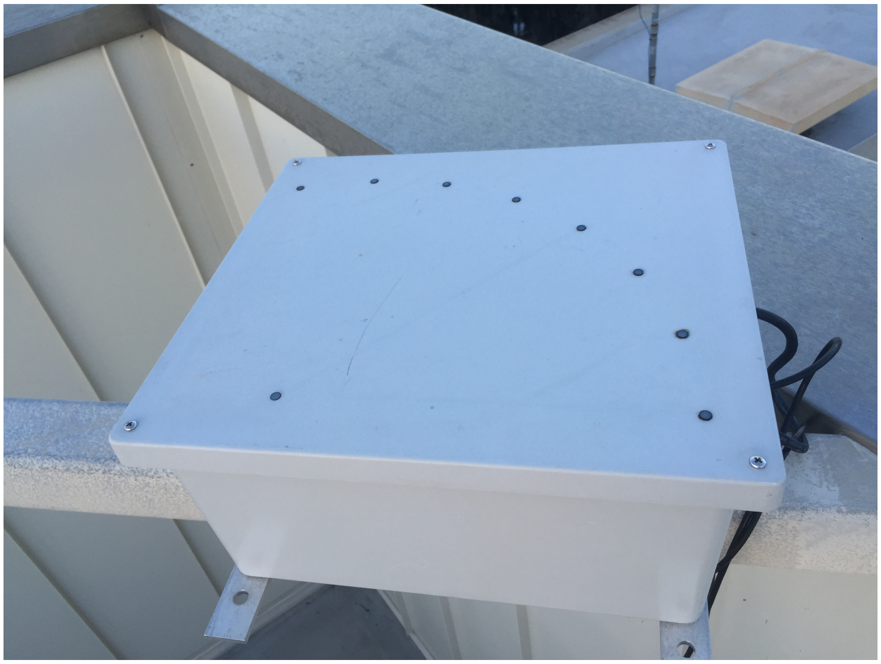

Figure 3.1: Cloud Shadow Speed Sensor (CSS) contained inside a weather-proof enclosure with dimensions 0.45 x 0.40 m. On the top of the enclosure is an array of nine phototransistors.

imaging technique. Independent measurements of CBH were taken by a Vaisala CT25K ceilometer co-located with the CSS. While all sensors report CBH above ground level (AGL), the elevation of the sensor was added to obtain CBH (AMSL).

3.2.2 Data Availability

Since USI data was available continuously, data availability was restricted by the CSS and ceilometers operational availability. The CSS was setup on Apr 4 2015. However; intermittent technical issues occurred until May 1, 2015, when it became fully operational. In order to comprehensively assess the performance of the CSS during a variety of sky conditions, April 5, April 20, and the period of May 1 through July 29 were selected for analysis. During this period, 35 of 92 days were clear or contained less than 4 hours of cloud cover per day, and there were 21 overcast or rainy days. Because clear and overcast days do not produce nearly as many ramp occurrences as partly cloudy days, our study rejects the days with clear or overcast conditions. Nine additional days had to be eliminated due to missing ceilometer measurements. The remaining 27 days contain partial cloud cover for at least 4 hours (except July 1 which contains unusually high clouds for the southern California region which lasted for 2 hours), which are the conditions of interest for testing CSS performance.

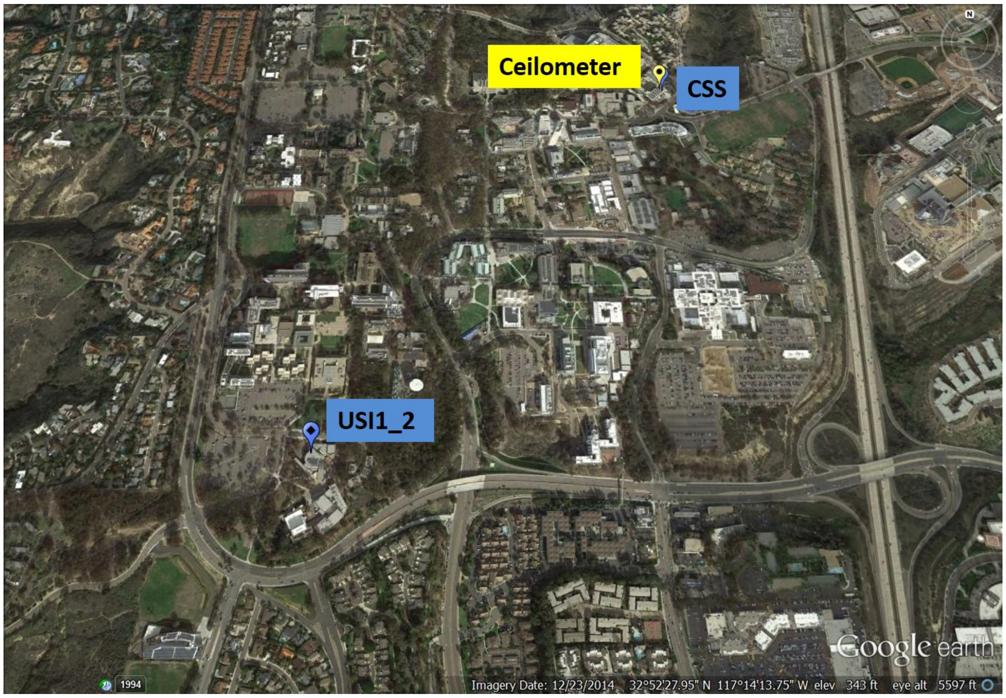

Figure 3.2: Locations of sky imager (USI1-2), ceilometer and Cloud Shadow Speed Sensor (CSS) on the UCSD campus. The straight-line distance between USI and CSS is 1.25 km. Map data ©2015 Google.

3.3 Cloud Speed Measurements

3.3.1 Prior Cloud Speed Sensor Algorithm: Most Correlated Pair Method

While the method of determining CBH is compatible with any measurement of cloud speed, we also present a new method to obtain cloud speed from the CSS as it had not been documented before. In the prior CSS algorithm proposed by [44], the CMVs were determined by the Most Correlated Pair Method (MCP). MCP assumes that due to heterogeneity in the cloud shadow over the area of the sensor, the pair of sensors that lie along the direction of cloud motion will experience the largest cross-correlation as they see the same transect of the cloud [44]. Thus, the pair with the largest cross-correlation coefficient is therefore used to determine the direction of cloud motion. The time shift of maximum cross-correlation between the selected pair is then used to calculate the cloud speed. The MCP method suffers from some deficiencies. Most importantly, for the ideal case of a linear cloud edge separating shadow from clear sky, each sensor would see exactly the same signal shape and there would be no single most correlated pair. Instead, the most correlated pair would simply result from arbitrary correlations from sensor noise. Scenarios close to this idealization were found to be common. Since clouds are typically much larger than the spacing between sensors, it seems intuitive that the cloud is nearly homogeneous over the area of the CSS. Thus, CMV results were highly variable. Bosch et al. [44] addressed the variability through statistical post-processing to determine the most common cloud direction and corresponding cloud speed. The post-processing was shown to be robust and accurate, but the temporal averaging reduced the response of the sensor to sudden changes in cloud velocity. The MCP method also had limited precision as the final direction could only be along individual sensor pairs.

3.3.2 Improved Cloud Speed Censor Algorithm: Linear Cloud Edge Curve Fitting Method (LCE-CFM)

The assumption in the MCP method is modified to enhance the accuracy and robustness of the method in an operational environment. Because the CSS is small compared to a typical cloud, we can reasonably assume the cloud edge to be linear (Fig. 3.3). The signal measured by each sensor is then identical except of the temporal deviation between the signals, resulting in a perfect cross-correlation $R_{ij} = 1$ (i and j refer to the sensors). Therefore, it is not the magnitude of the cross-correlation that distinguishes the sensor pair aligned with the CMV, rather the time lag associated with the maximum R_{ij} between different sensor pairs provides clues as to the relative alignment of each pair with respect to the CMV. Hence, we will fit a function to the time lag versus sensor-pair direction, and we term this method the "Linear Cloud Edge Curve Fitting Method".

As in the MCP method, the maximum cross-correlation coefficient R_{ij} of each pair of signals will be determined and the associated time shift Δt_{ij} for that pair will be recorded. Considering a linear cloud edge that is crossing the CSS moving in the direction of the sensor line (a-c), it is straightforward that:

$$r \cdot \cos\phi_{ij} = \mathrm{U}_{\mathrm{CSS}} \cdot \Delta t_{ij}, \tag{3.1}$$

where r is the radius of the sensor circle, ϕ_{ij} is the cloud edge direction that is defined as the angle between the line connecting sensors i and j and the line (a-c). i and j vary from 0 to 8, but only 12 sensor pair combinations i / j (0/1, 0/2, 0/3, 0/4, 0/5, 0/6, 0/7, 0/8, 1/5, 2/6, 3/7, 4/8) are used in our configuration. ϕ_{ij} can be expressed as $(360° - p_k \times 15°)$ where p_k is the sensor pair number $k = 0$ to 11 following the brackets in the previous sentence). U_{CSS} is the speed of the cloud edge, i.e. cloud speed. With distance r and cloud speed $\mathrm{U}_{\mathrm{CSS}}$ being constant for each pair, the time shift Δt_{ij} becomes a function of $\cos\phi_{ij}$. The trigonometricrelation holds for all cloud edge directions as the cloud velocity is assumed to be perpendicular to the cloud edge.

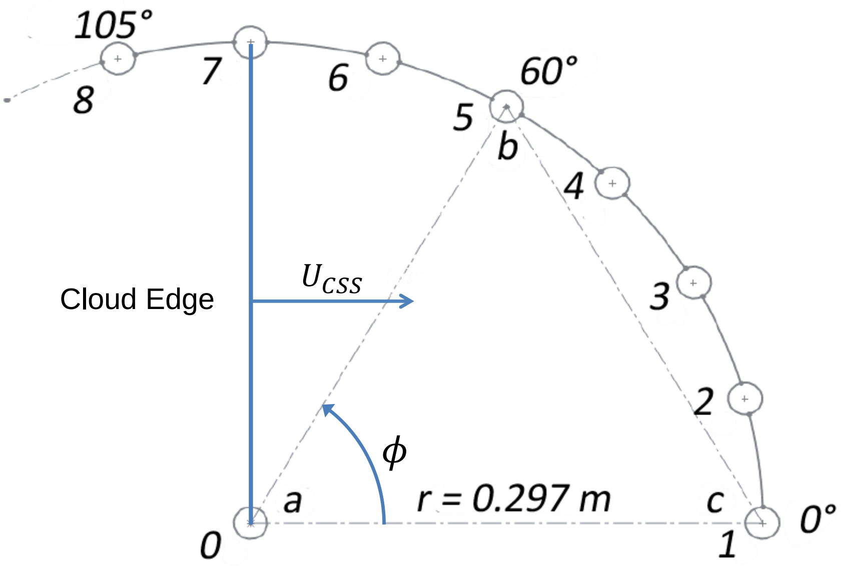

Figure 3.3: Illustration of the linear cloud edge assumption and LCE-CFM method on top of the CSS luminance sensor arrangement. Each circle represents a sensor arranged in a circular pattern with 15° spacing about the central sensor. The sensor pair combinations are constructed with the central sensor and one of the other sensors for angles from 0° to 105°, e.g., sensor pair combination 0/1 for 0°, 0/2 for 15°, etc. Additional angles from 120° to 165° are obtained through equilateral triangles between the central sensor and another sensor pair, e.g., sensor pair combination 1/5 for 120°, 2/6 for 135°, etc. The linear cloud edge is shown as a blue line and is assumed (for simplicity, but not limiting generality) to be advected along the line connecting sensors 0 and 1.

For the sensor pairs without the central sensor, Eqn. (3.1) still holds as long as the selected sensor i and the sensor j lie on one side of an equilateral triangle constructed from the central sensor, sensor i and sensor j. Because the time shift Δt_{ij} returned by the CSS can be either positive or negative depending on the cloud direction, 12 sensor pairs are sufficient to cover 360° in 15° increments.

For the ideal assumption of a linear cloud edge, plotting Δt versus would therefore be expected to produce a cosine function. For verification, the cosine function is used to fit the Δt_{ij} versus ϕ_{ij} points for each 9 sec measurement, and the R^2 value is employed to determine the goodness of the fit (Fig. 3.4). A high R^2 supports the linear cloud edge assumption. Since the linear cloud edge assumes that the velocity is perpendicular to the cloud edge, the sensor pair aligned with the CMV is farthest apart along the CMV at distance r. Thus, the maximum of the cosine function which represents the longest time shift Δt should occur at the CMV direction. While the side effect of LCE assumption is not explicitly visible in the results, the potential limitation and future improvement of LCE assumption are discussed in Section 3.4.3. The cloud speed then becomes the ratio of the distance r and the time shift Δt:

$$\mathrm{U_{CSS}} = \frac{r}{\Delta t}. \tag{3.2}$$

Note that the cosine model fit to Eqn. (3.1) should be constrained to return solutions with $\Delta t > 0$. Figure 3.4 illustrates this procedure using 9 sec of luminance data. The correlations between all sensor pairs are very large (>0.999), which causes issues in the robustness of the MCP method. On the other hand, the linear cloud edge assumption is validated through a high R^2 value (0.99) which indicates that the time shift is indeed a strong function of the cosine of the direction. As a result the CMV direction and speed can be obtained with confidence using Eqn. (3.2). In the example in Fig. 3.4 the time shift is determined as $\Delta t = 0.104$ s, and the corresponding direction is = 323° yielding a cloud speed $\mathrm{U_{CSS}} = 2.87$ m s^{-1} as per Eqn. (3.2).

A filter is applied for data quality control: If the average R_{ij} is less than 0.9 or R^2 of the cosine curve fit is less than 0.9, the CMV will not be computed. Small R_{ij} is likely a result of no cloud passage or dynamic clouds. A small R^2 indicates poor curve fitting and therefore an unreliable result. Generally partly cloudy conditions result in numerous valid CMVs while homogeneous cloud conditions (e.g., clear and overcast) result in infrequent valid CMV output due to small R_{ij}. Typically, 1700 raw data sets are recorded during an eight hour analysis day, and about 110 CMVs are delivered for an overcast day and less than 10 CMVs for a clear day. For partly cloudy days, about 400 CMVs pass the quality control, which is equivalent to one CMV value every 50 sec. The sampling rate is sufficient for cloud motion estimation.

Figure 3.5 shows a set of CMVs for one day together with filtered CMV direction determined by the USI as an independent validation. Clouds are moving northward at 1 to 6 m s^{-1} changing to eastward as the day progresses. The USI direction generally falls in the center of the CSS raw data points indicating good agreement. There is some variability in CSS raw data, which is likely a result of both physical cloud dynamics and sensor noise. The same trends are seen in the wind-rose plot for CSS data on this day in Fig. 3.6; most of CMVs cluster in the north-east-ward direction with an average speed range of 2 to 6 m s^{-1}. Additional validation of the LCE assumption is presented in Section 3.6.1.

In summary, compared to the prior MCP method, the LCE-CFM yields two distinct advantages: (i) more clustered, i.e. robust, CMV results without post-filtering, and (ii) continuous cloud direction output compared to the 15° (equivalent to the angular arrangement of the sensors) discretized output for the MCP method. To demonstrate the improvement of the LCE-CFM, an example of the prior MCP method is provided

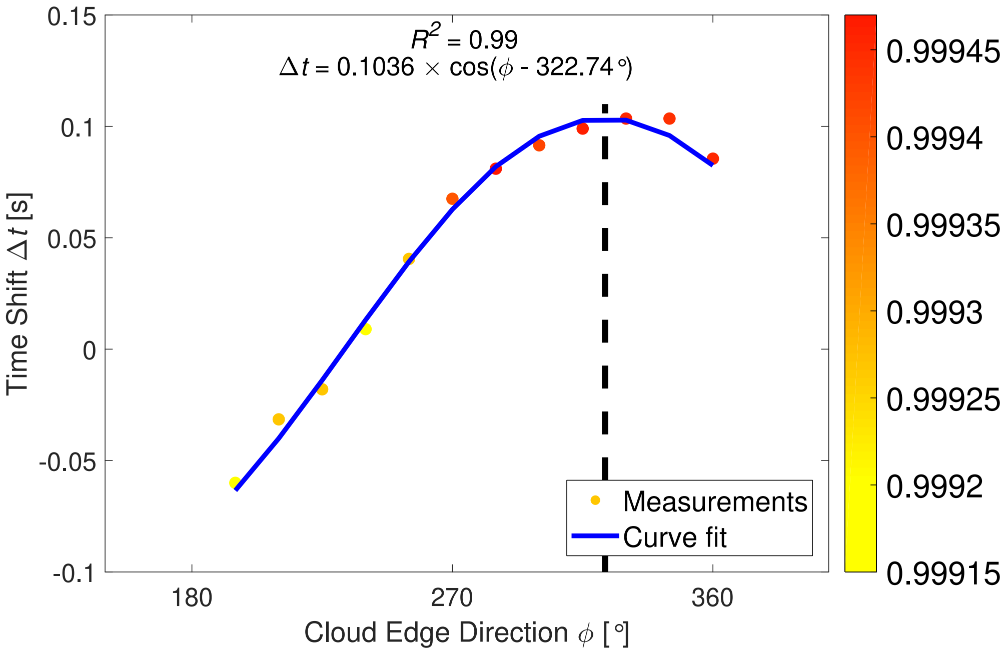

Figure 3.4: Illustration of the LCE-CFM to determine CMVs on May 31, 2015 at 17:16:19 UTC. The x-axis represents direction ϕ that is equal to $(360° - p_k \times 15)$, where p_k is the sensor pair number ($k = 0$ to 11). The y-axis represents the time shift Δt, and the color indicates the strength of correlation R_{ij}. The curve indicates the best fit of $\Delta t = 0.103 \times \cos(\phi - 322.7°)$. The maximum time shift of the cosine function is selected as the direction of cloud motion as indicated by the vertical dashed black line.

in Section 3.6.2. The disadvantage is that the LCE-CFM calculates correlation for all sensor pairs, whereas the MCP method can bypass the calculation for poorly correlated pairs. This triples the computational time on the CSS microcontroller to 40 sec. Therefore, for this application, the processing was performed on a remote Intel I5 workstation instead, which decreases computational time by more than an order of magnitude.

3.3.3 Cloud Pixel Speed from USI Data

In this section, we will first introduce the sky imager cloud motion algorithm, and based on that in conjunction with the CSS cloud speed, a local CBH will be determined. The USI can be used to detect clouds and obtain cloud pixel speed. These measurements yield forecasts of future cloud locations at high spatial and temporal resolutions and can improve forecast skill up to a 20 min forecast horizon. The benefit of using sky imager observations over a large ground sensor network is that only one or a few instruments deployed around the area of interest are capable of determining the current distribution of cloud cover at a high resolution. The forecast procedure is outlined in the flow chart in Fig. 3.7. The USI forecast procedure

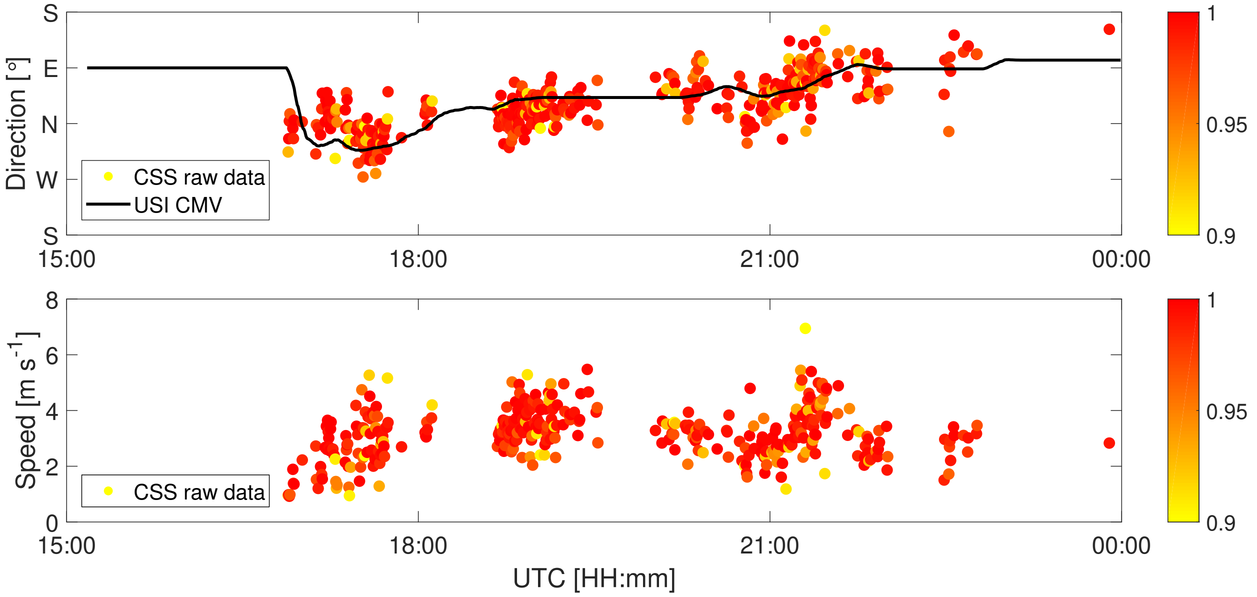

Figure 3.5: Cloud direction (defined as direction the cloud is moving towards) and cloud speed determined by the LCE-CFM using CSS data on May 31, 2015. Each circle represents one CMV computed based on 9 seconds of measurements and the color provides the R^2 value for the curve fit of 9-second measurements. The black line presents an independent validation of cloud direction using the CMV determined from the USI. Since cloud speed in m s^{-1} cannot be determined from the USI alone, there is no validation data is in the lower graph.

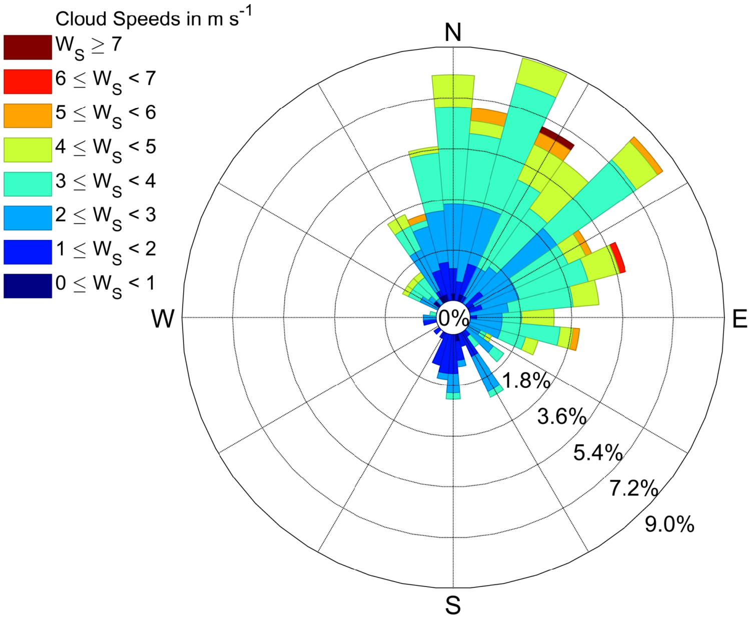

Figure 3.6: Wind-rose plot of cloud direction and cloud speed of the data shown in Fig. 3.5. The color bins show cloud speed range, and the values on concentric circles represent the frequency of appearance of each cloud speed bin.

is briefly explained within this section. It is very similar to other standard forecast procedures, such as those presented by [41] and [48]. For more details on the USI forecast, consult the references [35], [40], and [2].

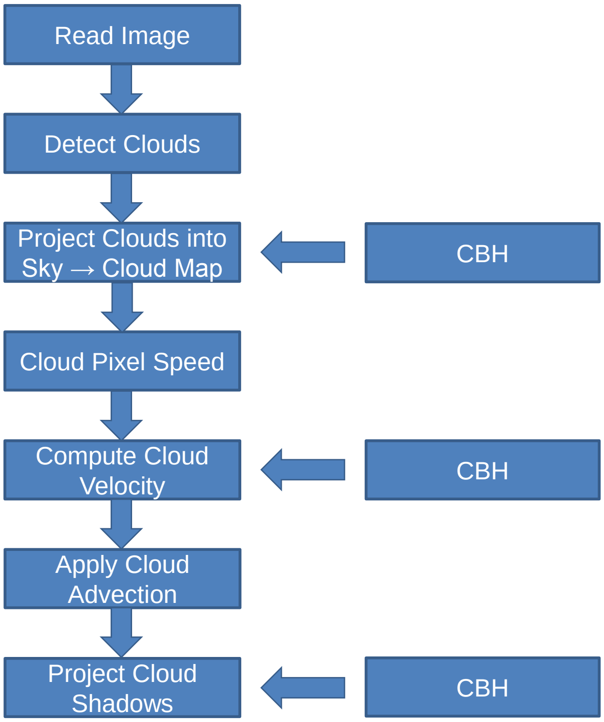

Figure 3.7: Flowchart of the sky imager solar forecast procedure. CBH is used to project clouds onto a Cartesian sky coordinate system, to obtain cloud speed, and to project the advected cloud shadows to the ground.

Cloudy pixels are detected using spectral information from the RGB images. CBH is then used in conjunction with lens geometry to map these clouds to a latitude-longitude grid at the CBH creating the cloud map. In absence of local data, CBH is taken from the closest METAR. Cloud pixel velocity is obtained by applying the cross-correlation method (CCM, [35]) to the RBR of two consecutive cloud maps. The cloud velocity [m s^{-1}] is then calculated by converting from cloud pixel speed [pixel s^{-1}] to cloud shadow speed using a velocity scaling factor which is a function of CBH (see Eqn. 3.3 later). Note that since the distance from sun to earth is much larger than the distance from cloud to earth, the cloud shadow speed is assumed to equal the cloud speed for all solar zenith angles.

3.3.4 Cloud Base Height Determination from CSS and USI

In this section, we introduce the mathematical algorithm (CSS+USI) that obtains the CBH for sky imager forecasting from CSS cloud speed measurements. Figure 3.8 introduces the geometrical terms on a cloud map. In the USI forecast, cloud velocity is calculated by converting from cloud pixel speed to equivalent m s^{-1} cloud speed as:

$$\mathrm{U_{USI}} = \mathrm{U_{pixel}} \times \frac{\mathrm{CBH} \times 2\tan\theta_m}{n_\mathrm{p}}, \tag{3.3}$$

where $\mathrm{U_{USI}}$ is cloud speed in units of m s^{-1} and $\mathrm{U_{pixel}}$ is image-average cloud pixel speed in units of pixel s^{-1} obtained through the cross-correlation method applied to two consecutive USI images. The last term in Eqn. (3.3) represents a velocity scaling factor, in which θ_m is the maximum view angle of the USI measured from zenith (here θ_m = 80°), $\mathrm{CBH} \times 2\tan\theta_m$ is the horizontal length of the sky imager view domain (termed “cloud map”), and n is the number of pixels of the cloud map in one dimension (Fig. 3.8). Therefore, the velocity scaling factor has units of m pixel^{-1}. Note that the pixel size of the cloud map is distinct from the pixel size in the original sky image.

In Fig. 3.8, the cloud observed by the USI moves from time $t = t_0$ to $t = t_1$ and $\mathrm{U_{pixel}}$ is computed from the number of pixels that the cloud moves during the period $t_1 - t_0$. The cloud map consists of $n_\mathrm{p} \times n_\mathrm{p}$ pixels, i.e. n_p is the number of pixels of the cloud map in one dimension. Its physical size is computed with the trigonometric expression $\mathrm{CBH} \times 2\tan\theta_m$. So the term $\mathrm{CBH} \times 2\tan\theta_m/n_\mathrm{p}$ refers to the physical distance per pixel of the cloud map. With the cloud speed expressed as the number of pixels per second, $\mathrm{U_{USI}}$ can be calculated according to Eqn. (3.3).

Eqn. (3.3) indicates how to obtain cloud speed in [m s^{-1}] from CBH and the USI derived cloud pixel speed. Conversely, with independent measurements of cloud speed from the CSS, $\mathrm{U_{CSS}}$, we can back-calculate the local CBH (labeled as $\mathrm{CBH_{CSS+USI}}$) by replacing $\mathrm{U_{USI}}$ with $\mathrm{U_{CSS}}$ in Eqn. (3.3) to yield:

$$\mathrm{CBH_{CSS+USI}} = \frac{\mathrm{U_{CSS}}}{\mathrm{U_{pixel}}} \times \frac{n}{2\tan\theta_m}. \tag{3.4}$$

It can be observed that CBH depends on the ratio of $\mathrm{U_{CSS}}$ and $\mathrm{U_{USI}}$. Eqn. (3.4) is implemented into the USI forecast algorithm to calculate local CBH at each step using the most recent CSS measurement. The method is called CSS+USI in the rest of the paper.

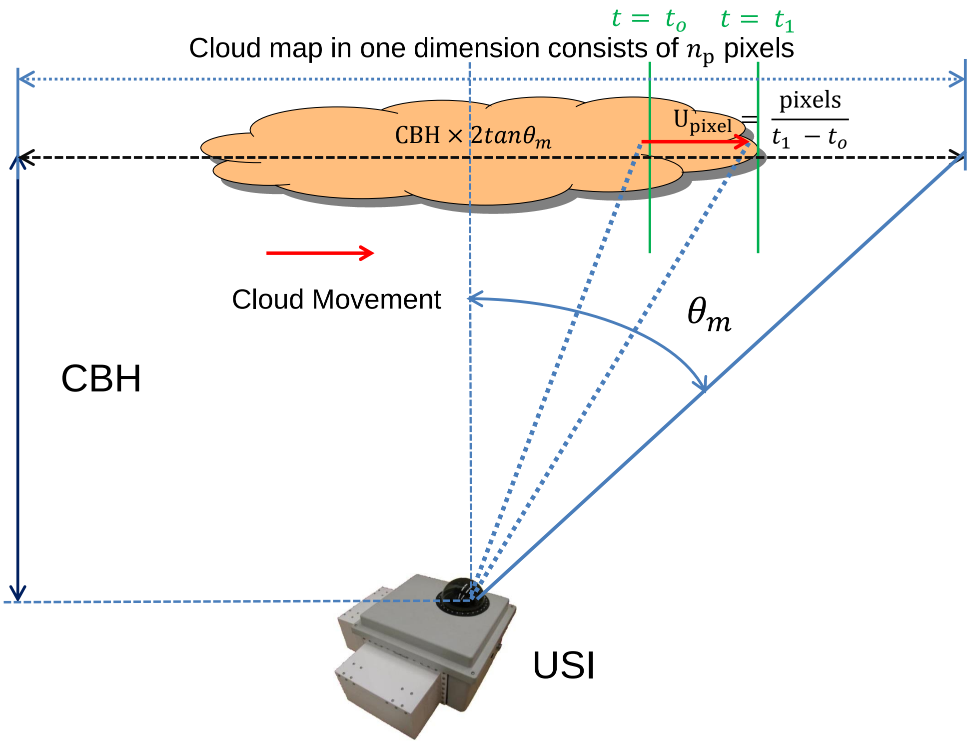

Figure 3.8: Illustration of the geometrical and kinematic relations between cloud pixel speed U_{pixel}, cloud speed determined by USI U_{USI}, maximum view angle of the USI θ_m and CBH.

A 10 min window median filter was applied to the time series of CBH from the CSS+USI method. Due to the small sampling area (a small cone above the ceilometer), heterogeneous cloud shapes, and cloud formation and movement, the raw 20 sec ceilometer data is too variable and is not representative of the CBH in the field of view of the USI. Therefore, consistent with [29] when the CSS+USI method yields a $CBH_{CSS+USI}$ at the USI timestamp, a 15 minute median filter centered on that timestamp is applied to ceilometer measurement. In this way, only the dominant ceilometer cloud layer is captured to compare with the filtered results of the proposed CSS+USI method.

3.4 Cloud Base Height Validation

3.4.1 Aggregate CBH Statistics

The CBH validation is presented in this section. The CSS+USI method is validated against METAR and an on-site ceilometer on the available days listed in Table 3.1. Two error metrics were used to characterize the performance of the method: root mean square difference (RMSD) and normalized RMSD.

$$\mathrm{RMSD} = \sqrt{\frac{1}{N}\sum_{k=1}^{N}(\mathrm{CBH}_{\mathrm{CSS+USI}} - \mathrm{CBH}_{\mathrm{ceilo}})^2}, \tag{3.5}$$

where N is the total number of data points. RMSD is divided by the daytime average CBH to obtain the normalized RMSD (nRMSD). Note that although both RMSD and nRMSD are used to evaluate the method, RMSD is relevant for the correct prediction of the timing of a ramp event.

The performance of the proposed method is summarized in Table 3.1 for a range of cloud types, cover fractions, heights, and layers that existed on these days. Generally low cumulus and low stratus clouds prevailed, but high cirrus clouds were observed on July 1, and May 22 featured altocumulus clouds. The best performance occurred on July 24 with the RMSD as low as 21 m (6.2% nRMSD), with the daily RMSD remaining below 130 m. The daily biases are usually less than 80 m and the overall bias is only 23 m indicating that most of the RMSD is driven by shorter-term random fluctuations that are difficult to model. Also, an unusual day with high cirrus for only two hours was observed on July 1, 2015, so we were able to demonstrate the performance of the method in different conditions. Thin clouds tend to have more diffused edges which may weaken the linear cloud edge assumption and the ability to obtain high correlations between different sensors. Nevertheless, the method still captures the CBH with a RMSD of 830 m that corresponds to an nRMSD of 14.2% given the large CBH. On the other hand, METAR delivers CBH with large differences to local CBH and ceilometer, which demonstrates the spatial variability in cloud coverage due to the climate difference as the METAR site is located 8.8 km further inland, while the CSS is only 1 km from the coastline (These spatial differences would likely be smaller at flat continental sites). In fact, the CSS+USI CBH delivers better CBH than METAR on all days in this study. The proposed CSS+USI method is therefore expected to be superior to METAR CBH in short term solar forecasting.

Note that the sky imager cloud pixel velocity represents all cloud edges in the entire sky image, while the CSS measurement represents a single cloud edge approaching the sun. However, we assume that those two measurements refer to the same cloud edge when applying Eqn. (3.4) and the effect of the assumption limits the CBH accuracy. In addition, the ceilometer measurement in our validation represents temporally averaged CBH at zenith, while CSS+USI CBH represents spatially averaged CBH. Therefore, random differences between ceilometer CBH and CSS+USI CBH are expected. In summary, the method was generally accurate for low clouds and although it is rare to observe alto-cumulus and cirrus clouds in coastal southern California, May 22 and July 1 confirmed the robustness of the method under those conditions.

Table 3.1: Daytime average ceilometer, METAR, and CSS+USI cloud base height and difference metrics between ceilometer and CSS+USI. The last line provides the average of the entries in the 27 rows.

Date	METAR [m]	Ceilometer CBH [m]	$\mathrm{CBH}_{\mathrm{CSS+USI}}$ [m]	RMSD [m]	nRMSD [%]
04/05/15	3536	788	848	108	13.7
04/20/15	793	650	707	76	11.7
05/02/15	2101	424	490	80	18.9
05/04/15	1346	782	841	230	29.4
05/10/15	4577	441	495	73	16.6
05/20/15	4904	851	1013	170	20.0
05/22/15	1107	1421	1500	305	21.5
05/29/15	6631	350	443	100	28.6
06/02/15	504	450	498	55	12.2
06/04/15	670	849	948	145	17.1
06/05/15	740	595	680	145	24.4
06/07/15	460	359	385	41	11.4
06/16/15	2840	355	420	80	22.5
06/18/15	365	288	320	38	13.2
06/25/15	1759	390	386	30	7.69
07/01/15	2438	5864	5245	830	14.2
07/03/15	498	345	398	55	15.9
07/08/15	708	736	841	200	27.2
07/09/15	4676	979	976	192	19.6
07/13/15	374	348	358	22	6.32
07/16/15	609	494	521	39	7.89
07/17/15	806	450	452	46	10.2
07/22/15	965	411	420	117	28.5
07/23/15	500	415	355	68	16.4

Continued on next page.

Table 3.1 – Performance metrics of CBH estimates, continued from previous page.

Date	METAR [m]	Ceilometer CBH [m]	$\text{CBH}_{\text{CSS+USI}}$ [m]	RMSD [m]	nRMSD [%]
07/24/15	550	340	332	21	6.21
07/26/15	470	400	458	73	18.3
07/29/15	540	444	495	70	15.8
All Days	1683	748	771	126	16.9

3.4.2 CBH Validation Examples for Two Days

Two detailed examples are analyzed in this section to further illustrate and explain the performance of the CSS+USI method. Fig. 3.9 shows the CBH comparison of ceilometer measurements, METAR, and the CSS+USI method for May 22, a day with different cloud types and multiple cloud layers. The period from 16:00 to 17:30 UTC is characterized by nearly overcast stratus clouds at 2,000 m AMSL that turn into alto-cumulus at the same altitude. During 18:30-21:45 UTC, scattered cumulus dominate, while after 21:45 UTC, broken cumulus are observed. UTC lags local standard time (PST) by 8 hours.

In the middle plot of Fig. 3.9, both CBH from local ceilometer measurements (the ground truth) and the CSS+USI method yield the same trend. For example, between 16:00-18:30 UTC, the CSS+USI method produces similar CBHs as the local ceilometer at about 2,000 m, while METAR reports 800 m which substantiates the concerns about using off-site METAR CBH data. At 18:30 UTC, ceilometer measurements indicate a CBH transition from about 2,000 m to 750 m, and the CBH from the CSS+USI method follows this transition, although with about a 300 m offset. After 20:00 UTC, an additional cloud layer with a different direction and variable speed, temporarily confuses the $\text{CBH}_{\text{CSS+USI}}$, as evident in a briefly elevated CBH around 20:15 UTC, 21:15 UTC and 22:15 UTC. However, the CSS+USI method still captures the CBH transition detected by the ceilometer from 800 m to 1,500 m at 22:00 UTC, and follows the ceilometer measurement until the end of the day. Again, METAR CBHs differ after 22:00 UTC indicating spatial heterogeneity in CBH. In summary, the CSS+USI method is accurate on this day especially in the morning. The daily RMSD is 305 m and nRMSD is 21.5%.

July 8 is analyzed in Fig. 3.10 as an example of a day with one of the largest observed nRMSD (27.2%). On this day, there are unusual fluctuations in cloud pixel speed reported from 19:30 UTC to 22:00

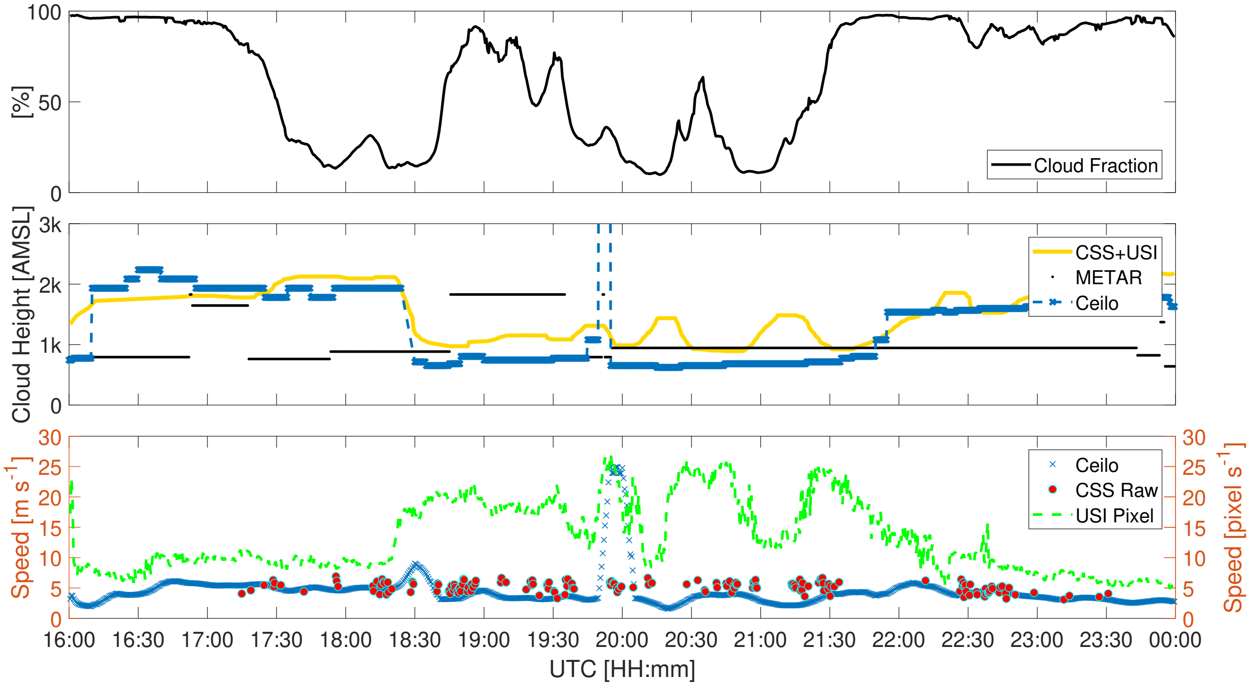

Figure 3.9: Sample comparison among different CBH measurements during the daytime of May 22, 2015. See Fig. 3.2 for locations of the instruments. Top: USI cloud fraction in units of %. Middle: CBH comparison between local ceilometer measurements (blue crosses), and the proposed CSS+USI method described in Section 3.3 (yellow line). The black dots indicate the measurement from airport METAR at Miramar Naval Air Station (KNKX), 8.8 km to the east of ceilometer. Bottom: Cloud speed determined by the CSS and USI. The green dashed line shows U_{pixel} (right *y*-axis). The blue line represents the cloud speed U_{USI} m s^{-1} calculated by Eqn. (3.3) with the CBH input from the local ceilometer measurements, while the red dots show the raw measurements from CSS. The USI pixel speed is not expected to match, but the other two methods are expected to match. Note that the brief period of ~ 25 m s^{-1} USI+Ceilo cloud speed at 20:00 UTC is a result of ceilometer measurements of CBH = 7,500 m which are cut off the middle graph for readability of the CBH variation.

UTC, especially a brief period of significantly smaller pixel speeds around 20:30 UTC, which causes a large CBH peak at that time. Visual inspection of the cloud images indicates that these fluctuations are not representative of the actual cloud motion, though the exact reason that the USI motion algorithm performs poorly is unclear. Regardless, this illustrates again that the accuracy of the CBH estimate depends on the quality of cloud vectors from both the USI and the CSS.

3.4.3 Assumptions and Limitations

In this section, the improvement and possible reasons for CBH errors are further discussed. Its performance is further compared to a prior method introduced by [44].

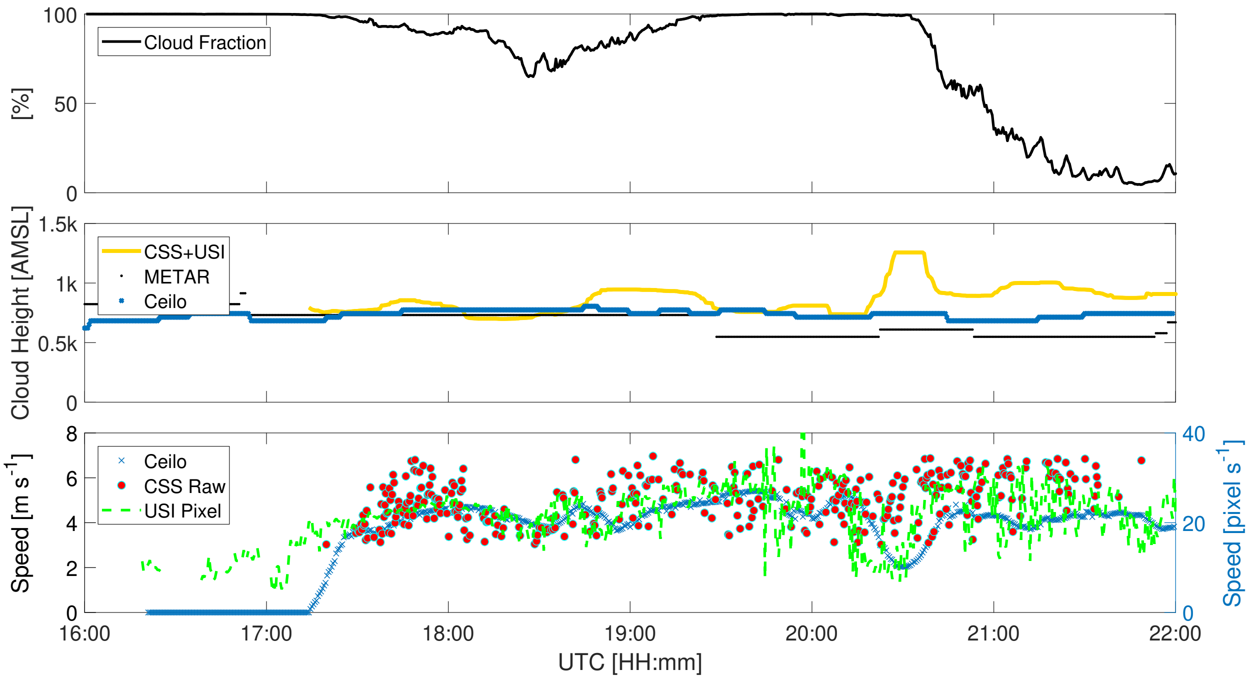

Figure 3.10: Same as Fig. 3.9, but for July 8 illustrating a case when unstable cloud pixel speed determination causes a large offset of local CBH estimates.

As implemented in Section 3.3.2, the LCE assumption implies that only the component of the velocity that is perpendicular to the cloud edge is detected. This assumption can cause offsets in determining CMVs, which is illustrated in Fig. 3.11. The cloud edge initially shades the central sensor at $t = t_0$, and then moves in one of two ways until it shades sensor 6. (i) It moves perpendicular to the cloud edge with speed v_1 and reaches sensor 6 at $t = t_1$, which is consistent with LCE assumption. (ii) It moves in a non-perpendicular direction with speed v_2 whose component normal to the cloud edge is v_1, and also reaches sensor 6 at $t = t_1$. In these two cases the signal measured by sensor 6 would be identical. Therefore, no matter what the direction of the CMV, the LCE-CFM will only detect the cloud speed component perpendicular to the cloud edge (here v_1). Thus, if the CMV is not perpendicular to the cloud edge, the cloud speed is underestimated, and subsequently, the lower CSS measurements causes a lower local CBH according to Eqn. (3.4). This is the main limitation of the linear cloud edge assumption.

For an infinite linear cloud edge, the cloud positions resulting from v_1 and v_2 in Fig. 3.11 are indistinguishable, while for real (finite) clouds, the cloud positions will be different. Bosch et al. [44] addressed this ambiguity by assuming that successive clouds passing the sensor move with the same CMV as they are transported by air at the same height in the boundary layer. Two successive clouds that pass the

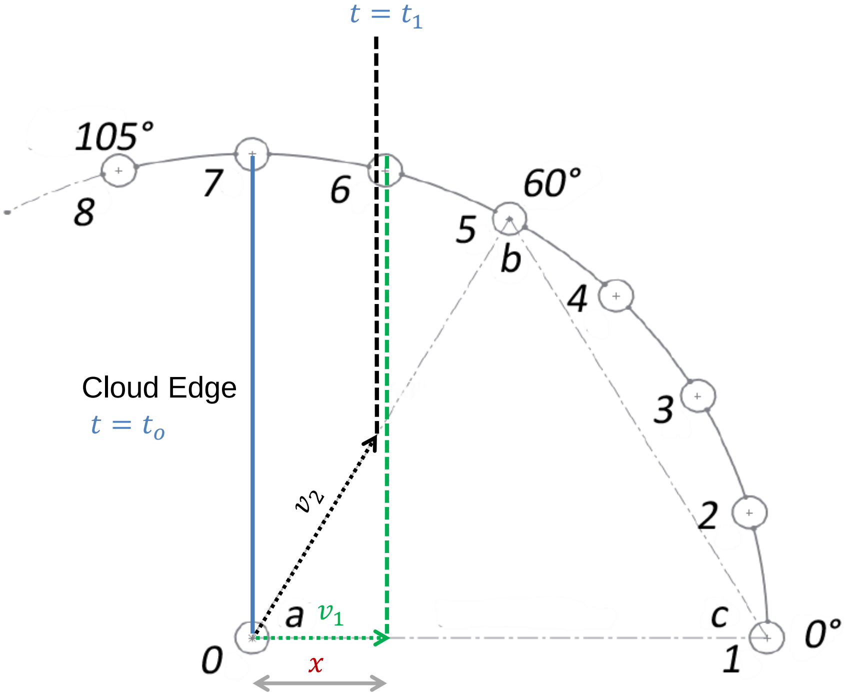

Figure 3.11: Illustration of a thought experiment that shows LCE-CFM method can only measure the velocity component perpendicular to the cloud edge due to a limitation of the linear cloud edge assumption. The blue line represents the original cloud edge and the vertical green dashed line represents the future position associated with the CMV v_1, while the black line indicates the future position associated with the CMV v_2.

sensor array with CMV v_{real} and different edge orientations will record velocities $v_{\perp_1}$ and $v_{\perp_2}$, at angles $\phi_{\perp_1}$ and $\phi_{\perp_2}$ as shown in Fig. 3.12. The true velocity v_{real}, can then be found as:

$$|v_{\text{real}}| = \frac{|v_{\perp_1}|}{\cos\left(\phi_{\perp_1} - \beta\right)} = \frac{|v_{\perp_2}|}{\cos\left(\phi_{\perp_2} - \beta\right)}, \tag{3.6}$$

which requires the angle of the true velocity β:

$$\tan\beta = -\frac{|v_{\perp_1}|\cos\left(\phi_{\perp_2}\right) - |v_{\perp_2}|\cos\left(\phi_{\perp_1}\right)}{|v_{\perp_1}|\sin\left(\phi_{\perp_2}\right) - |v_{\perp_2}|\sin\left(\phi_{\perp_1}\right)}. \tag{3.7}$$

However, as can be seen in Eqns. (3.6) and (3.7), v_{real} and β are sensitive to noise when $\phi_{\perp_1}$ is approximately equal to $\phi_{\perp_2}$. We have therefore opted to leave a more complete implementation of this method as future work. For the present analysis, we assume $v_{\text{real}} = v_{\perp_1} = v_{\perp_2}$ and use temporal averaging of motion vectors. This is expected to produce approximately correct direction vectors, since detected velocities are distributed about v_{real}, but systematically underestimates the speed (vector magnitude) slightly, because all potential $v_\perp$ are shorter than v_{real}. The underestimation varies quantitatively depending on the cosine of the cloud edge orientation bias as per Eqn. (3.6).

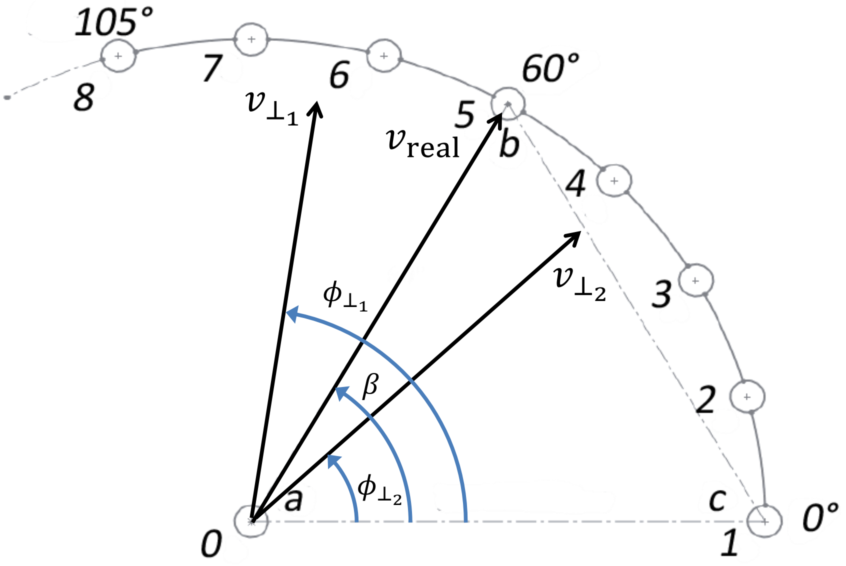

Figure 3.12: Determining real cloud velocity from perpendicular components. v_{real} is real cloud speed with angle of β in reference to horizontal line (a-c). $v_{\perp_1}$ and $v_{\perp_2}$ are the CMVs perpendicular to the detected cloud edge from two different passing clouds, and their angles are $\phi_{\perp_1}$ and $\phi_{\perp_2}$ in reference to line (a-c), respectively.

The original LCE method was developed by [44] for a sensor triplet in any non-linear configuration and spacing and CMVs are solved by geometric-kinematic equations based on the cloud arrival time at different sensors. While the sensor setup differs, the basic kinematic analysis of the original LCE method and the LCE-CFM that relies on LCE assumption is similar; a linear cloud edge passes over the sensors and causes different arrival times based on sensor arrangements relative to the CMVs. But two main differences do exist between two methods. i) The original LCE method develops equations to solve two unknownsspeed and directionusing two data points. In contrast, the LCE-CFM uses 12 data points to solve for the same two unknowns. The resulting system is over-defined and therefore more tolerant to signal noise. This also explains why the original LCE method requires low noise signals and multiple quality controls to produce less scattered results but the LCE-CFM has more clustered CMV raw measurements without post-filtering. ii) As discussed above, the original LCE method provides a mechanism to account for the impact of CMV not being perpendicular to the cloud edge, while the LCE-CFM method returns the CMV perpendicular to the cloud edge. The difference is summarized in Table 3.2.

Table 3.2: Performance comparison between the original LCE and proposed LCE-CFM method.

	Original LCE	LCE-CFM
CMV distribution	High noise and scattered raw data	Low noise and clustered raw data
CMV limitation	None	Detect the CMV only perpendicular to the cloud edge

3.5 Discussion and Conclusions

The principal objective of this research is to introduce a combination of sensors and an algorithm to provide an accurate local CBH for sky imager solar forecasting. The combination of a cloud speed sensor and sky imager makes measurements of CBH more affordable and convenient compared to a ceilometer. Ceilometers cost about US$ 20k while the bill of materials for the CSS is less than US$ 400. Furthermore, a CSS could be directly integrated into the enclosure of a sky imager avoiding the need for a separate setup site, power and Ethernet connectivity. In contrast, a ceilometer is bulky and requires separate infrastructure.

Firstly, the linear cloud edge assumption of [44] is leveraged to propose a method (LCE-CFM) for CSS measurements. The method analyzes the similarity, i.e. the correlation, of luminance signals between pairs of sensors aligned in different directions. Unlike the original CSS method that only considered the time delay of the most correlated pair, all 12 pairs of sensors are utilized to fit a cosine function of cross-correlation time delay versus sensor pair direction. The approach is motivated by assuming a linear cloud edge passing over the array of sensors. If a good fit is observed, the cloud direction is determined as the angle with the maximum time delay of the cloud passage on the cosine curve fit. The cloud speed is then equal to the sensor spacing divided by that time delay. The advantages and limitations of the LCE-CFM are illustrated. The method is also compared to a prior LCE method proposed by [44].

CBH is derived by comparing CSS cloud speed measurements in [m s^{-1}] to cloud pixel speed in [pixel s^{-1}] from a single sky imager. Over 27 days, the CSS+USI method shows promising CBH results with average RMSD of 126 m and nRMSD of 16.9% compared to on-site ceilometer measurements. The CBH accuracy depends on the accuracy of both the CSS cloud speed and the USI cloud pixel speed, as well as their mutual agreement. While the cloud pixel velocity is identified based on CMVs in the entire sky image, the CSS measures the CMVs just of the clouds approaching the sun. This discrepancy limits the CBH accuracy. Also, multiple layers of cloud with different direction and/or speed could degrade the performance because both CSS and USI are only able to determine cloud speed of a single cloud layer. In addition, the accuracy is restricted by the fact that the linear cloud edge assumption requires that the cloud motion vector be perpendicular to the cloud edge, which causes an underestimation of cloud speed. Lastly, the validation suffers from inconsistent measurement areas: (i) the ceilometer measures clouds straight overhead, (ii) the CSS detects the clouds that obscure the sun, and (iii) the USI analyzed clouds within its field of view that is typically about 10 km^2. This could result in inconsistencies between the ceilometer and the CBH from the CSS+USI method.

Future efforts will focus on implementation of real cloud velocity estimates from perpendicular components of two different passing cloud edges. USI cloud speed detection could also be improved. For example, a CMV field derived from optical flow ([39]) could provide the localized information to associate the CMV of the cloud passing the CSS. Optical flow also enables detection of multiple cloud layers as well as their respective cloud pixel speeds. Finally, validation under different meteorological conditions more relevant to continental climates would further substantiate the general applicability of the method.

3.6 Appendix

3.6.1 Validation of the LCE method

Figure 3.13 illustrates the direction offset between the direction of 0 s time shift ($\Delta t_{ij} = 0$) and the direction that is determined by the LCE-CFM. For example, in Fig. 3.4, the direction determined by the LCE-CFM method is 322.7°, while the direction closest to 0 s time shift is 240°, so the offset is -82.74°. Under the LCE assumption, these two directions should always be at right angles to each other; if the cloud edge is not linear, the offset will be larger or smaller depending on the shape of the cloud edge. The calculation is applied to all 27 days analyzed in this paper and the results are plotted in form of histogram in Fig. 3.13. Most of the angle offsets are clustered around -90° and +90° which indicates that the data are consistent with the LCE assumption.

3.6.2 Prior MCP Method Performance

Figure 3.14 illustrates that the prior MCP method suffers from some deficiencies as a result of arbitrary correlations from sensor noise, resulting in scattered CMVs outputs. Filtering can address the CMVs variability issue, but at the same time reduces the response of the sensor to sudden changes in cloud velocity.

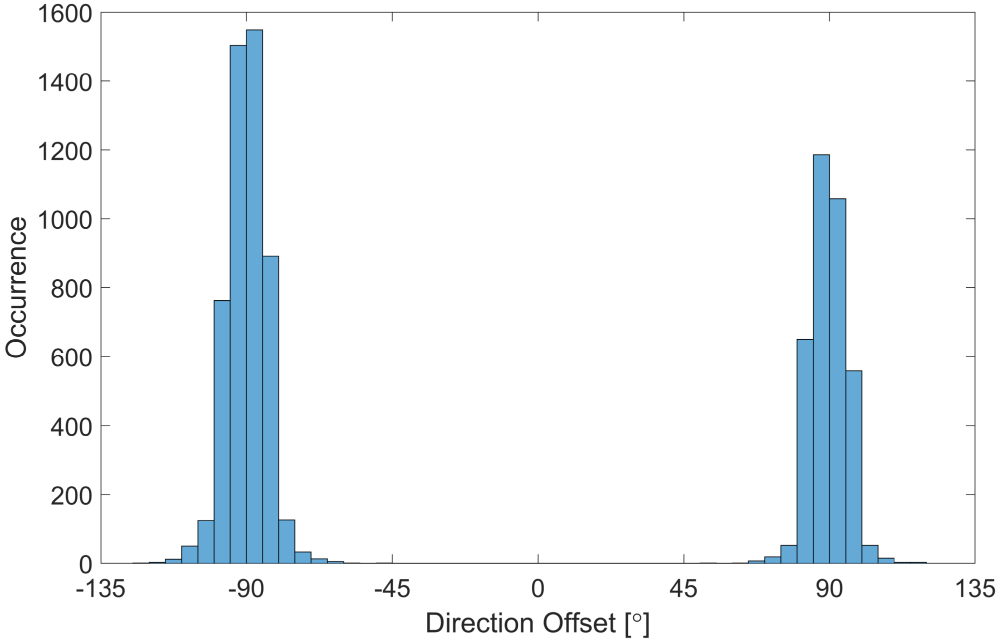

Figure 3.13: Histogram of LCE assumption validation on all 27 days analyzed in this paper. The *y*-axis represents the number of CMVs determined by the LCE-CFM using 9-sec segments of CSS data, and the *x*-axis represents angle offsets between the cloud direction from the LCE-CFM and the direction from the sensor pair which has a time shift closest to zero.

Also, the cloud direction outputs are not continuous as the final direction can only lie along individual sensor pairs.

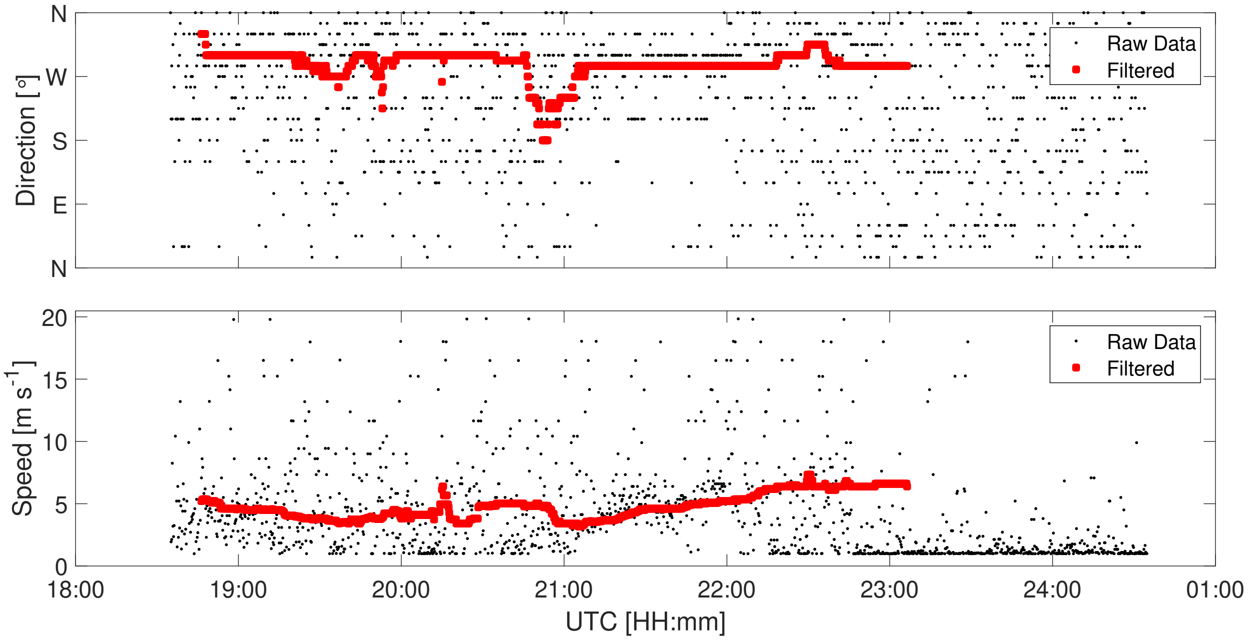

Figure 3.14: An example of the MCP method on July 24, 2013. Black dots show the raw measurement, and red dots show the filtered measurements after moving median filtering.

3.7 Acknowledgments

We acknowledge the donation of a ceilometer from Lawrence Livermore National Laboratory facilitated by Julie Lundquist. We also thank Victor Fung, Juan Luis Bosch, and Dominic Fong for assisting with CSS maintenance and operation.

This Chapter, in full, is a reprint of the material as it appears in “Wang, G., Kurtz, B. & Kleissl, J. (2016): Cloud base height from sky imager and cloud speed sensor, *Solar Energy*, 131, 208-221”. The dissertation author was the primary investigator and author of this paper.

Chapter 4

Maximum Expected Ramp Rates Using Cloud Speed Sensor Measurements

4.1 Introduction

The power output variability by large-scale grid-connected photovoltaic (PV) systems can negatively affect power quality and grid network reliability. Regulations have been introduced to restrict the maximum power ramp rates for PV plants on 1 min timescales [49]. These restrictions typically invite one of two approaches: (1) Compensate the power variability through energy storage systems (ESS). The storage requirements and strategies to comply with the regulations, considering capacity losses and cycling degradations, have been comprehensively studied [50–53]. (2) Curtail the PV output to smooth up-ramps reactively and provide a buffer for smoothing down-ramps proactively [54, 55]. For example, short-term forecasts for future cloud arrivals allow a system operator to meet ramp rate restrictions with less battery reserve or curtailment [56–58]. If all ramp rates (except plant outages) were to be mitigated, approach (1) would require knowledge of the worst-case ramp rate to determine the power and energy rating of the ESS. Given perfect forecasts, approach (2) could mitigate all power ramps without ESS. In practice, however, significant errors in the short-term ramp forecasts combined with restrictive ramp rate compliance requirements typically still require an ESS to mitigate worst-case ramp rates, but accurate forecasts can reduce the number of charge/discharge cycles of the ESS.

Although the nature of PV power variability has been well-studied [15, 59–61], existing approaches to simulate power fluctuations are expensive or cumbersome. For example, Kuszmaul et al. [62] created a network with synchronized sensors throughout the plant and approximated power output fluctuations as the average of the sensor readings. Marcos et al. [63] studied the smoothing effect of power fluctuations over the area of the power plant by low-pass filtering irradiance measurements at a single point, but the model needs to be empirically tweaked for each PV plant. Lave et al. [64] proposed a wavelet variability model to simulate the reduction in power output fluctuations of a plant or a fleet of dispersed plants. The correlation scaling coefficient introduced in the model is universal and a function of cloud speed [65]. Marcos et al. [66] simulated the power output by a fleet of plants, using only irradiance measurements at a single location and the smoothing effect due to geographical dispersion and plant size.

To the authors knowledge, all solar power high frequency variability models in the literature require high frequency (O(1s)) solar irradiance measurements [67, 68]. The sensitivity of solar variability model performance influenced by low frequency solar irradiance data is demonstrated in [15]. One-minute resolution data was found to be barely acceptable for rough simulation with the induced errors within the margin of other modeling errors. While high frequency solar irradiance data are rarely available, some applications, such as PV plant design and ESS sizing, require only the worst power fluctuations. For example, the worst ramp rate determines the required power and energy capacity rating for the ESS to buffer all down-ramps. Motivated by this, we propose a novel analytical approach, where the maximum expected PV ramp rate is computed in a process-based model using: 1) low-resolution point irradiance or PV power measurements, 2) the geometrical layout of PV plant, and 3) cloud velocity. During a cloud passage, the power ramp amplitude depends on the cloud optical depth, and the ramp rate is a function of both the optical depth and how fast a moving cloud passes over the plant [53]. This paper utilizes cloud motion vectors (CMVs) at ground level measured by our unique and recently upgraded cloud speed sensor (CSS) [31, 34].

The main contribution of this paper is the derivation and demonstration of a simple analytical model that bounds the maximum ramp rate. As CMVs are a key input of the model, we also take this opportunity to test and validate the performance of the improved version of the CSS, which measures CMVs more accurately in real-time and at a higher sampling rate.

The remainder of this paper is organized as follows. Sections 4.2.1 and 4.2.2 present the ramp rate model as an analytical relation between cloud velocity, PV plant dimension, recent PV power measurements / cloud optical depth, and expected maximum ramp rate. Section 4.2.3 introduces metrics to evaluate the model. Section 4.3.1 describes the experimental setup and data. Section 4.3.2 introduces recent CSS advancements in hardware design and real-time data processing, and Section 4.3.3 introduces the process to derive the data input for the model. Sections 4.4.1 and 4.4.2 examine the model performance through detailed analysis for an example day and validation over a 10-month period. Sections 4.4.3 to 4.4.7 discuss the performance, application, and benefits of the model in detail. Lastly, Section 4.5 provides the conclusions.

4.2 Mathematical Derivation and Problem Formulation

4.2.1 A Geometric Ramp Rate Model

The model is based on the following three assumptions:

1) A known irradiance field moving at a constant cloud motion vector impacts a PV plant. The irradiance field is larger than the size of the PV plant. The irradiance field contains a homogenous cloud layer with a stationary optical thickness over the time period when the ramp occurs. These assumptions are needed to apply a simple geometric ramp rate model.

2) No mismatch losses in the PV plant. Thus the PV plant power is proportional to spatially averaged irradiance. In reality, PV system efficiency in partly cloudy conditions would be expected to be reduced due to partial shading and mismatch between the output of different cells on a string;

3) Constant PV efficiency, i.e. no cell temperature or inverter effects. Assumptions 2) and 3) simplify the irradiance-to-power model.

With these assumptions, a power ramp results solely from the movement of the irradiance field over the PV plant (of size S), as conceptualized in Fig. 4.1. For simplicity, but without loss of generality, the irradiance field is represented in Fig. 4.1 as a cloud (bordered by the square grey rectangle) embedded within a clear sky (round rectangle). The ramp rate is modeled based on the interaction between the irradiance field (round rectangle) advected by the CMV and the PV plant geometry. At the initial time t_0 (Fig. 4.1a), the cloud (grey rectangle) covers a part of the PV plant. The irradiance field (and the cloud embedded in it) is then advected during a small time interval Δt into a new position (Fig. 4.1b). During the advection, part of the irradiance field (the grey area, denoted as the outgoing portion ΔS_{out} moves off the plant while the complementary upwind part moves in (the red area, denoted as the incoming portion ΔS_{in}). ΔS_{out} and ΔS_{in} are of the same size ΔS which can be computed as:

$$\Delta S = (\mathrm{L}v|\cos\alpha| + \mathrm{W}v|\sin\alpha|)\,\Delta t - (v\Delta t)^2\,|\sin\alpha\cos\alpha|, \tag{4.1}$$

*w*here L and W are the dimensions of the PV plant, and v is the speed of the irradiance field. The direction of the irradiance field α is defined with respect to the boundary of the PV system, which in this case is aligned with the North-South direction. The difference in the averaged cloud optical thickness between ΔS_{out} and ΔS_{in} induces a power ramp event. In this case, since the average kt over ΔS_{in} is constant at the cloudy kt, a downramp is induced by the fact that a part of ΔS_{out} that was previously clear is being covered by the cloud.

For reference, the area-normalized clear sky power production $\hat{P}_{\mathrm{cs}}$ is calculated by (Eqn. 4.2) given the power production under clear sky condition P_{cs}:

$$\hat{P}_{\mathrm{cs}} = \frac{P_{\mathrm{cs}}}{\mathrm{LW}}. \tag{4.2}$$

The solar power that would be produced by the net area ΔS under clear sky condition can be expressed as:

$$P^* = \hat{P}_{\mathrm{cs}}\Delta S. \tag{4.3}$$

The solar power change ΔP caused by a change in cloud optical thickness between ΔS_{in} and ΔS_{out} is described using the clear sky index kt [45] as:

$$\Delta P = (kt_{\mathrm{in}} - kt_{\mathrm{out}})P^*, \tag{4.4}$$

where kt_{in} and kt_{out} respectively represent the average cloud optical thickness for ΔS_{in} and ΔS_{out} of the irradiance field. Since the irradiance field is a mix of cloudy and clear sky conditions, $kt_{\mathrm{in}} - kt_{\mathrm{out}}$ can be positive or negative. Finally, the ramp rate RR becomes:

$$\begin{aligned}\mathrm{RR} = \frac{\Delta P}{\Delta t} =&((\mathrm{L}v|\cos\alpha| + \mathrm{W}v|\sin\alpha|)\Delta t \\ &- (v\Delta t)^2|\sin\alpha\cos\alpha|)\frac{|kt_{\mathrm{in}} - kt_{\mathrm{out}}|P_{\mathrm{cs}}}{\Delta t\mathrm{LW}}.\end{aligned} \tag{4.5}$$

4.2.2 The Worst-Case Scenario Ramp Rate (WCS-RR)

Equation (4.5) is not intended for operational ramp rate forecasts since without a sky imager the upwind irradiance field that would be needed to quantify kt_i is generally not available. Instead, we consider

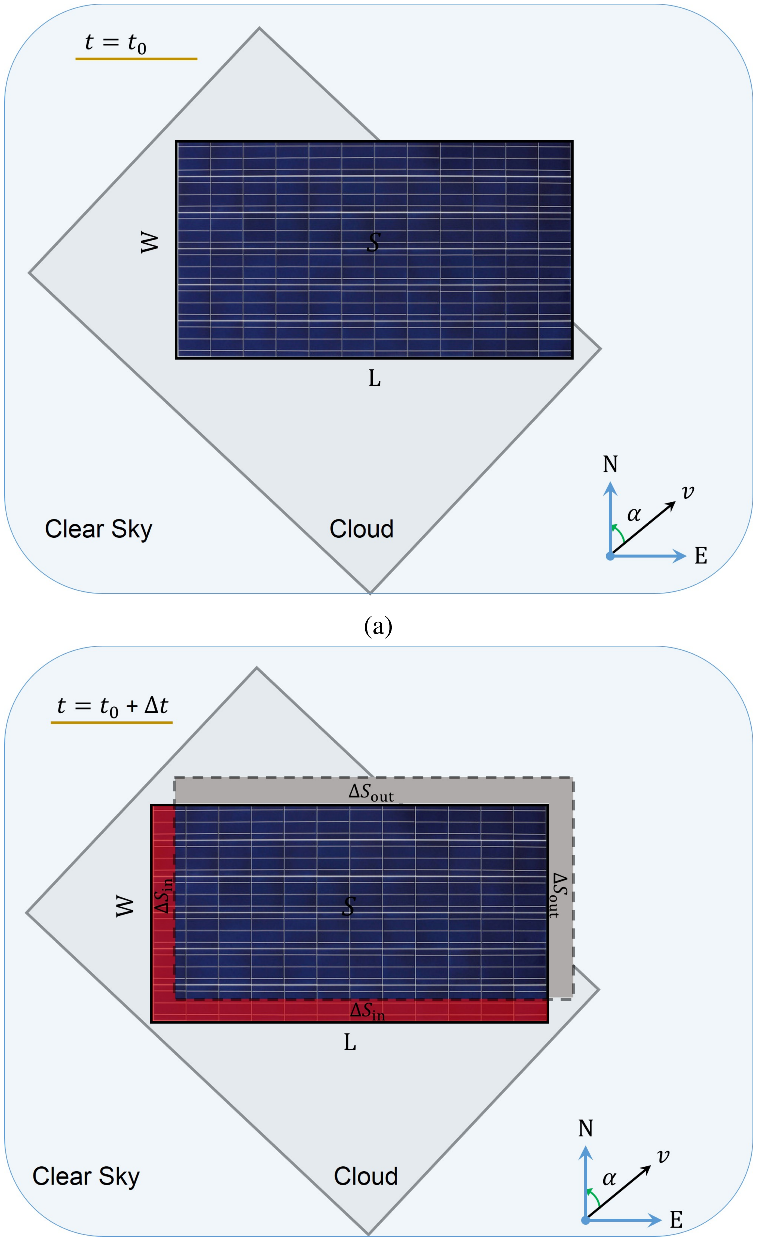

(a)

(b)

Figure 4.1: The irradiance field is defined by the outer round rectangle containing a cloud layer (inner rectangle) and clear sky (spaces in between). The cloud field over the PV plant of size S at t_0 is advected by speed v and direction α. The portion of the irradiance field that previously covered the plant (a) moves to a new location (b) during time interval Δt, resulting in incoming ΔS_{in} and outgoing ΔS_{out} portions of the irradiance field with $\Delta S_{\text{in}} = \Delta S_{\text{out}}$. The PV plant measures L by W and for convenience aligns with the East and North directions.

the worst-case scenario, where a clear sky gives way to an overcast sky. We estimate the largest ramp rate by picking the largest and smallest kt from recent history (e.g. 30 min) $kt_{\max}$ and $kt_{\min}$. The WCS-RR can be expressed as:

$$\begin{aligned} \text{WCS-RR} = \pm \Big(& (\mathrm{L}v|\cos\alpha| + \mathrm{W}v|\sin\alpha|)\Delta t \\ & - (v\Delta t)^2 |\sin\alpha\cos\alpha| \Big) \frac{|kt_{\max} - kt_{\min}| P_{\mathrm{cs}}}{\Delta t \mathrm{LW}}. \end{aligned} \tag{4.6}$$

Note that $kt_{\max}$ is not necessarily the kt in clear sky condition. For example, due to additional diffuse irradiance, cloud edge enhancement can cause irradiance to exceed the clear sky irradiance [60]. Equation (4.6) presents an analytical estimate of expected maximum ramp rate given cloud velocity, largest and smallest kt, and PV plant dimension.

4.2.3 Performance Evaluation

The WCS-RR estimate is evaluated by the following performance metrics. First, we define the compliance indicator σ by dividing the actual ramp rate by the corresponding WCS-RR estimate, as in:

$$\sigma(t) = \frac{\mathrm{RR}^{\mathrm{actual}}(t)}{\mathrm{RR}^{\mathrm{estimate}}(t)}. \tag{4.7}$$

When $\sigma \leq 1$, the actual ramp rate complies with the WCS-RR estimate. The maximum σ in each non-overlapping evaluation window of length m minutes is:

$$\mu_j = \max\Big\{\sigma(i)\,; i \in [nj, n(j+1)-1]\Big\}, j \in [1, N] \tag{4.8}$$

$n = (m \times 60)\,\mathrm{sec}/R$ represents the number of σ in the j^{th} evaluation window with a temporal resolution of R sec, and N is the total number of evaluation windows computed by rounding up the expression T/m, in which T is the overall daily time window of WCS-RR in minutes. The selection of m is somewhat arbitrary: a shorter window length results in more windows with exclusively clear or overcast conditions which are not of concern for ramp rates while longer window lengths tend to evaluate σ too infrequently. Since transmission system operators are typically required to counteract power fluctuations with load following at a time scale of less than 30 min, we apply window lengths of $m = 2$, 10, and 30 min. The compliance rate ξ is defined as:

$$\xi = \frac{N_{\text{cpl}}}{N}. \tag{4.9}$$

The number of compliance events N_{cpl} indicates the number of windows that satisfy $\mu_j \leq 1$. Subsequently, the noncompliance rate becomes:

$$\epsilon = (1 - \xi) \times 100\%. \tag{4.10}$$

While risk-adverse actors would prefer that the WCS-RR always envelopes the observed ramps, excessive WCS-RR may result in an over-sized energy storage system. To quantify the extent to which the WCS-RR overpredicts the actual ramp, all compliance events are further evaluated by the degree of overestimation δ:

$$\delta = \left(\frac{1}{N_{\text{cpl}}} \sum_{j \in N_{\text{cpl}}} 1 - \mu_j \right) \times 100\%. \tag{4.11}$$

4.3 Validation Setup and Data

To validate the proposed method, we set up an experiment at the University of California San Diego (UCSD) campus testbed. The validation is inconsistent with the stated objective of the method, and this paragraph clarifies and resolves these inconsistencies. The objective of the WCS-RR method is to size ESS systems that can mitigate the worst-case situations to ensure a 100 % compliance rate with ramp rate restrictions. To satisfy this objective, only the single worst ramp rate is needed, and time-resolved ramp rates are unnecessary. While the WCS-RR method would only be validated with a single data point per site, we instead validate a time-series implementation of the method where the WCS-RR are compared within evaluation windows against observed worst ramp rates. While not necessary to the objective of the method, our present time-resolved validation approach provides a larger validation sample size and allows relational analysis between the time-resolved physics variables, such as cloud speed and cloud optical depth, and the modeled and measured ramp rates.

A note on the time resolution of the input data is also in order. As advertised, the advantage of the method is that it can be applied to coarse resolution (e.g., hourly) input data. Such data is commonly available: for example, hourly output of cloud optical thickness (or GHI) and CMV from Numerical Weather Prediction (NWP) or reanalysis could be input to Equation (4.6) to create a time-series of WCS-RR. Alternatively, hourly GHI from ground stations and CSS data or NWP CMV could be used. However, higher resolution data will benefit the accuracy of the method as some short-lived cloud phenomena, such as irradiance enhancement or deep convection, may generate worse ramp rates (larger cloud optical depth and/or larger cloud speed) than in the hourly data. Safety factors should be considered if hourly or coarser data is used. Lastly, if very high resolution (O(1s)) data is available, then power variability models, such as those presented in the introduction, are recommended as they provide more accurate information on the worst-case ramp rates and their time-series ramp rate estimates can be directly applied to energy storage simulation models.

4.3.1 PV data

Figure 4.2 illustrates the layout of two existing PV systems located on the EBU2 building at UCSD (32°52′53.1”N, 117°13′59.2”W) with a tilt angle of 20° and an azimuth angle of 225° east of north. The system under study is arranged in 3 arrays, consisting of a total of 181 PV panels (marked in red in Fig. 4.1 with overall dimensions of 33.5 m x 16 m. The total nominal power is 37.1 kW DC. The PV power was measured at a 2 sec sampling rate from September 29, 2017 to October 1, 2018 by 5 inverters, including two SMA Sunny Boy 5000US and three SMA Sunny Boy 7000US, with a total rated power of 31 kW AC. The PV database consists of 10 months of power output measurements. Excluding server shutdowns, rainy and overcast days, and clear days (defined as less than 30 min of cloud cover), 90 partially cloudy days remain. Only partial cloud cover is of interest to the experiment because it causes the largest power ramps. Note that the production field includes the ground area in between the rows, which mathematically enters (Eqn. 4.2) through the power plant dimensions. To avoid errors from clear sky and PV performance models, the power produced on the most recent clear day is used as the clear sky power.

4.3.2 Cloud Speed Sensor (CSS) and Recent Improvements

Cloud speed and direction are required to relate cloud field to ramp rates. While the proposed model is compatible with any type of cloud speed measurement, we obtain cloud speed measurements from our in-house CSS, which provides an accurate yet affordable means to measure local CMVs. Refer to Wang et al.

[31] for detailed design specifications and features. The instrument is installed on the same rooftop as the PV system and marked in Figure 4.2.

The CSS consists of a set of nine phototransistors arranged about a semi-circle. In the original CSS design, 666 Hz raw data were measured in 18 sec segments. For processing, data had to be sent to a remote server via an Ethernet card. In the upgraded CSS, the existing microcontroller Max32 still collects raw data, but now transmits the data to an attached Raspberry Pi through a much faster serial connection. This Raspberry Pi directly computes the CMV based on the Linear Cloud Edge (LCE) assumption, using a 9 sec set of 6,000 samples from the phototransistors. The updated hardware design processes each 9 sec batch of sampled data in only 2 sec, outputting the CMVs at 11 sec resolution. This resolution is sufficient for cloud motion estimation in power variability modeling. Continuous sampling would be feasible with additional software upgrades. The final processed CMV measurements are stored on-board along with the raw sensor data and timestamps. To support external real-time use of the data, CMV data is also published to the local network via the MODBUS TCP protocol. The upgraded CSS has been fully operational since October 8, 2017.

Figure 4.2: Aerial view of the PV system installed on EBU2 at UCSD and the PV production field (red) considered in this paper. The yellow star indicates the CSS located 10 meters north of the upper corner of the PV system. ©Drone Photo

Enhanced CSS Software to Extract Real Cloud Velocity

An enhanced CMV calculation algorithm addressing limitations of the LCE assumption is introduced in this subsection. The CSS algorithm in [31] assumes a linear cloud edge passing over the array of sensors given the fact that clouds are typically much larger than the spacing between sensors (0.078 m), and therefore only detects the component of the velocity that is perpendicular to the cloud edge, which systematically underestimates the speed (vector magnitude) slightly. Figure 4.3 illustrates that under the LCE assumption, the CSS would report different CMVs from a single cloud. To resolve this ambiguity that was left for future work in [31], we compute actual cloud velocity using the reported perpendicular cloud speeds $v^i_\perp$ and directions $\phi^i_\perp$ following Equations 6 and 7 in [31] through a weighted non-linear regression of $|v|$ and α to the N CMVs collected in a recent period Section 4.3.3 using:

$$\begin{gathered} w_i|v^i_\perp| = |\boldsymbol{v}|\cos\left(\phi^i_\perp - \alpha\right), \\ \text{where } w_i = \frac{1}{t_o - t_i}, \forall i \in N \end{gathered} \tag{4.12}$$

$|\boldsymbol{v}|$ represents the magnitude of the vector of actual cloud velocity $\boldsymbol{v}$. w_i represents the weighting factor computed as the inverse of the time difference between the end point of the time window (i.e. present time t_o and the timestamp of the i^{th} CMV t_i. As a result, the most recent CMV has the highest weight.

While v becomes unreliable when $\phi_\perp$ exhibits only small variations, small $\phi_\perp$ variation suggests only minor changes in the LCE orientation of the cloud field. Hence, if the collected N raw CMVs differ by less than 20°, the reported CMV is reasonably considered as the true CMV and the regression in Eqn. (4.12) is not conducted. Instead, the CMVs are decomposed into horizontal and vertical directions, and the median value of each is then used to recompose one median filtered CMV.

4.3.3 Search Time Window and Data Processing

Since CMVs are only available in irregular intervals, a search window length needs to be defined to average the CMV data and select the largest and smallest *kt*. Longer windows challenge the assumption of cloud field homogeneity, causing older cloud fields that are likely different from the ones at present to be counted. Shorter windows may not contain sufficient cloud cover events and falsely suggest that clear conditions will persist. Based on our experience, the cloud field in coastal Southern California is typically steady over a few hours, so we consider a 30 min window centered at the time of interest as a conservative

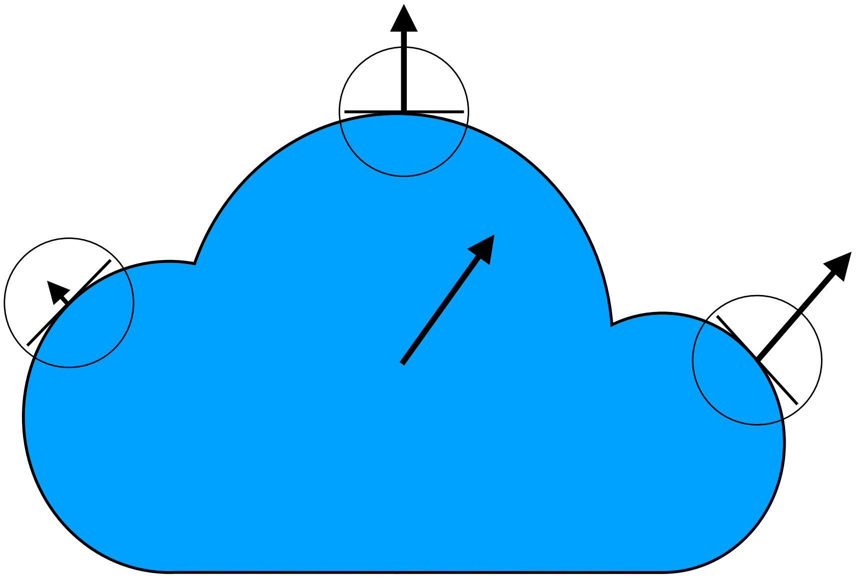

Figure 4.3: Schematic depiction of actual cloud velocity (centered arrow) vs. perpendicular components (other arrows) derived from different cloud edge orientations in the prior CSS algorithm.

upper bound. For areas with faster changing cloud pattern and more short-lived cloud phenomena, a smaller time window will is preferred.

The search window is also needed to determine appropriate kt values in Equation (4.6). kt is obtained by normalizing measured power output from the past 30 min is first normalized to kt using clear sky power output. The largest and smallest kt are selected for kt_{max} and kt_{min} respectively. If CMVs are not available in the 30 min search window, typically suggesting (near) clear or overcast conditions, the WCS-RR is not computed. Fortunately, in uniform sky cover conditions, the PV output variation is small and the associated ramp events are not important for PV planning applications. Lastly, since PV data is sampled every 2 s, the WCS-RR is also computed in 2 s intervals for consistency, but the analysis could be conducted with 15 min or hourly data.

4.4 Results and Discussion

4.4.1 Performance on a Sample Day

In this section, one detailed example is analyzed to illustrate and explain the performance of the proposed model. The WCS-RR estimate on June 21, 2018, one of the best days (100% compliance across all evaluation windows, refer to Table 4.2 in the Appendix), is validated against the actual ramp rates in Figure 4.4. The top plot shows real PV power from 10:30 to 15:15 PDT and clear sky power output from 2 days earlier. The clouds are observed to move eastward to southward over the day with speeds ranging from 2 to 5 m s^{-1}, as illustrated in the middle plot. The bottom plot illustrates the observed ramp rate and the WCS-RR estimates.

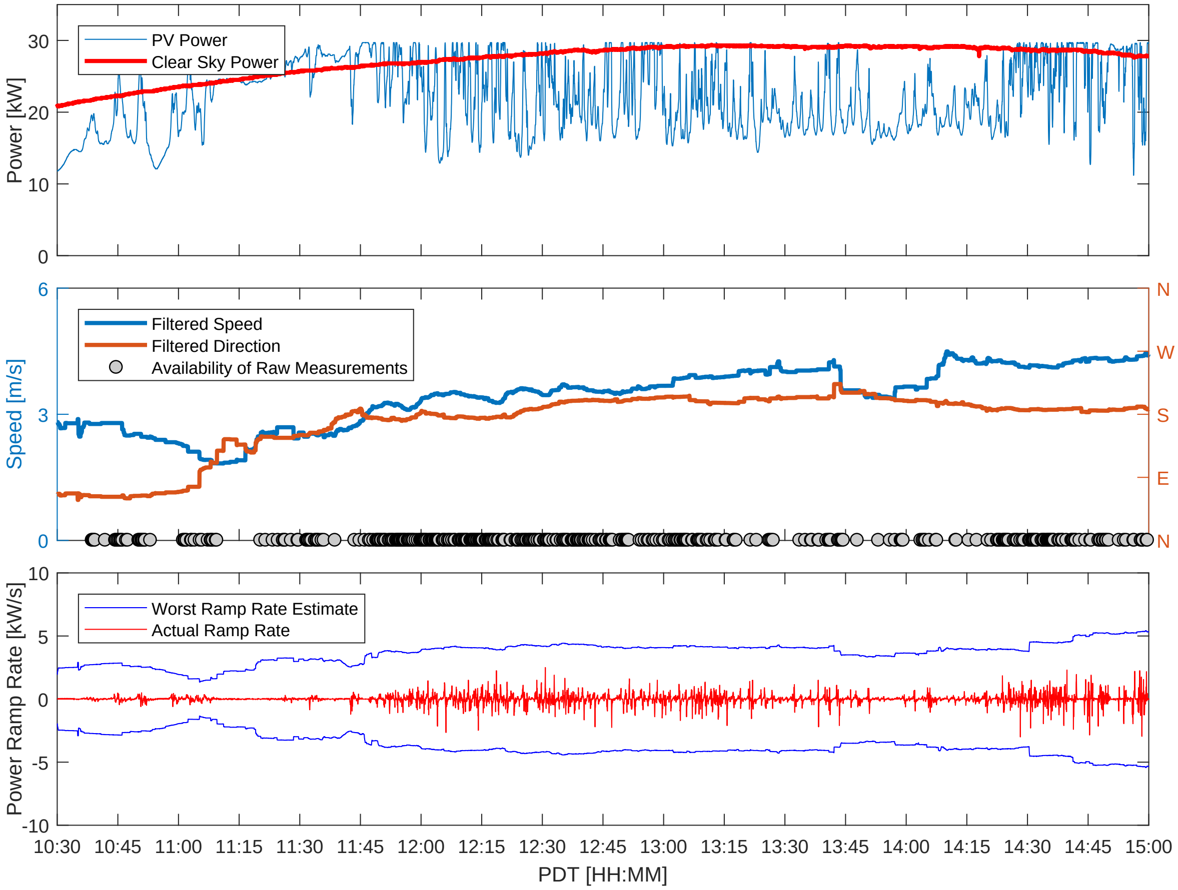

Figure 4.4: Example validation of the proposed method on June 21, 2018. Top: Actual PV power on June 21, 2018 (blue) and on the most recent clear day (June 19, 2018, red). Middle: The true velocity and directions (lines) are derived from non-linear regression (Equation 4.12) on the raw CSS CMVs measurements (their availabilities are indicated by dots) collected in the centered 30 min time window. Bottom: Comparison between actual ramp rate and WCS-RR estimate.

The WCS-RR generally is positively proportional to clear sky power and cloud speed, completely enveloping the actual ramp rate magnitude and timing. Since the cloud optical thickness is steady over the 5-hour period, kt variation does not have a strong impact on WCS-RR. There is a 30 min exception around 11:45 PDT when the sky was briefly clear, lowering the WCS-RR estimate and the difference between $kt_{\max}$

and kt_{min}, but the ramp events in that period are still enveloped successfully. The largest observed ramp rate of the day occurred at 14:28 PDT, during partial cumulus cloud cover (Figure 4.5). Because this cloud condition change (clear to worst-case thick clouds) is the exact circumstance modeled in Equation (4.6), the WCS-RR estimate accurately captures the ramp magnitude with only a 7% overestimate.

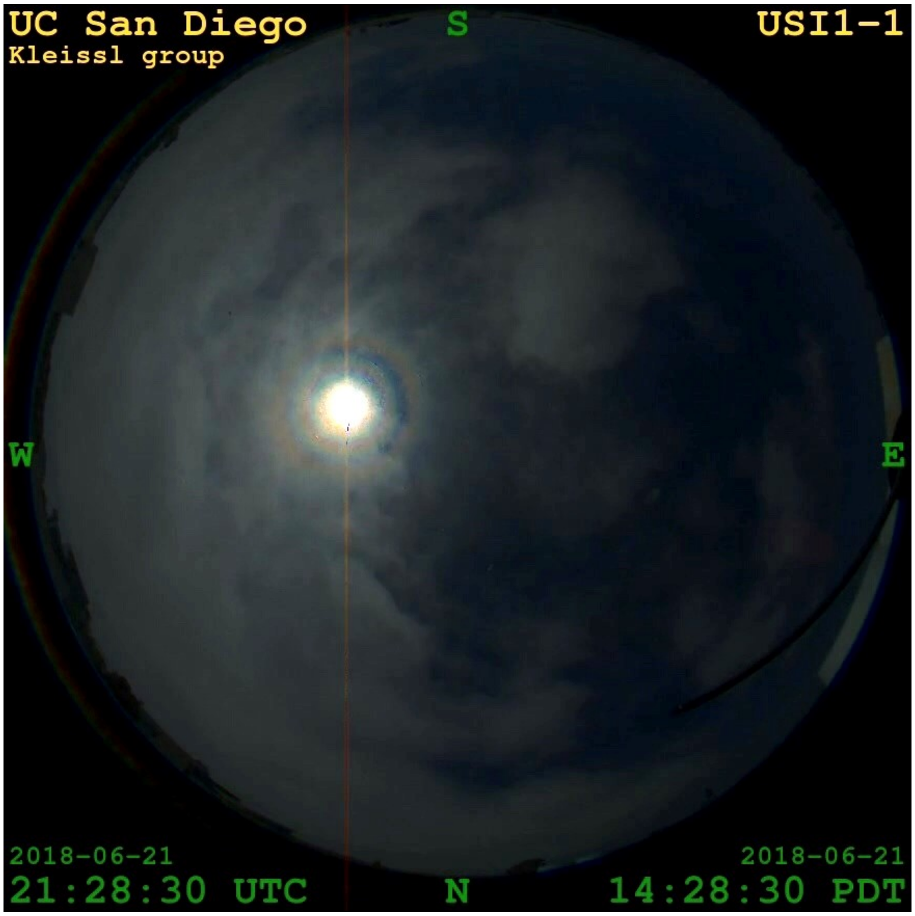

Figure 4.5: Sky image of a cloud shading event at 14:28 PDT on June 21, 2018 with the largest ramp of the PV plant. Figure 4.4 demonstrates that this largest ramp is accurately bracketed by the WCS-RR.

The daily noncompliance rate of 0% across all evaluation windows in this day confirms that the observed ramp rates are perfectly enveloped by the WCS-RR, at a cost of averaged ramp rate overestimates over the designated time window from 45% (30 min window) to 62% (2 min window). The ideal metrics would be 0% noncompliance and 0% overestimate. In reality, there is a trade-off between the noncompliance and the overestimate metrics: to bracket all large power ramps, the model conservatively assumes a cloud condition change from clear sky to thick clouds at all times, including periods with clear or overcast skies, which inevitably over-predicts the ramp rates. The only hypothetical scenario with 0% noncompliance and 0%

overestimate would be for a series of thick clouds of the same size as the PV plant to pass the plant along L (or analogously W); in those conditions, the plant would continually ramp up or down with a ramp rate equal to the WCS-RR. For June 21, 2018, the minimal overprediction of the largest ramp event (on 14:28 PDT) proves that the over-prediction is not excessive. The tradeoff between the overestimate and noncompliance rate will be further discussed in Section 4.4.5.

4.4.2 Aggregate Ramp Rate Statistics

The evaluation of the proposed method over an extensive set of 90 days is summarized in Table 1 and the performance metrics are illustrated by the box plot in Figure 4.6. A more detailed day-by-day performance comparison is tabulated in the Appendix. Overall, the method shows promise: for the shortest 2 min window, the average noncompliance rate is only 1.1% at a cost of a 64.3% overestimate. The noncompliance rate slightly worsens with wider evaluation windows, which is expected as the chance that a noncompliance event ($\sigma > 1$) is included in the evaluation window increases with wider windows. Nevertheless, even under the longest 30 min evaluation window, the average noncompliance rate is only 6%. The largest-ever observed ramp rate of 9.2 kW (or 29.7% PV capacity) per second on February 27, 2018 is successfully captured, and on this day, the noncompliance rate of 0% is achieved across all evaluation windows, further demonstrating that the proposed WCS-RR model functions as designed. The degree of overestimation worsens with shorter evaluation window length (greater number of windows). Because WCS-RR estimates are generally conservative at all times except the time of the daily largest ramp events (e.g., 14:28 PDT in Figure 4.4), the degree of overestimate would be minimal when only a single time window (i.e. the daily largest ramp event) is considered over the day but tends to be larger when more evaluation windows are considered.

Table 4.1: Noncompliance rate ϵ and the degree of overestimation δ for 2 min, 10 min, and 30 min evaluation time windows averaged over all data points in 90 days.

	$\epsilon_{2\mathrm{min}}$ [%]	$\delta_{2\mathrm{min}}$ [%]	$\epsilon_{10\mathrm{min}}$ [%]	$\delta_{10\mathrm{min}}$ [%]	$\epsilon_{30\mathrm{min}}$ [%]	$\delta_{30\mathrm{min}}$ [%]
Average	3.1	71.8	7.6	63.4	11.9	56.4

As the primary goal of the WCS-RR model is to estimate the largest possible ramp rate, the observed maximum ramp rates and associated binary daily success/failure flags are also tabulated for each day in Table 4.2. The WCS-RR successfully envelopes the maximum daily observed ramp events on 83 out of 90 days.

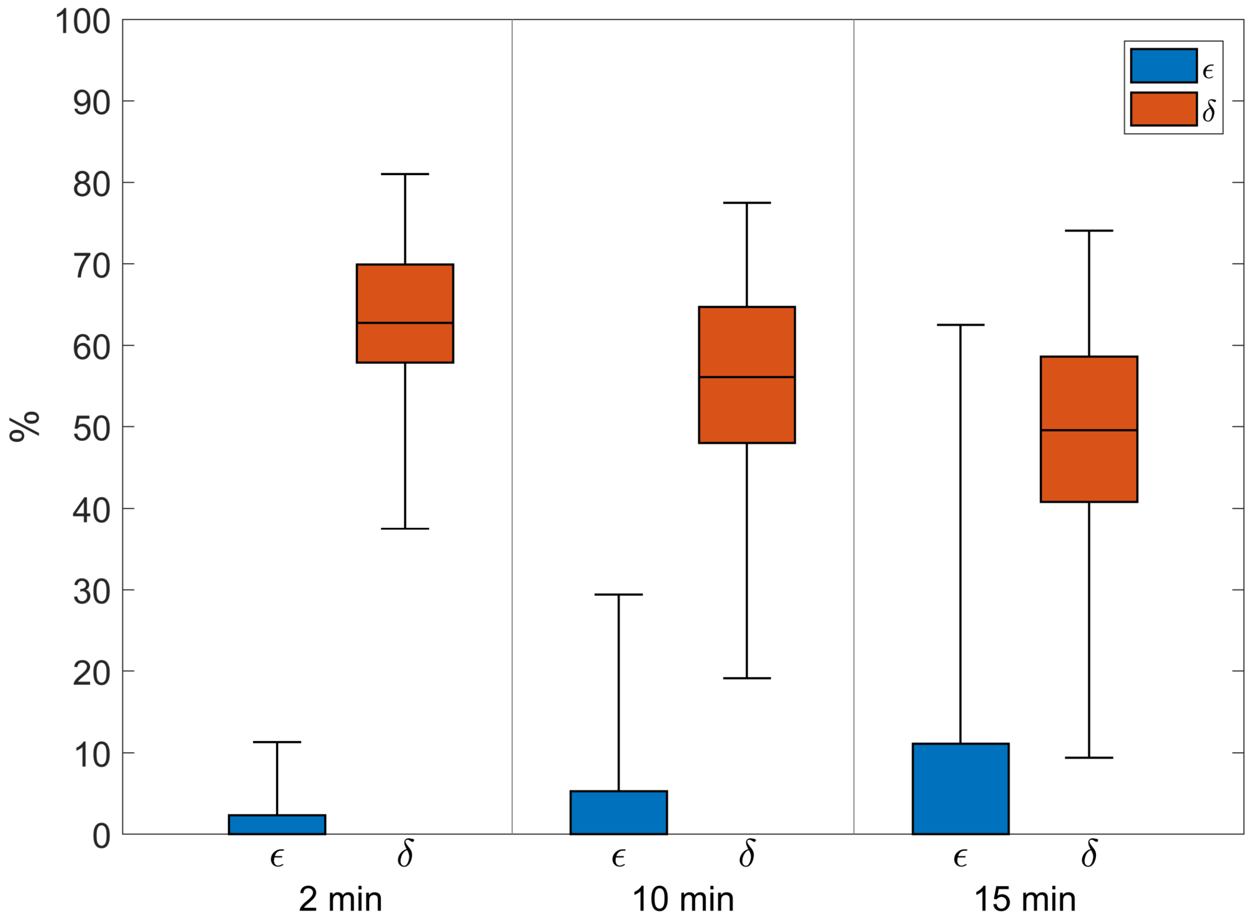

Figure 4.6: Box plot of noncompliance rate and associated overestimate for 2 min, 10 min, and 30 min evaluation time windows over 90 days. This figure is a visual representation of the performance metrics in Table 4.2.

For one of the remaining 7 days, the WCS-RR was not computed because CMV measurements were not available at the moment of the largest ramp rate. While ramp rate violations are observed in the other 6 days, the actual ramp rate exceedances are relatively small at [0.7, 0.2, 0.1, 0.8, 0.2, 0.3] kW/s (refer to the value in the parenthesis after the “F” flag in Table 4.2). The causes for these violations are discussed in the next section.

For PV system planning applications, statistics about the magnitude and frequency of noncompliance events are of interest. The distributions of ramp size and exceedance of all 129 noncompliance events ($\sigma > 1$) over the 90 days are displayed in Figure 4.7. The left histogram indicates that 92% of the missed ramps are less than 6% s^{-1}. The right histogram indicates that most ramp rate exceedances normalized by the rated AC power are less than 2% s^{-1}. Thus, Figure 4.7 illustrates that even when the real ramps exceed the WCS-RR estimate, the difference is small and most of the noncompliance events only involve small ramps compared to the largest observed ramp of 29.7% s^{-1} (February 27, 2018). While the ramps may seem large for utility-scale plant operators, relative ramp rates decrease with the size of the PV plant. Since Eqns. (4.5) and (4.7) have the solar power plant area LW in the denominator for the same irradiance field for example ramps for a 371 MW plant would be 1/10,000 those of our 37.1 kW plant. Therefore for example the 2% s^{-1} ramps should not be judged as an absolute number, but rather relative to the 29.7% maximum ramp rate for our particular PV plant.

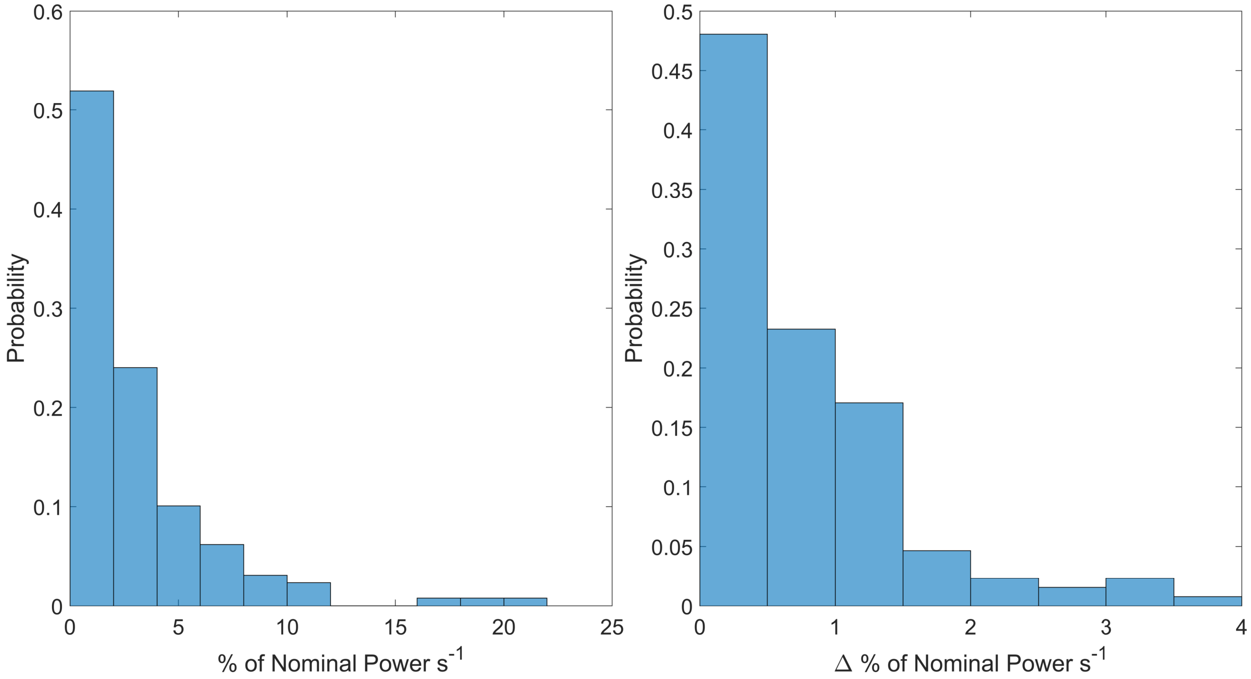

Figure 4.7: Distribution of noncompliance events by ramp size (left) and exceedance (right) over 90 days.

4.4.3 Example Noncompliance Events

In this section, the causes of noncompliance events are further investigated, and the limitations of the model and experimental setup are discussed. June 27, 2018 is analyzed as a representative day with the noncompliance events detailed in Figure 4.8. Three ramps exceeded WCS-RR from 16:00 to 16:30 PDT. The WCS-RR does not produce an estimate for the first ramp event and does not bracket the other two. Visual inspection of sky images reveals that the three power ramps were induced by small, dissipating stratocumulus clouds (e.g., red circle in Figure 4.9). The cloud cover fraction is likely too small to induce frequent illuminance fluctuations in the CSS, resulting in only two CMV measurements (dots in the middle plot, Figure 4.8) in the 30 min window, and no measurement at all within ± 15 min of the first ramp event at 16:00 PDT. Furthermore, the cloud is dissipating as it crosses the measurement site, and as a result, the cloud edge is not sharp enough to satisfy the assumption of LCE, which degrades the goodness of the curve fit [31]. Subsequently, the CMV measurement is likely inaccurate, which in turn degrades the accuracy of WCS-RR. Since these shortcomings in the CSS are specific to thin and sporadic clouds, they are not of concern for the main application of the WCS-RR model, which is to estimate the possible maximum ramp events for PV planning. In fact, the maximum ramp event of the day at 14:15 PDT was still successfully bracketed. This example day demonstrates that the accuracy of the proposed method primarily relies on the quality and availability of CMVs.

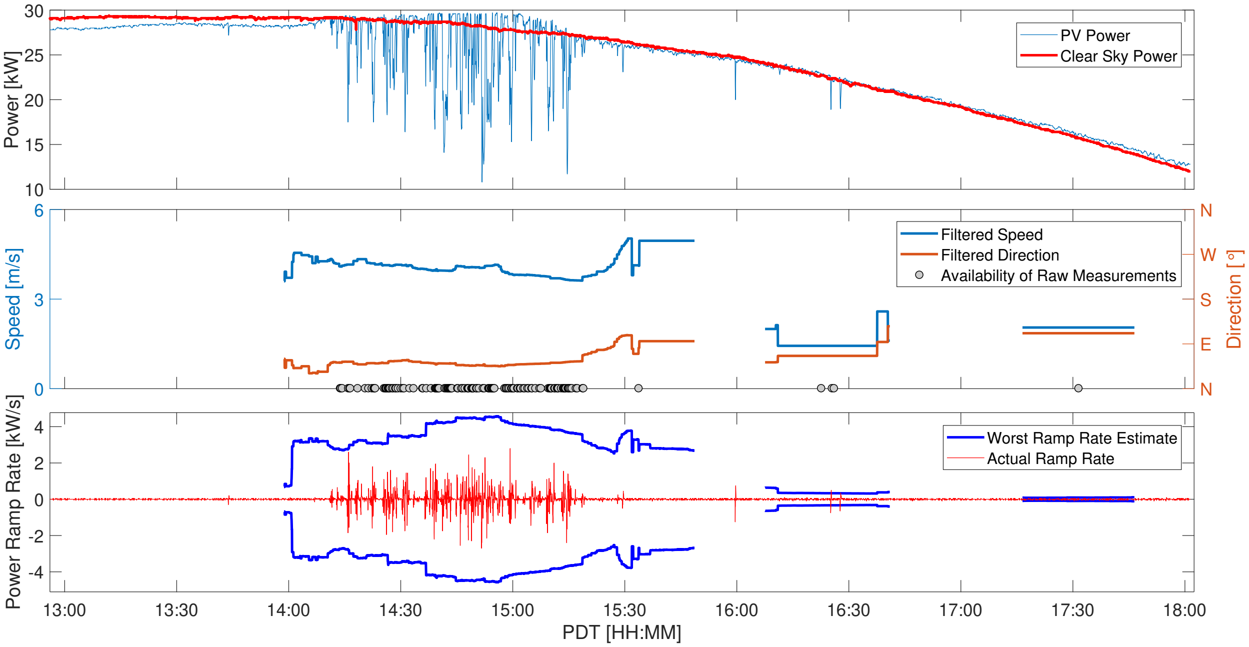

Figure 4.8: Example day when the observed ramp does not comply with the computed WCS-RR at 16:25 and 16:28 PDT. Top: real PV power on June 27, 2018 and the most recent clear day. Middle: Cloud speed and direction from the CSS. Bottom: Comparison between actual ramp rate and WCS-RR estimate.

By examining the sky images for all noncompliance events, we conclude that the primary reasons for noncompliance events are few and/or inaccurate CSS measurements. For example, for many noncompliance events, only one single CMV measurement appeared in the 30 min search window, which is an insufficient sample size. Sparse CSS measurements are typically caused by sparse cloud fields and are associated with significant uncertainty, which affects the performance of the WCS-RR method because CMV are key inputs to the WCS-RR (as per Eqn. 4.1). Sparse and/or small clouds may even go undetected by the CSS and

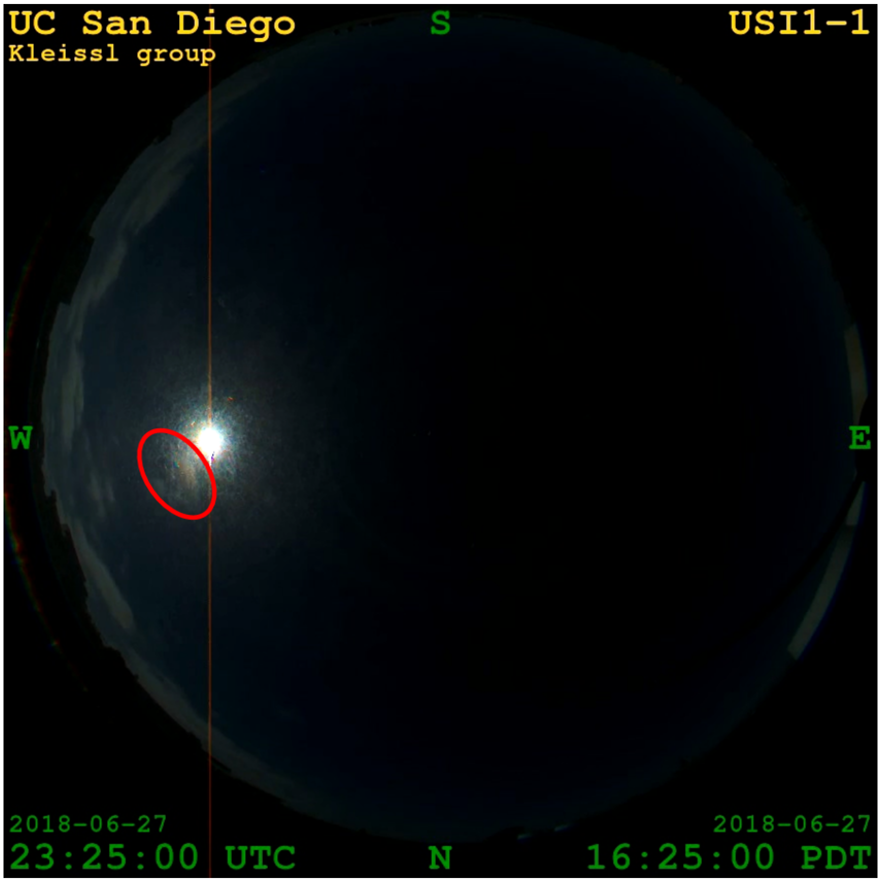

Figure 4.9: Example day when the observed ramp does not comply with the computed WCS-RR at 16:25 and 16:28 PDT. Top: real PV power on June 27, 2018 and the most recent clear day. Middle: Cloud speed and direction from the CSS. Bottom: Comparison between actual ramp rate and WCS-RR estimate.

thus, a WCS-RR estimate would not be produced. However, these clouds may not be large enough to cause meaningful ramps. Figure 4.10 shows the distribution of ramp rates for times with one or less CMV measurements in the 30 min search window over the 90 days. All ramp sizes are smaller than 10% PV capacity per second even for this small PV system, which confirms our expectation. For large PV systems, we expect sparse or small clouds to be even less relevant as they cover only a fraction of the PV plant.

4.4.4 Comparison to Other Simple Ramp Rate Estimates

The method provides a reasonable maximum RRs through a simple process model framework which characterizes the maximum RRs better than other even simpler methods (e.g. a constant ratio). For example, Figure 4.11 shows the distribution of observed daily maximum power ramp rate over 90 days normalized by PV systems nominal power. Based on this histogram a constant ratio of 30% s^{-1} ramp would be a good

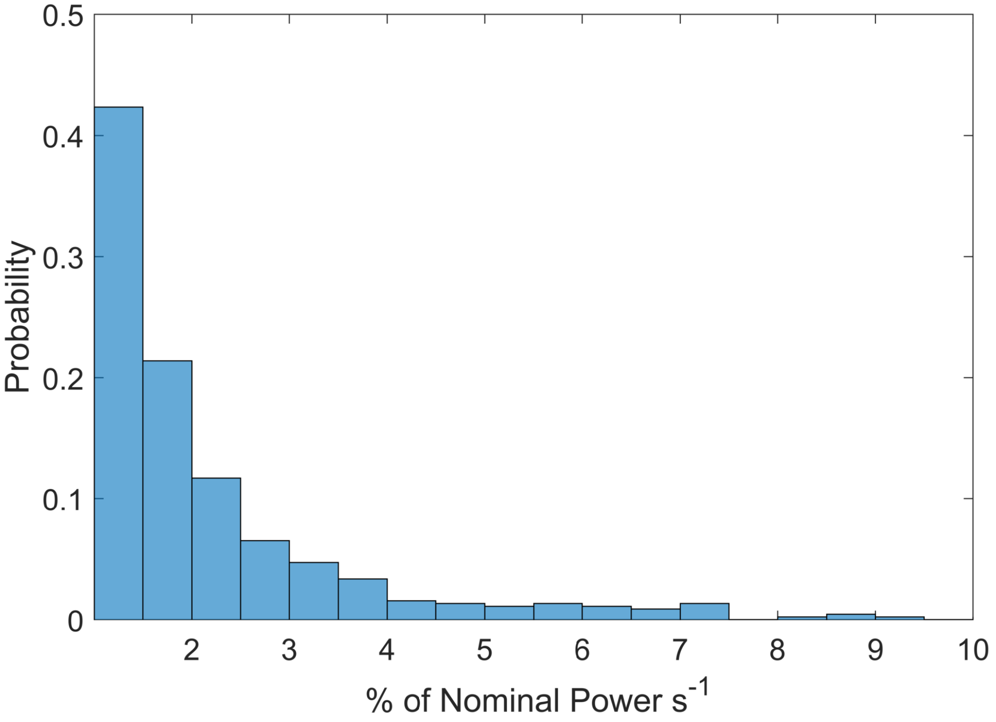

Figure 4.10: Distribution of ramp rates over 90 days when none or only one CMV measurement is available in the 30 min search window. Most of these ramp events do not comply with the respective WCS-RR estimates.

assumption for this specific PV system and climate. However, because the maximum ramp rate depends on the system size, orientation, prevailing wind direction and speed, and optical depth of the atmosphere, a universal RR assumption for a PV system is a crude assumption and not likely close to the ground truth unless it is prior known. For example, the PV system dimensions are directly accounted for in our model and they can vary by a factor of 100 more between a commercial rooftop PV system and a utility-scale PV system resulting in ramps on the order of 10% s^{-1} and 10% min^{-1}, respectively. Thus a constant maximum ramp assumption will result in large and unacceptable errors. The advantage of our method is that as long as basic historical data for a given site is available, we can provide an estimate of the possible maximum ramp rate.

On the other hand, if a simpler threshold was chosen, the energy storage system would be linearly over/undersized according to the ratio of the chosen maximum ramp rate to the optimal maximum ramp rate. With our method the energy storage system can be sized appropriately.

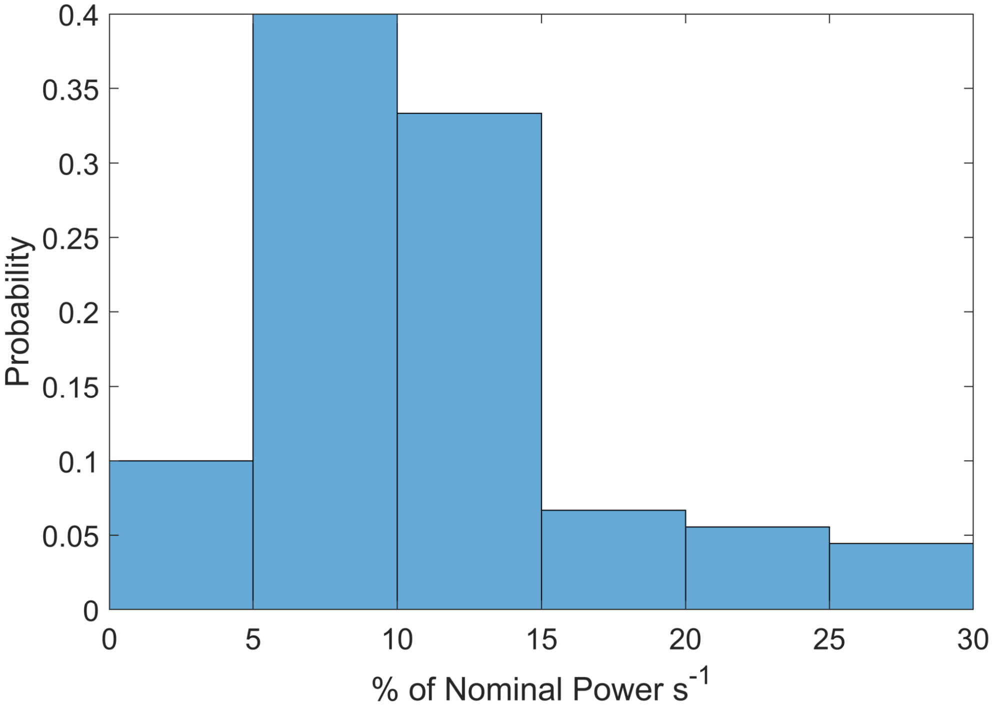

Figure 4.11: Distribution of daily maximum 1 sec ramp rates in percentage of nominal power per second over 90 days

4.4.5 Tradeoff between Noncompliance and Degree of Overestimation

The tradeoff between noncompliance rate and degree of overestimation deserves further discussion. For example, an energy storage system that can mitigate larger ramp rates is more costly, but a smaller energy storage system may not mitigate all ramps and result in noncompliance penalties and/or curtailed solar energy. Ultimately, economic modeling specific to a project is needed to answer these questions. In such conditions, a safety factor could be introduced in Equation (4.5) to determine appropriate kt_i and kt_o and accordingly adjust WCS-RR to fit the risk profile of the investor. Last, the possible ways to reduce the amount of prediction overhead include 1) PV performance models will be needed for more accurate power kt measurements. 2) The power mismatch needs to be considered. For example, the wiring diagram of the PV system can be integrated in the model to count for the extra power loss depending on in which direction the cloud shadow is intercepting the solar arrays.

4.4.6 Applications with Low-frequency Solar Irradiance Data

The proposed method is universally applicable even in the absence of local cloud speed or *kt* measurements, as since the cloud velocity can be derived from NWP data, such as the North American Mesoscale Forecast System (NAM), which contains surface solar irradiance, cloud cover (or at least relative humidity), and wind vectors at different pressure levels, among other variables (see e.g. [65]). The Solar Integration National Dataset (SIND) from NREL is also an alternative source for solar irradiance data [69]. Finally, as discussed earlier, for higher accuracy, power *kt* can be obtained from PV performance models versus the simpler heuristic method used in this paper.

4.4.7 Benefits for PV Plant Sizing and Energy Storage Sizing and Scheduling

The proposed WCS-RR model reveals that, in addition to the meteorology (cloud motion and optical depth), the PV plant shape and orientation with respect to the prevailing winds at the cloud level can significantly impact the power ramp rate. Our results may inspire PV plant designers to preferentially select plant orientations that result in smaller ramp rates at a given location. For example, if consistent prevailing winds at the cloud level are observed as westward or eastward, then extending the PV plant in the east-west direction reduces the relative ramp rates. With an optimal PV plant orientation, the PV system operator can smooth the power output variability by using a smaller ESS. Note that prevailing winds may change seasonally, resulting in potential trade-offs between peak and off-peak production.

During operation, the WCS-RR would be useful to ESS operation for benefit stacking (e.g. frequency regulation and ramp rate mitigation), as less energy and power capacity of the ESS would have to be reserved for ramp rate compliance. For example, Figure 4.12 shows ESS operation for ramp-rate control during a solar power down-ramp. A steeper down-ramp increases both the maximum power and the energy requirement of the ESS. Therefore if solar forecasts suggested thin clouds, slow clouds, or a favorable cloud movement direction, less of the ESS power and energy would have to be reserved for ramp rate control and more could be utilized to monetize other value streams.

Similarly, during ESS planning, if the WCS-RR are smaller at a site given historical weather data, the ESS size could be reduced without risking a ramp rate violation.

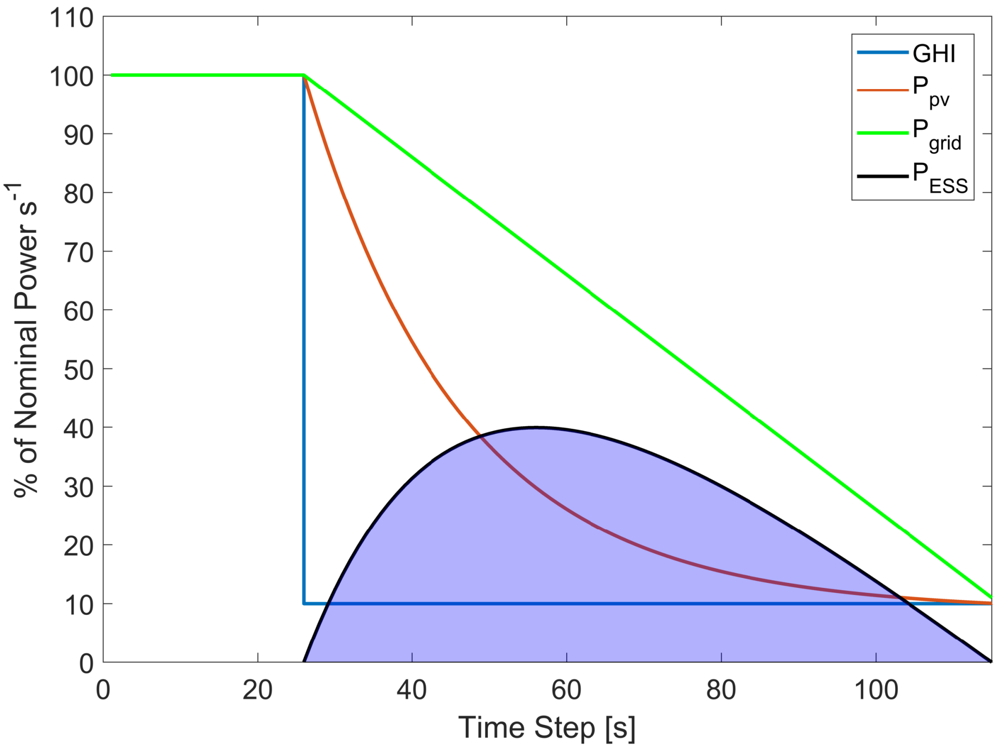

Figure 4.12: Energy storage system discharge for controlling down-ramps in solar power. Blue: Solar irradiance production. Red: Solar power production. Green: Allowable ramp rate from the grid operator. Black: Power shortage between grid operator requirement and solar power production (green minus red). Shaded area: Minimal state-of-charge requirement. *Reproduced from* [51].

.

4.5 Conclusions

Knowing the maximum expected photovoltaic (PV) production ramp rate proves useful for the design of PV and energy storage system (ESS) by determining the ESS energy reserve required to offset power fluctuations. The main goal of this paper is to provide a method to inform optimal design of a solar power plant with ESS by estimating the worst-case scenario ramp rates (WCS-RR) in the design phase, prior to the PV installation and in the absence of local high-frequency irradiance data. A method to estimate WCS-RR for a PV plant, based on cloud speed and direction, solar irradiance (or power), and geometric PV plant layout is developed and validated. The principal assumption is that the cloud field properties are stationary during the cloud passage over the plant. WCS-RR is validated against a PV system during a 10-month period, using cloud motion vectors (CMVs) measured by a cloud speed sensor (CSS).The largest observed ramp of 29.7% s^{-1} is contained with the worst case estimate of 34.3% s^{-1}. The actual ramp rates comply with the calculated maximum ramp rates 98.9% of the time with 2 min evaluation time window. The remaining 1.1% of times can be primarily ascribed to inaccurate cloud velocity measurements in conditions with sparse and/or thin clouds. The high compliance rate also indirectly demonstrates the accuracy of our recently remodeled CSS. The WCS-RR method will be helpful for both PV plant planning and operation.

4.6 Appendix

4.6.1 Output Sampling Rate Requirement

A physical limit for the validation of the WCS-RR method (the application of the WCS-RR does not require high sampling frequency data) is that its sampling time needs to be such that the cloud velocity times the sampling time step does not exceed the dimension of the PV system in the direction of cloud motion. Geometrically, Equation (4.1) in the paper is only valid if the following constraint is met:

$$\Delta t \leq \min\left\{\frac{\mathrm{W}}{|\cos\alpha|v}, \frac{\mathrm{L}}{|\sin\alpha|v}\right\}.$$

For any given sampling rate and dimension of the PV system, there is a cloud velocity for which that sampling rate starts to become too slow, and the WCS-RR then incorrectly predicts the ramp rate over the sampling time step. Since the cloud velocity has a physical limit, a conservative sampling rate can be derived given the dimensions of the PV system. Large PV systems can tolerate larger sampling rates. For example, assume a typical PV panel with a power rating of 208 watts measures 1.6 m x 1 m for 1.6 m^2 of area. A 5 MW PV system with 24,000 PV panels and a typical ground coverage ratio of 0.6 would then cover a square measuring 320 m on each side (including 200 m spacing). With α ranging from 0 to 90 degrees, v = 25 m/s, W = L = 320 m, 13 s to 18 s (or shorter) sampling time step will be needed depending on the specific cloud direction. In summary, a sampling time step of 13 s will guarantee that our model is geometrically valid for a 5 MW PV system.

4.6.2 Day-by-day Performance Comparison Table

Table 4.2: Daily ramping duration (when more than 1% s^{-1} power variation is present), absolute value of daily observed maximum ramp rate in kW s^{-1}, success of bracketing that maximum ramp rate, noncompliance rate ϵ and associated degree of overestimation δ for 2 min, 10 min, and 30 min evaluation time windows for 90 days. The average in the last row is computed over all data points. Small values of ϵ and δ indicate better performance.

Date	Cloudy Period [min]	Maximum RR [kW/s]	Success (S)/Failure (F)	$\epsilon_{2\text{min}}$ [%]	$\delta_{2\text{min}}$ [%]	$\epsilon_{10\text{min}}$ [%]	$\delta_{10\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]
10/8/17	275	3.7	S	0.0	64.4	0.0	57.4	0.0	52.2
10/10/17	140	4.2	S	11.3	37.5	29.4	19.1	62.5	9.4
10/11/17	455	4.4	S	0.5	58.8	2.1	46.2	5.9	35.8
10/13/17	115	1.4	S	4.5	47.7	13.3	39.0	28.6	40.3
11/26/17	125	2.4	S	3.7	56.9	13.3	48.6	28.6	44.3
12/3/17	230	2.1	S	0.0	61.3	0.0	52.4	0.0	42.4
12/16/17	180	3.1	S	0.0	66.2	0.0	61.4	0.0	54.3
1/9/18	170	9.0	S	1.4	68.8	5.0	60.8	14.3	49.6
1/10/18	55	1.8	F(0.7)	11.1	68.9	22.2	67.6	20.0	62.7
1/12/18	35	2.2	S	0.0	49.8	0.0	43.0	0.0	37.1
1/16/18	230	2.3	S	1.2	74.7	3.7	67.9	9.1	56.9
1/18/18	20	2.8	S	11.1	55.1	25.0	46.3	33.3	50.7
1/20/18	155	4.3	S	0.0	69.9	0.0	64.8	0.0	59.9
2/5/18	90	2.8	S	0.0	65.3	0.0	56.8	0.0	47.7
2/10/18	145	3.3	S	0.0	68.4	0.0	64.7	0.0	53.8
2/11/18	85	3.1	S	0.0	68.5	0.0	61.3	0.0	56.7
2/13/18	210	5.6	S	5.7	65.9	12.5	58.8	27.3	54.7
2/14/18	100	2.9	S	0.0	76.7	0.0	69.2	0.0	64.2
2/22/18	60	8.8	S	0.0	67.1	0.0	60.9	0.0	50.6
2/23/18	200	6.9	S	0.0	67.1	0.0	61.1	0.0	51.2

Continued on next page.

Table 4.2 – Day-by-day performance comparison table, continued from previous page.

Date	Cloudy Period [min]	Maximum RR [kW/s]	Success (S)/Failure (F)	$\epsilon_{2\text{min}}$ [%]	$\delta_{2\text{min}}$ [%]	$\epsilon_{10\text{min}}$ [%]	$\delta_{10\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]
2/27/18	185	9.2	S	0.0	71.7	0.0	64.8	0.0	60.7
3/3/18	200	3.2	S	0.0	81.0	0.0	76.1	0.0	71.7
3/4/18	130	2.6	S	2.6	61.0	5.0	56.6	9.1	52.7
3/8/18	275	3.2	S	0.0	75.8	0.0	72.6	0.0	65.1
3/9/18	200	2.6	S	0.0	75.0	0.0	70.7	0.0	63.0
3/11/18	295	6.4	S	0.0	69.6	0.0	61.2	0.0	57.3
3/12/18	190	1.4	S	0.0	77.4	0.0	75.2	0.0	72.1
3/13/18	60	1.1	S	0.0	74.1	0.0	71.9	0.0	72.3
3/14/18	315	3.4	S	0.0	76.3	0.0	70.0	0.0	65.4
3/17/18	295	6.3	S	0.0	73.3	0.0	65.4	0.0	56.7
3/18/18	295	2.2	S	0.0	76.2	0.0	71.2	0.0	65.2
3/20/18	340	5.3	S	0.0	78.4	0.0	72.7	0.0	65.6
3/23/18	210	3.2	S	0.0	70.0	0.0	65.7	0.0	59.1
3/24/18	265	6.6	S	0.0	67.1	0.0	59.7	0.0	51.0
3/25/18	120	3.2	S	2.3	65.8	6.7	63.6	14.3	49.8
3/30/18	180	2.9	F(0.2)	5.2	58.9	15.8	50.5	37.5	37.5
3/31/18	195	0.9	S	0.0	74.6	0.0	72.7	0.0	68.9
4/1/18	170	2.2	S	0.0	66.4	0.0	59.1	0.0	51.3
4/3/18	250	3.4	S	0.0	57.7	0.0	48.3	0.0	45.2
4/12/18	60	3.3	F(0.1)	5.6	74.4	11.1	67.2	16.7	59.0
4/19/18	345	8.5	S	5.2	64.2	7.0	56.2	11.1	49.6
4/24/18	125	2.8	S	2.6	55.3	5.3	51.5	7.1	49.3
5/8/18	105	1.7	N/A	0.0	67.8	0.0	61.9	0.0	58.6
5/9/18	150	2.5	S	5.3	53.0	15.8	41.8	28.6	28.9

Continued on next page.

Table 4.2 – Day-by-day performance comparison table, continued from previous page.

Date	Cloudy Period [min]	Maximum RR [kW/s]	Success (S)/Failure (F)	$\epsilon_{2\text{min}}$ [%]	$\delta_{2\text{min}}$ [%]	$\epsilon_{10\text{min}}$ [%]	$\delta_{10\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]
5/10/18	100	2.4	S	0.0	70.7	0.0	63.7	0.0	58.8
5/15/18	130	2.7	S	2.6	51.0	5.3	42.1	8.3	40.8
5/17/18	395	5.3	S	0.0	60.2	0.0	47.9	0.0	40.5
5/22/18	225	2.8	S	0.0	71.5	0.0	66.7	0.0	59.7
5/23/18	200	4.0	S	0.0	56.8	0.0	42.4	0.0	33.1
5/25/18	320	3.6	S	0.0	63.6	0.0	53.1	0.0	46.7
5/26/18	335	5.8	S	0.8	65.4	2.2	57.2	5.0	49.5
5/27/18	160	2.2	S	0.0	66.0	0.0	58.8	0.0	54.6
5/28/18	165	4.6	S	0.0	59.8	0.0	48.5	0.0	34.2
5/30/18	215	4.4	S	1.3	62.8	3.7	55.5	7.7	44.4
5/31/18	225	3.9	S	3.5	56.8	10.7	48.0	25.0	47.9
6/2/18	95	1.4	S	0.0	52.4	0.0	43.6	0.0	35.0
6/3/18	260	3.3	S	2.9	58.2	6.5	48.8	14.3	41.9
6/5/18	245	4.0	S	0.0	58.9	0.0	48.1	0.0	42.7
6/10/18	95	3.0	S	0.0	57.9	0.0	44.5	0.0	30.4
6/21/18	245	3.0	S	0.0	62.0	0.0	53.0	0.0	45.2
6/27/18	85	2.8	S	5.4	60.0	10.0	48.9	20.0	44.1
6/28/18	140	3.1	S	0.0	60.6	0.0	48.8	0.0	34.7
6/29/18	380	4.6	S	0.0	58.6	0.0	47.0	0.0	41.0
6/30/18	320	3.8	S	1.6	59.7	5.1	51.3	5.9	48.2
7/4/18	70	2.9	S	0.0	55.0	0.0	37.3	0.0	24.9
7/13/18	150	2.1	S	1.7	59.2	5.3	53.6	11.1	49.5
7/14/18	200	3.5	S	6.2	60.5	12.0	48.3	18.2	45.0
7/15/18	120	2.5	S	0.0	52.8	0.0	39.5	0.0	27.6

Continued on next page.

Table 4.2 – Day-by-day performance comparison table, continued from previous page.

Date	Cloudy Period [min]	Maximum RR [kW/s]	Success (S)/Failure (F)	$\epsilon_{2\text{min}}$ [%]	$\delta_{2\text{min}}$ [%]	$\epsilon_{10\text{min}}$ [%]	$\delta_{10\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]	$\epsilon_{30\text{min}}$ [%]
7/17/18	95	1.1	S	0.0	79.5	0.0	77.5	0.0	74.1
7/18/18	265	2.9	S	0.0	74.0	0.0	68.1	0.0	60.8
7/20/18	365	3.6	S	0.0	61.0	0.0	52.9	0.0	46.8
7/21/18	465	6.1	F(0.8)	0.5	61.1	1.9	54.2	5.0	46.4
7/22/18	115	3.5	S	0.0	60.8	0.0	52.9	0.0	41.1
7/26/18	165	2.4	S	1.7	55.5	5.0	42.0	10.0	33.1
7/27/18	75	2.0	S	0.0	67.4	0.0	63.5	0.0	61.8
7/28/18	130	2.7	S	0.0	56.9	0.0	44.5	0.0	38.7
7/29/18	65	1.2	S	5.6	55.7	9.1	53.6	12.5	52.0
7/31/18	40	1.6	F(0.2)	10.0	46.2	14.3	34.9	16.7	40.2
8/12/18	80	1.7	S	5.3	62.7	10.0	54.1	16.7	48.6
8/13/18	190	6.6	F(0.3)	2.5	53.5	10.0	41.1	28.6	31.6
8/14/18	390	5.8	S	1.9	60.6	7.1	47.9	13.3	36.2
8/20/18	145	3.1	S	0.0	74.9	0.0	69.6	0.0	64.5
8/21/18	35	1.0	S	0.0	62.1	0.0	59.4	0.0	59.1
8/23/18	290	4.1	S	0.0	58.7	0.0	45.6	0.0	36.5
8/24/18	185	2.2	S	0.0	70.9	0.0	63.1	0.0	53.3
8/25/18	140	2.6	S	7.3	50.8	16.7	39.6	27.3	40.8
8/26/18	105	2.6	S	0.0	52.7	0.0	46.5	0.0	40.0
8/27/18	245	1.5	S	0.0	75.0	0.0	65.1	0.0	54.8
8/28/18	145	2.2	S	0.0	56.6	0.0	56.0	0.0	48.5
8/29/18	185	3.1	S	0.0	69.0	0.0	63.0	0.0	58.0
Average	N/A	N/A	83S/90	1.1	64.3	2.9	56.8	5.9	50.4

Acknowledgement

This Chapter, in full, is a reprint of the material as it is expected to appear (accepted pending major revisions) in “Wang, G. C., Kurtz, B., Bosch, J. L., de la Parra I. & Kleissl, J. (2020): Maximum expected ramp rates using cloud speed sensor measurements. *Journal of Renewable and Sustainable Energy*.” The dissertation author was the primary investigator and author of this paper.

Chapter 5

Corrective Receding Horizon EV Charge Scheduling Using Short-Term Solar Forecasting

5.1 Introduction

5.1.1 Problem Statement and Literature Review

The variable nature of solar power is of concern to electric grid operators, where there is substantial growth in photovoltaic (PV) installations. Variable power flow occurs primarily due to passing clouds. As a result, PV generation exhibits high variability, leading to power quality issues such as flicker, power imbalance, reverse power flow, and increased wear on conventional voltage regulator equipment at the distribution level [56, 70]. With high PV penetration, these effects aggravate and can cause challenges to grid operations.

Accurate solar forecasts help grid operators integrate increased levels of solar generation while maintaining power quality. However, forecast errors are inherent to any forecasting technique, and non-linear atmospheric dynamics make it challenging to reduce or correct solar forecast errors. Forecast errors can be detrimental to reduce peak loads using energy storage [71–73], or maximizing PV self-consumption [74]. While the characteristics of forecast errors vary with time scale and location [75, 76], grid integration studies typically model forecast errors by two main approaches:

1) Generate an imperfect forecast by adding a synthetic error to the actual solar generation. The synthetic error is often sampled from univariate distributions such as Gaussian, uniform, and Weibull [77–82]. Such error timeseries have zero autocorrelation, failing to capture the autocorrelation properties of real forecast errors. Days with large forecast error autocorrelation, especially persistent over-forecasts for several hours, usually present the most challenging conditions for energy storage to reduce peak load [16]. On such days the energy storage system needs to discharge continuously to make up for the shortfall in PV generation which may cause premature discharge, demand peaks, and associated demand charges.

2) Probabilistic modeling of errors: sophisticated statistical methods such as copula functions [83], Markov model [84], enhanced Markov chain model [85] kernel density estimation [86] are utilized to model the time series forecast. However, these models do not fully capture the autocorrelation of forecast errors due to nonlinear and higher order dependencies. For example, the transformation-based heuristic methodology in [87] captures the spatio-temporal correlation properties of forecast errors on the day-ahead time scale, but not intra-day. Other forecast error modeling considering autocorrelation include [16] where a parametric first-order autoregressive process is developed to generate autocorrelated time series forecasts and [88] where a simulated autocorrelated PV forecast error through a transformed multivariate ARMA model is presented.

5.1.2 Objectives and Contribution

While modeling forecast errors to support grid planning and operation studies has received much attention, corrections of the impact of forecast errors in real time are typically not applied. Motivated by this, we propose a corrective approach with the main contributions listed as follows:

1) We avoid statistical models of forecast errors by using a day-ahead persistence and 15 min-ahead sky imager forecast to produce real forecast data which inherently preserves the temporal dependence structure of forecast errors.

2) We apply a corrective optimization framework. The baseline day-ahead persistence forecast is corrected by three short-term forecasts showcasing different levels of realistic forecast errors: day-ahead perfect forecasts, day-ahead persistence forecasts corrected by imperfect short-term sky imagery forecasts, and

day-ahead persistence forecasts corrected by perfect short-term forecasts. This approach captures the real forecast property of improved accuracy with shorter forecast horizon.

The proposed case study to implement and validate the above framework uses EVs as mobile energy storage systems (ESS). Adopting ESS to compensate the mismatch between variable PV output and grid load [50, 89] has been extensively studied in the past decade. For example, Nottrott et al. [90] modeled ESS dispatch schedules for peak net load minimization by linear programming. Mixed-integer linear programming and quadratic programming (QP) are also commonly utilized to solve the ESS scheduling problem at the distribution level [91–94]. The adoption of EV as an alternative to ESS has received significant attention [95] because of low acquisition cost. The interaction between EVs and the power grid is comprehensively discussed in [12]. High EV penetration can provide grid services [96] such as valley filling defined as increasing load demand during the load depression [97] using unidirectional EV charging management (V1G) [98] or vehicle-to-grid (V2G) [99] schemes. We formulate a typical valley filling problem through centralized [100] EV charge scheduling in a realistic, quasi-operational case study. Since V2G still faces challenges such as market barriers and limited commercial availability we focus on V1G here. The impact of realistic forecast errors is quantified as the deviation of the resulting valley filling solution from the ideal solution. Last, we illustrate how correcting day-ahead forecasts in real time with more accurate short-term forecasts benefits the valley filling solution.

5.1.3 Paper Organization

The remainder of this paper is organized as follows. In Section 5.2 we introduce the methodology for determining load demand, solar forecasts, and EV charge events. Section 5.3.1 defines the problem statement, and sections 5.3.2 and 5.3.3 introduce a QP optimization algorithm based on [94]. We extend the work by supplementing constraints for start and end time of EV charging and energy demand satisfaction, and integrating a receding horizon framework. Section 5.3.4 introduces error metrics. Section 5.4.1-5.4.3 provide a validation on a sample day, and statistical results from one month of valley filling results, and the discussion of the results and limitations of the proposed methodology are given in Section 5.4.4-5.4.5. Lastly, Section 5.5 provides conclusions and future work.

5.2 Problem Setup and System Data

5.2.1 Geographic Setup

We select the region of San Bernardino located in Southern California, where Southern California Edison (SCE) installed 125 MW of multiple rooftop photovoltaic (PV) systems under the Solar Photovoltaic Program (SPVP). This region is home to many commercial buildings, large warehouses, and abundant solar resources, which makes large PV rooftop arrays common. Local distribution feeders are therefore prone to solar variability issues, making this region an ideal location for solar power integration studies.

Figure 5.1 shows an overview of the area served by the 66/12 kV substation in San Bernardino, CA. The substation load demand was simulated by EnergyPlus, a building energy simulation tool developed by the U.S. Department of Energy. Power output from two SPVP systems were provided by SCE. To produce short-term solar forecasts, we use a UC San Diego Sky Imager system (USI). The USI is rooftop-mounted and its field-of-view covers all buildings within the substation service territory. It can be used to geolocate clouds, measure cloud velocity, and track cloud motion [35, 39]. These measurements are then used to predict future cloud shadow locations and solar irradiance up to 15 min ahead. For more detailed information of solar resource assessment and forecasting using the USI, refer to [2, 37].

5.2.2 Solar Forecasts

To demonstrate the net load flattening improvement through correcting day-ahead persistence forecast by short-term forecast updates, we construct four different solar generation forecasts G in 24-hour time windows representing different forecast accuracy through a combination of persistence G_{p}, USI forecast G_{USI}, and perfect forecast G_{perfect}.

1) Base Forecast (G_{p}): As conventionally adopted [71, 101] as a baseline for load forecasting, a 24-hour persistence forecast is defined as solar power at the same time of the previous day. Persistence forecasts are expected to have the largest forecast error and largest autocorrelation of forecast errors.

2) Operational forecasts ($G_{\mathrm{p+USI}}$): The base forecast is continually updated with the most recent USI forecast. Since USI forecast horizons are limited to 15 min, only the G_{p} in the first 15 min of the 24-hour time horizon is replaced with G_{USI} while the G_{p} in the remaining time horizons are left

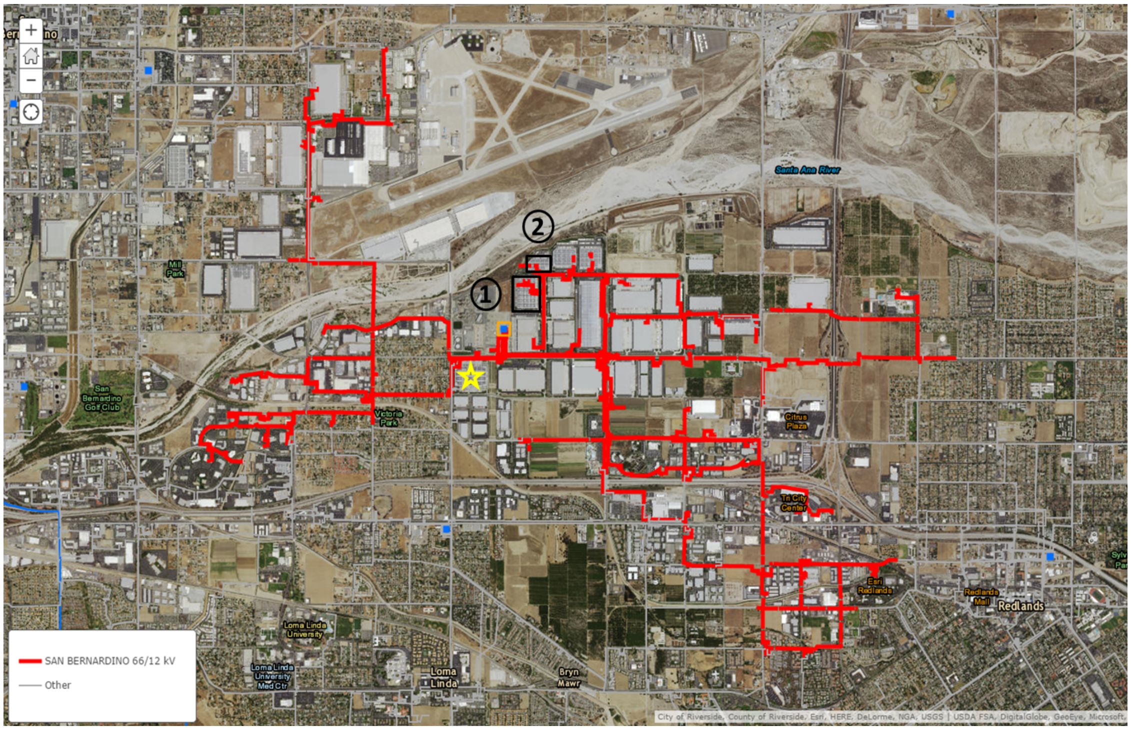

Figure 5.1: Southern California Edison's (SCE) Distributed Energy Resource Interconnection Map (DERiM) showing the San Bernardino substation service territory. The substation is marked by a blue box with orange border and feeder lines emanating from that substation are in red. Rooftop PV systems (black boxes) and sky imager (yellow star) are located up to 1,100 m apart. The map spans 12 x 6.5 km.

unchanged. As G_{USI} is more accurate than G_{p} (refer to Table 5.4 in the Appendix), $G_{\text{p+USI}}$ is expected to exhibit smaller forecast errors than G_{p}.

3) Benchmark forecast ($G_{\text{p+perfect}}$): Similar to the operational forecast, but the first 15 min of G_{p} is replaced with G_{perfect}. $G_{\text{p+perfect}}$ is expected to exhibit even smaller forecast errors than the operational forecast. Thus it elucidates whether further improvements in the accuracy of short-term forecast would result in better mitigating the impact of forecast errors on the net load.

4) Perfect Forecast (G_{perfect}): The entire 24-hour time horizon of G_{p} is replaced with G_{perfect}, yielding zero forecast error. The perfect forecast brackets the net load flattening that is achievable. If EV availability was unconstrained, G_{perfect} would yield a flat net load curve.

Table 5.1: Catalog of the datasets and their native temporal resolution. All data are interpolated to 15 min temporal resolution. The rated EV charge capacity is determined by the sum of the charge rates in Table 5.3.

Data Source	**Rated Capacity**	**Scale Factor**	**Average Daily Power**	**Type**	**Original Time Resolution**
Loads	5.4 MW	150	23 kW	Simulated	60 min
PV Generation	7.5 MW	75	13 kW	Measured	30 s to 2 min
USI PV Forecast	7.5 MW	75	13 kW	Measured Sky Images	30 sec
EV Charge Needs	0.23 MW	1	19 kW	Simulated	5 min

5.2.3 EV Fleet

We focus on workplace charging of EVs during regular business hours coinciding with times of peak PV production. A fleet of EVs connected to workplace charging stations is simulated. The EV make and models and their battery capacity and charge rates are selected based on the EV market share in the US as of 2015 (Table 5.3).

Arrival time, layover duration, and the initial state of charge were sampled from the following distributions: 1) EV arrival time varies between 06:00 PST and 10:00 PST centered on a mean arrival time of 07:30 PST with a standard deviation of 1 h. 2) layover duration spans from 6 to 11 hours centered on 8 hours with a standard deviation of 1 h. 3) initial state of charge (χ_0, in %) ranges from 0-100% centered on 60% with a standard deviation of 10%. The energy demand is derived assuming a full charge by the departure time. The resulting EV charging events are replicated for all 30 days of analysis. Since most employee day-to-day schedules are repetitive, the persistence of daily charging events is a reasonable assumption.

5.2.4 Summary of Data Sources and Availability

Table 5.1 provides an overview of the load, PV generation, and EV energy demand datasets. Because PV generation data is complete only for April 2013, the full month is selected for our analysis. During this month, there were 2 overcast days, 13 clear days, and the remaining 15 days were partly cloudy. The simulated substation loads were scaled down by a factor of 150x from 5.4 MW to 36 kW and PV data were also scaled down by 75x from 7.5MW to 100 kW so that the 31 EVs in Table 5.3 are able to fill the entire energy valley on a clear day.

5.3 Methodology

5.3.1 Problem Setup

To test our proposed approach, we formulate a valley filling optimization problem to schedule EV charging. Figure 5.2 schematically illustrates the problem configuration. The load demand estimate L for all buildings connected to the distribution feeder is offset by PV systems injecting power G. The substation is assumed to allow bidirectional power flow (i.e. net load NL, positive if delivering power and negative for reverse power flow). N EVs draw charge power $\beta^{[N]}$ from the distribution grid. The arrowheads indicate positive power flow. As V2G (i.e. Vehicle to Grid charging) still faces challenges such as market barriers and limited commercial availability, we focus on V1G (unidirectional EV charging) in this study. However, EV discharging functionality is supported in the optimization framework.

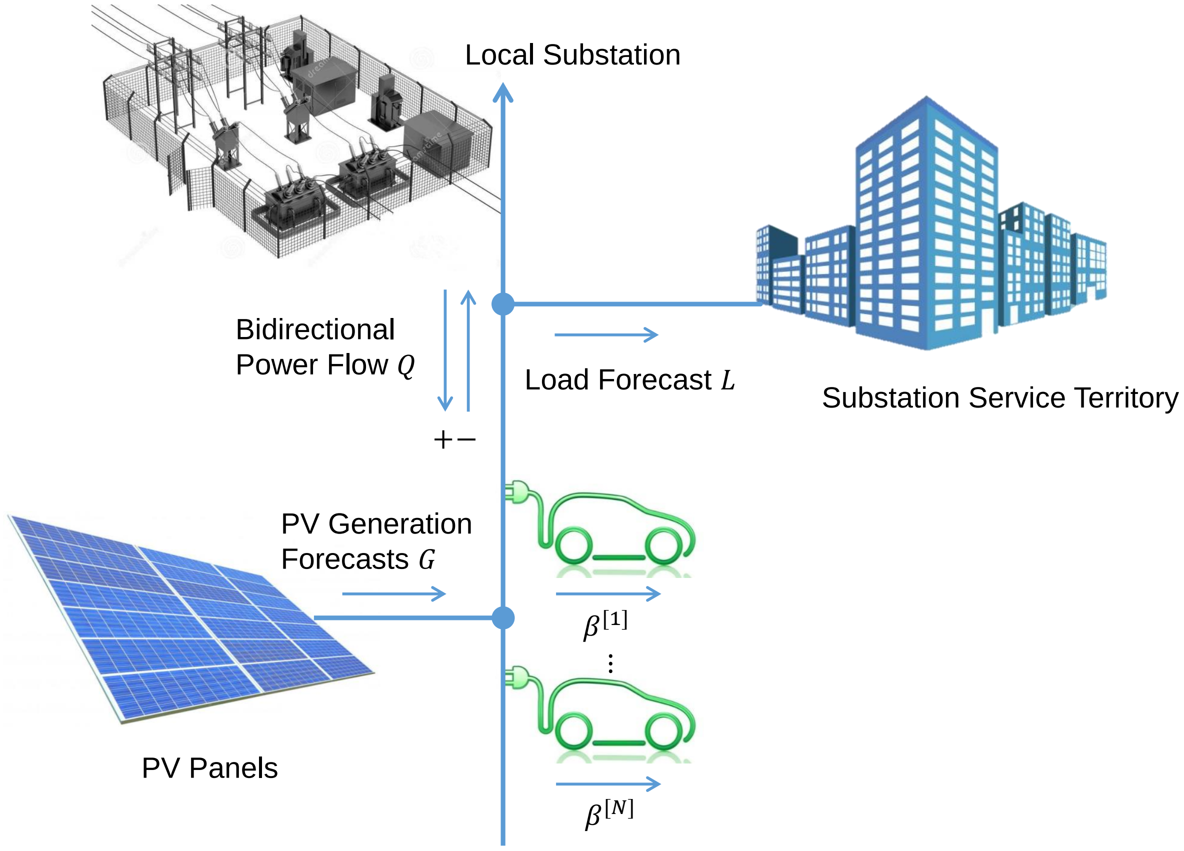

Figure 5.2: Notation and topology of the proposed optimization problem.

From Figure 5.2, the following power balance between the four power variables always holds true:

$$\mathrm{NL} = L + \boldsymbol{\beta} - G, \tag{5.1}$$

where EV charge power $\boldsymbol{\beta} = \left[\beta^{[1]} \ldots \beta^{[N]}\right]$. As we are interested in a day-ahead optimal valley filling solution, the total length of the planning horizon is 24 hours as in $s = 96$ steps at a temporal resolution Δ of 15 min (0.25 h), consistent with the USI forecast horizon. The time index $T_k = k\Delta$. The power variables in Eqn. (5.1) at each time index represent the average power output over the period from a time index ahead to present, expressed as $((k-1)\Delta,\ k\Delta)$, where $k \in [1 \ldots s]$. For example, the power balance at the first time index T_1 considers the average power output between midnight 0Δ = 00:00 h and 1Δ = 00:15 h. Refer to Figure 5.4 for detailed illustration of planning horizon, temporal resolution and time step.

Figure 5.3 illustrates the schematic overview of the proposed study. The study solves the valley filling optimization problem with individual input of four solar forecast scenarios (Section 5.2.2) G_{p}, $G_{\mathrm{p+USI}}$, $G_{\mathrm{p+perfect}}$, and G_{perfect}, and compares the resulting bidirectional net load NL_{p}, $\mathrm{NL}_{\mathrm{p+USI}}$, $\mathrm{NL}_{\mathrm{p+perfect}}$, and $\mathrm{NL}_{\mathrm{perfect}}$ respectively.

5.3.2 Mathematical Formulation of the Optimization Algorithm

Battery scheduling to reduce peak load and minimize energy bills has been implemented through a variety of optimization algorithms [102]. Considering potential future applications of implementing financial incentives and prioritizing EVs, we prefer a framework which naturally supports weighting different objectives. Therefore, QP is selected owing to its flexibility:

$$\min_{x} h(x - f)^2, \tag{5.2}$$

where x is a vector of grid net load NL and EV charge power β of each EV (see Eqn. 5.6), and h is the corresponding matrix of weighting factors. While the weighting factor is not activated in this paper, it enables future implementation of economic objectives that allow trade-off between the decision variables (e.g., price, charge power). Thus, the weighting matrix h is set to a $(1+N)s \times (1+N)s$ matrix with an $s \times s$ identity matrix $\boldsymbol{I}_s$ included in the top left corner for grid and zero elsewhere for EV.

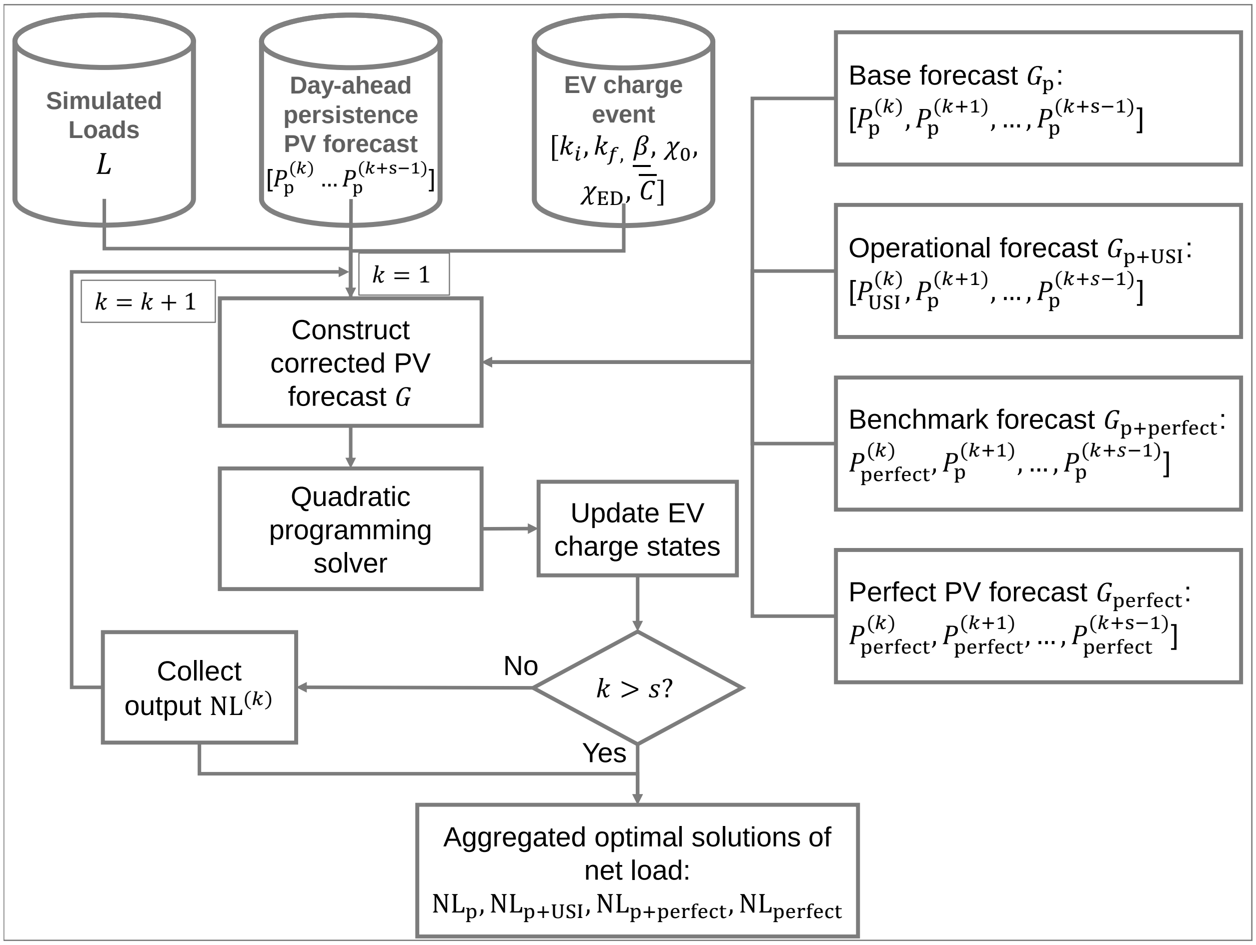

Figure 5.3: Flowchart of the proposed corrective approach. Refer to Table 5.3 for the EV charge events.

$$h = \begin{bmatrix} \boldsymbol{I}_s & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 0 \end{bmatrix} \in \mathbb{R}^{[(1+N)s \times (1+N)s]}. \tag{5.3}$$

In Eqn. (5.2), f is a vector composed of the preferred grid net load profile followed by preferred EV charge power. It defines the power target for the objective function. Ideally, the optimized grid net load profile should be a flat line in case of perfect forecasts (i.e. load forecast and solar forecast) and unrestricted EV availability. However, as discussed earlier, real solar forecasts have errors, and also EV availability can be restricted depending on EV owners work schedule. The resulting mismatch between the EV energy demand and the energy valley magnitude forces the optimized net load to deviate from a flat line (see Section 5.4). The preferred grid net load f should therefore be dynamically updated to align the future charge schedule with the remaining PV generation and EV charging needs, and is defined as

$$f_{(k)} = \begin{bmatrix} \boldsymbol{r} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \in \mathbb{R}^{(1+N)s}, \tag{5.4}$$

where $\boldsymbol{r} = [r \dots r]^T \in \mathbb{R}^s$ is the grid net load target with identical scalar element r, which is determined from the expected net load $(L - G)$ and the cumulative EV energy demand χ_{ED} for the time period from the current time step (t_k) to when the last EV disconnects (t_{end}):

$$r = \frac{\sum_{t_k}^{t_{\text{end}}} L - G + \chi_{\text{ED}}}{t_{\text{end}} - t_k}. \tag{5.5}$$

To proceed, x in Eqn. (5.2) is the decision variable consisting of temporal aggregated grid power load and all EV charge power.

$$x = \begin{bmatrix} \text{NL} \\ \beta^{[1]} \\ \vdots \\ \beta^{[N]} \end{bmatrix} \in \mathbb{R}^{(1+N)s}. \tag{5.6}$$

$\text{NL} = [\text{NL}_{(1)} \dots \text{NL}_{(k)}]^{\text{T}}$ and EV charge power follows $\beta^{[N]} = \left[\beta_{(1)}^{[N]} \dots \beta_{(k)}^{[N]}\right]^{\text{T}}$ where $k \in [1 \dots s]$, N is the number of EVs, and parenthesis () indicate the time index.

The charge power β of each EV in Eqn. (5.6) is subject to the following five constraints (i.e. Eqns. 5.7-5.11). First, the EV charge power is limited to its maximum charge capacity. Because discharging is not considered, the EV charge rate constraint for all time steps k is

$$0 \leq \beta_{(k)} \leq \overline{\beta}. \quad k \in [1 \dots s] \tag{5.7}$$

Second, the state of charge (SOC) of the battery is constrained by:

$$C \leq \chi_0 \cdot \overline{C} + \sum_{k=1}^{s} \beta_{(k)} \cdot \Delta \leq \overline{C}, \quad k \in [1 \dots s] \tag{5.8}$$

where χ_0 is the initial SOC in %, $\overline{C}$ is maximum SOC conventionally defined as the EV battery capacity and the minimum SOC C is assumed to be 0. The EV battery capacity is assumed to equal the amount specified by the vehicle manufacturer. Battery degradation is assumed to be negligible.

Third, the charging has to occur within the constraints of the EV layover time. If the EV connects at time index k_i, no charging can occur during the period $[T_1 T_{k_i}]$. In other words, EV charge power equals to 0 prior to connecting at $k = k_i$:

$$\sum_{k=1}^{k_i} \beta_{(k)} \cdot \Delta = 0. \tag{5.9}$$

Fourth, similarly given EV departure time index k_f, the EV battery charge is constrained to 0 after disconnecting at $k = k_f$:

$$\sum_{k=k_f}^{s} \beta_{(k)} \cdot \Delta = 0. \tag{5.10}$$

Last, the power supplied to the EV during the layover period equals the energy requested by the EV owner. We assume that all EV owners request a fully charged battery.

$$\sum_{k=1}^{s} \beta_{(k)} \cdot \Delta = \chi_{\mathrm{ED}}. \tag{5.11}$$

5.3.3 Receding Horizon Optimization Algorithm

To support the proposed corrective approach, a RHC algorithm is required [103–105]. The RHC algorithm modifies the control action with respect to predicted solar energy generation while satisfying constraints over a time window of fixed length 24 h. We use a moving time window resembling the receding horizon mechanism of RHC. The simulations are initialized at midnight and the time window moves forward one step at a time, as presented in Figure 5.4.

The iterations in RHC demand a fast optimization solver and real time application can be limited by the number of constraints and EVs. With 31 EVs, each iteration takes 1.5 s on an Intel I5 workstation, and a full day completes in 2.5 min. The computation cost scales with the number of variables (i.e., the number of EVs).

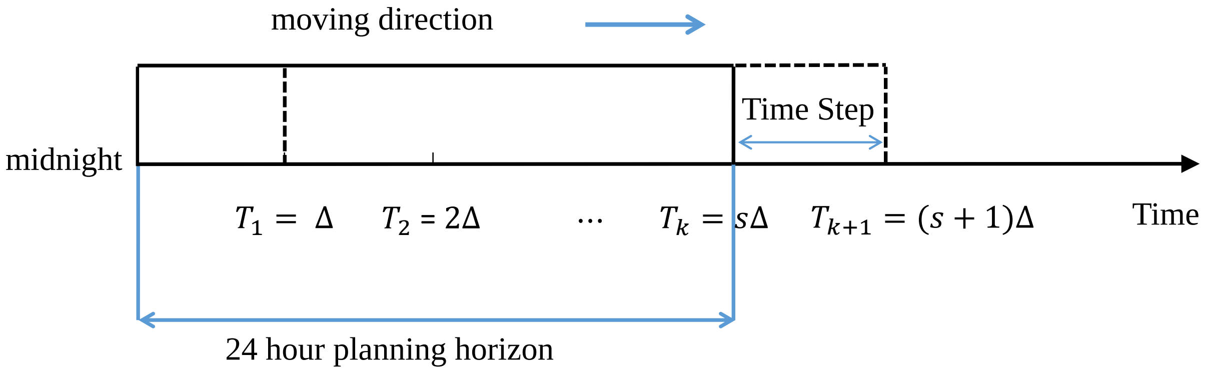

Figure 5.4: Illustration of receding horizon time window. Refer to Section 5.3.1 for definitions of the notations.

5.3.4 Forecast and Valley Filling Benchmark

The valley filling performance is evaluated as follows. The accuracy of the solar forecasts is characterized by nRMSD e^G, which is the root mean square difference normalized by the average solar power measurements:

$$e^G(\text{day}) = \frac{\sqrt{\frac{1}{s}\sum_{k=1}^{s}\left[G_{\text{forecast}}(k) - G_{\text{perfect}}(k)\right]^2}}{\frac{1}{s}\sum_{k=1}^{s}\left[G_{\text{perfect}}(k)\right]} \times 100\%. \tag{5.12}$$

Similarly, the nRMSD for optimized net load e^{NL} is defined as:

$$e^{\text{NL}}(\text{day}) = \frac{\sqrt{\frac{1}{t_{\text{end}}-t_{\text{start}}}\sum_{k=t_{\text{start}}}^{t_{\text{end}}}\left[\text{NL}_{\text{imperfect}}(k) - \text{NL}_{\text{perfect}}(k)\right]^2}}{|\max\left(\text{NL}_{\text{perfect}}(k)\right)|} \times 100\%. \tag{5.13}$$

Thus, the nRMSD of valley filling using imperfect solar forecasts NL_{p}, $\text{NL}_{\text{p+USI}}$, and $\text{NL}_{\text{p+perfect}}$ can be evaluated by e^{NL}_{p}, $e^{\text{NL}}_{\text{p+USI}}$, and $e^{\text{NL}}_{\text{p+perfect}}$ respectively. Note that only the net load at the first time step (15 min) of all time windows is evaluated in Eqn. (5.13) which means that at each index k, the solar forecasts for 15+ min horizon do not influence e^{NL}.

$$\sigma^{\text{NL}}(\text{day}) = \sqrt{\frac{1}{N}\sum_{t_{\text{end}}}^{t_{\text{start}}}\left[\text{NL}(k) - \text{NL}\right]^2}, \tag{5.14}$$

where NL is the average net load from t_{start} to t_{end}, and N is the number of time indices in between. Small σ indicates smoother net load profiles, and $\sigma = 0$ indicates a completely flat net load profile. The net load standard deviations after valley filling using solar forecasts NL_{p}, $\text{NL}_{\text{p+USI}}$, $\text{NL}_{\text{p+perfect}}$, and $\text{NL}_{\text{perfect}}$ are evaluated by $\sigma^{\mathrm{NL}}_{\mathrm{p}}$, $\sigma^{\mathrm{NL}}_{\mathrm{p+USI}}$, $\sigma^{\mathrm{NL}}_{\mathrm{p+perfect}}$, and $\sigma^{\mathrm{NL}}_{\mathrm{perfect}}$, respectively. Eqns. (5.13) and (5.14) are evaluated only during the maximum layover period (i.e. from t_{start} to t_{end}) when the energy valley can be filled.

5.4 Results and Discussion

5.4.1 Case Study

The methodology discussed in Section 3.3 is first applied to a single day with the perfect forecast scenario G_{perfect}. Apr 1, 2013 was a day with broken cumulus clouds, producing high PV output variability. Figure 5.5 shows the valley filling results by smart charging EV 1 and EV 3 defined in Table 5.3. The algorithm dynamically schedules charging during energy valleys while obeying all constraints. Specifically, the grid net load profile on the top (subtracting actual PV power G_{perfect} from the feeder load profile L) shows an energy valley during midday interrupted by episodes of cloud cover. The area between blue and red indicates the share of the energy valley that is absorbed by coordinating charging of 2 EVs.

EV 3 has a small energy demand of 7.7 kWh, but the layover period of 9 hours is comparably long since EV 3 only requires 2 h 20 min of charging at the highest charging rate to reach full charge. Thus, the algorithm can schedule charging to occur only during the four major energy valleys (clear periods during midday) while reducing charging power to zero in cloudy conditions and at the beginning and end of the day. EV 3 is fully charged 2 h prior to its planned departure.

On the contrary, EV 1 has much less charging flexibility with a larger energy demand (76.5 kWh or 7 h 39 min of charging at maximum capacity), which is spread over a similar layover period of 9 hours 40 min. Thus, the algorithm has limited flexibility (2 hours) to shift the battery charge schedule and/or reduce the charge rate. EV 1 connects to charge at 08:00 PST, and charging occurs from 08:15 PST when PV generation starts to depress the net load profile. The 2 h flexibility is used to charge at a rate below the maximum capacity from 08:15 to 13:30 PST. By lowering the charge rate, the algorithm shifts EV charging from times of relatively large net load to times with a larger energy valley later in the day. The algorithm boosts charge power of EV 1 to maximum capacity after 13:30 PST so that EV 1 can reach full charge right at its planned departure. Note that limited EV flexibility prevents the algorithm from completely shifting the EV charging to periods with larger energy valleys and it schedules charging also during off-peak solar generation (08:00 10:45 PST).

For this case, the power reference is unachievable due to the small number of EVs and restricted

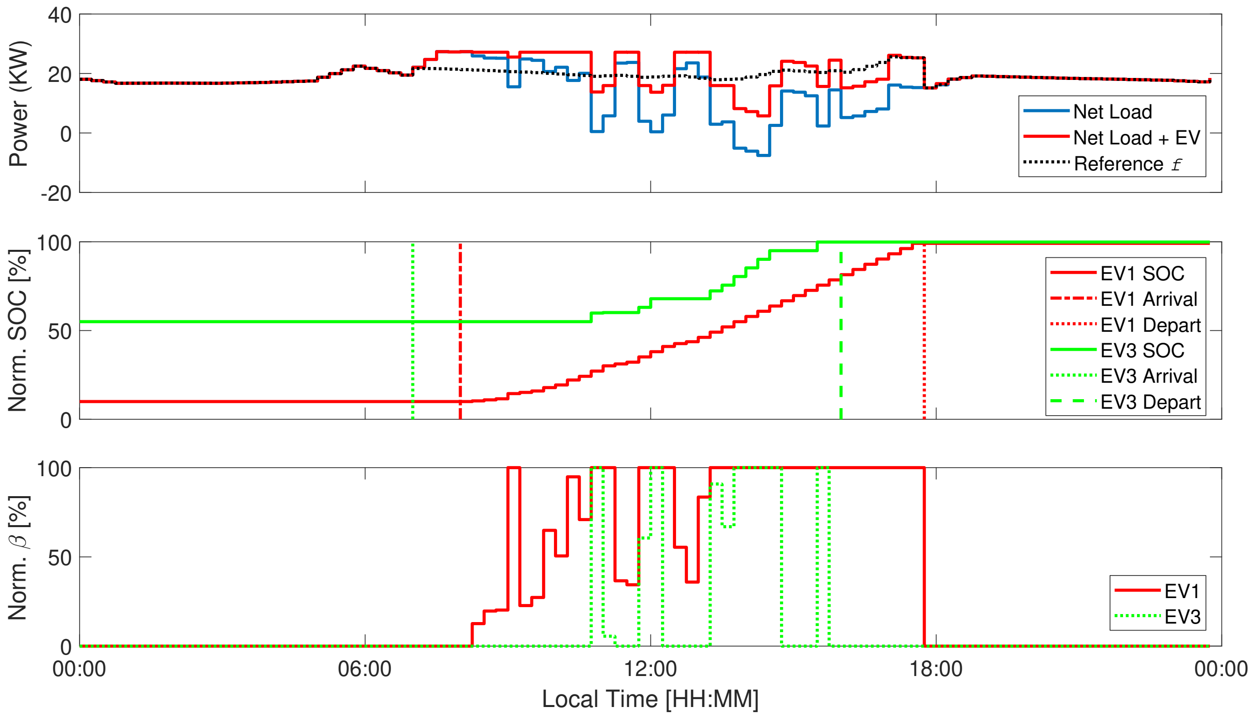

Figure 5.5: Sample optimized EV charge schedule for valley filling with 2 EVs. Top: Original grid net load profile showing an energy valley (blue), filling the valley towards a reference power (f, black) yields optimized net load (red). The PV output profiles are scaled down by a factor of 260x to create a reasonable energy valley for just two EVs. Middle: EV state of charge in %. Colors distinguish EV with their arrival time (dashed) and departure time (dotted). Bottom: EV charge power normalized by its maximum charge capacity.

charge schedule for EV 1. But the algorithms functionality and ability to reduce net load variability by shifting the charge schedule with respect to the energy valley availability is confirmed.

5.4.2 Illustration of Autocorrelated Solar Forecast Errors

The impact of autocorrelated solar forecast errors are thoroughly studied and illustrated in Figure 5.6. First, Fig. 5.6a presents two valley filling solutions by using persistence solar forecasts that naturally have autocorrelated errors and perfect solar forecasts.

With perfect forecasts, the exact energy valley is accurately estimated upfront (dashed orange) so the algorithm is able to fully flatten the grid net load during EV's layover, and the consumed energy is illustrated by the orange shaded area. In comparison, persistence forecast on this day significantly over-estimates the energy valley by producing an "imaginary" energy valley in the mid-day (dashed blue). While the resulting valley filling solution also meets the charge needs indicated by the blue shaded area of the same size to the orange shaded area, it introduces a large net load peak to the grid network. To proceed, we make up a "synthetic" solar forecasts by adding random noises to the perfect forecasts, with the forecast errors in the same magnitude to the forecast errors of the persistence forecasts (denoted by "persistence forecasts with random error"). As shown in Figure 5.6b, with the same amount of valley filling (shaded purple area), the synthetic solar forecasts result in a "noise-looking" but peak-free net load profiles, which lowers the net load variability to the level that using the realistic persistence forecasts would not be able to achieve. Figure 5.6 demonstrates that with a fixed magnitude of forecast errors, the autocorrelation structure can severely worsen the net load power quality by introducing power peaks to the grid network, and the simple statistical modeling of the solar forecast errors can cause unrealistic valley filling results. The quantitative impact of this autocorrelated solar forecast errors is further characterized in Section 5.4.4.

5.4.3 Monthly Results and Statistics of Forecast Errors Correction

Now we consider a full month, 31 EVs (Table 5.3 in the Appendix), and four solar forecast scenarios. For readability, only scenario $\mathrm{NL}_{\mathrm{p+USI}}$ is presented in Figure 5.7. On most days, valleys are completely filled, and optimized net load is closely aligned with the reference f indicating that the method works as designed. On clear days (e.g., Apr 21 and Apr 23), the energy valley is sufficiently large to charge all EVs, yielding a flat net load. On partly cloudy days (e.g., Apr 1 and Apr 4), solar variability is large, and less energy is available in the valley, resulting in increased and fluctuating optimized net load.

Table 5.2 summarizes the monthly performance and a detailed day-by-day performance comparison is presented in Table 5.4. The statistics of PV forecast errors and resulting valley filling performance are further investigated in the following sub-sections.

Table 5.2: Monthly summary of solar forecasts accuracy and optimized net load variability under different forecast error scenarios.Normalized RMSD for solar forcast e^G and optimized net load e^{NL}, and standard deviation of the optimized net load σ^{NL} under scenarios of base forecast, operational forecast, benchmark forecast, and perfect forecast. Superscripts indicate the type of validations and subscripts indicate the solar forecasts scenarios used in the valley filling problem.

	Persistence Forecast	USI Forecast	NL_{p}	$NL_{\mathrm{p+USI}}$	$\mathrm{NL}_{\mathrm{p+perfect}}$	$\mathrm{NL}_{\mathrm{perfect}}$
nRMSD [%]	68.2	37.2	23.1	17.5	14.1	N/A
Std. σ [kW]	N/A	N/A	11.7	8.8	7.3	1.1

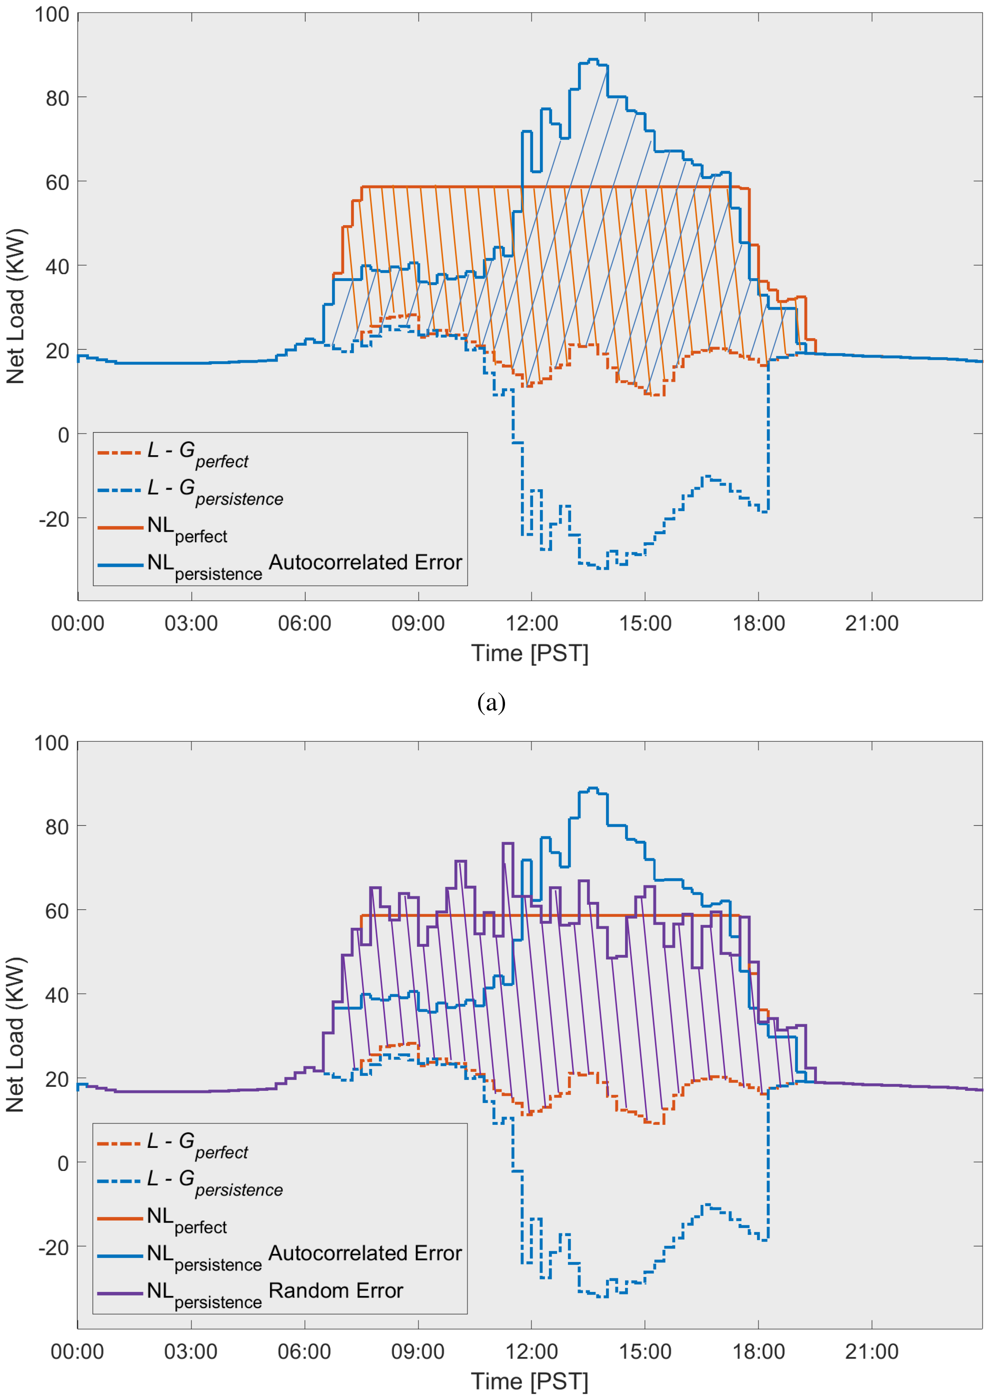

(a)

(b)

Figure 5.6: (a) Valley filling results produced by using perfect forecasts and persistence forecasts exhibiting autocorrelated errors. (b) Valley filling results produced by using perfect forecasts, persistence forecasts exhibiting autocorrelated errors, and persistence forecasts exhibiting random errors generated by adding white noises to the perfect forecasts.

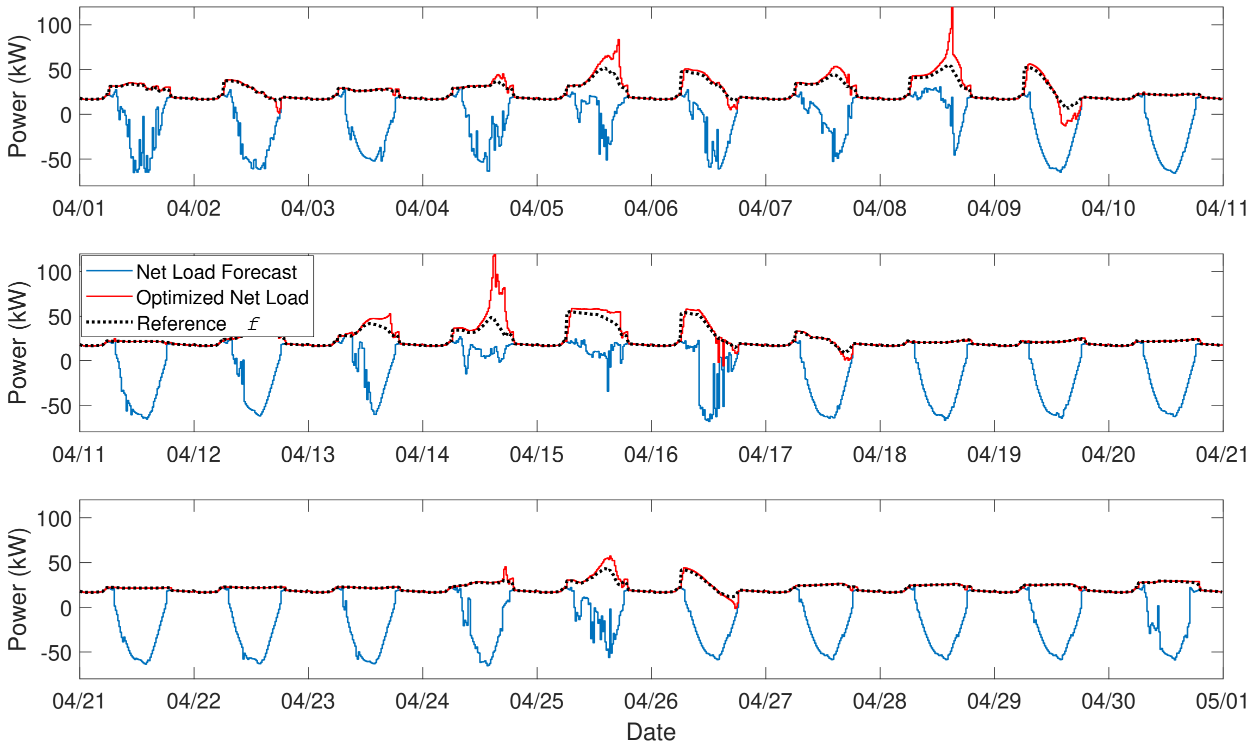

Figure 5.7: Results for one month of valley filling with forecast errors corrected by USI (using $G_{\mathrm{p+USI}}$). The original grid net load without EV (blue) is flattened by scheduling EV charging to achieve a preferred net load profile (black), resulting in reduced net load variability (red).

Persistence and USI PV Forecast Accuracy

Overall, the USI solar forecast outperforms 24-hour persistence on 17 out of 30 days. While the persistence forecast outperforms USI by an average nRMSD of 6.2 percentage points (30.8%) on the remaining 13 days, the USI solar forecast lowers monthly average nRMSD by 31.0 percentage points (45.4%) placing it about halfway between persistence and perfect forecast. In general, the solar forecast results confirm our expectation that correcting persistence forecast by USI forecast reduces forecast error.

Valley Filling with USI Forecast Correction

Overall, $\mathrm{NL}_{\mathrm{p+USI}}$ delivers monthly nRMSD and σ averages of 17.5% and 8.8 kW, which are 24.2% and 24.8% below NL_{p}, respectively. The optimization for NL_{p} and $\mathrm{NL}_{\mathrm{p+USI}}$ forecasts performs similarly for mostly clear days and errors are small; specifically, both NL_{p} and $\mathrm{NL}_{\mathrm{p+USI}}$ yield low error and variability averages less than 5.0 percentage point and 2.0 kW, respectively. On the 15 partly cloudy days $\mathrm{NL}_{\mathrm{p+USI}}$ outperforms NL_{p} by 33.0% and 32.7%. While $\mathrm{NL}_{\mathrm{p+USI}}$ performs worse than NL_{p} on some of the days when sky conditions in present day significantly differ from a day earlier (Apr 5, 8, 9, 14, 16, 25 and 26, see Section 5.4.4), the average error reductions of $\mathrm{NL}_{\mathrm{p+USI}}$ over NL_{p} on these days are still 7.9% and 11.2% for nRMSD and σ, respectively. Note that the large improvement in solar forecast accuracy by the USI over persistence forecast (45.4%) does not translate to an equal improvement in optimized net load variability, which will be discussed in Section 5.4.5.

Valley Filling with Perfect Forecast Correction

The error reductions by $\mathrm{NL}_{\mathrm{p+perfect}}$ over $\mathrm{NL}_{\mathrm{p+USI}}$ are relatively smaller (19.4% and 17.0% less with respect to nRMSD and σ). The fact that $\mathrm{NL}_{\mathrm{p+USI}}$ is closer to $\mathrm{NL}_{\mathrm{perfect}}$ than NL_{p} demonstrates the feasibility of operational forecast deployment of USI forecasts.

Valley filling with perfect 24-hour forecasts performs the best as expected with monthly average $\sigma_{\mathrm{perfect}}$ of 1.1 kW. Although $\mathrm{NL}_{\mathrm{perfect}}$ may be expected to be perfectly flat with zero variance, the optimized net load in the morning and evening usually ramps up and down over a finite time period (e.g. cyan in Figure 8a), causing non-zero variance. This ramp is a result of limited EV availability just after the first EV connects and before the last EV disconnects. For example, just the first EV by itself is not capable of dispatching sufficient charging power to follow a desired power r that is determined by spreading the energy of all EVs over the entire charging interval.

5.4.4 Characterization of Solar Forecast Error Impacts

The proposed receding horizon optimization with forecast error corrected by more accurate short-term forecasts ($G_{\mathrm{p+USI}}$) is generally effective in filling the energy valley. Exceptions are observed on Apr 5, 8, 9, 14, 16, 25, and 26 with a large peak in the optimized net load (Figure 5.7). To understand the challenges of the optimization on those days and further investigate the negative impact of solar forecast errors, we showcase a detailed example for Apr 14 in Figure 5.8.

On Apr 14, the baseline forecast G_{p} predicts large PV generation and net load valley (dotted blue) while this day is actually overcast with limited excess energy (solid blue). Between 06:30 PST and 12:00 PST, operating under the assumption of persistence forecast for >15 minute horizons, the algorithm expects a large energy valley later in the day. Therefore, it delays most of the non-critical EV charging while maintaining a flat net load profile. Because G_{p} is close to G_{perfect} (i.e. dotted blue follows solid blue) until 12:00 PST, the

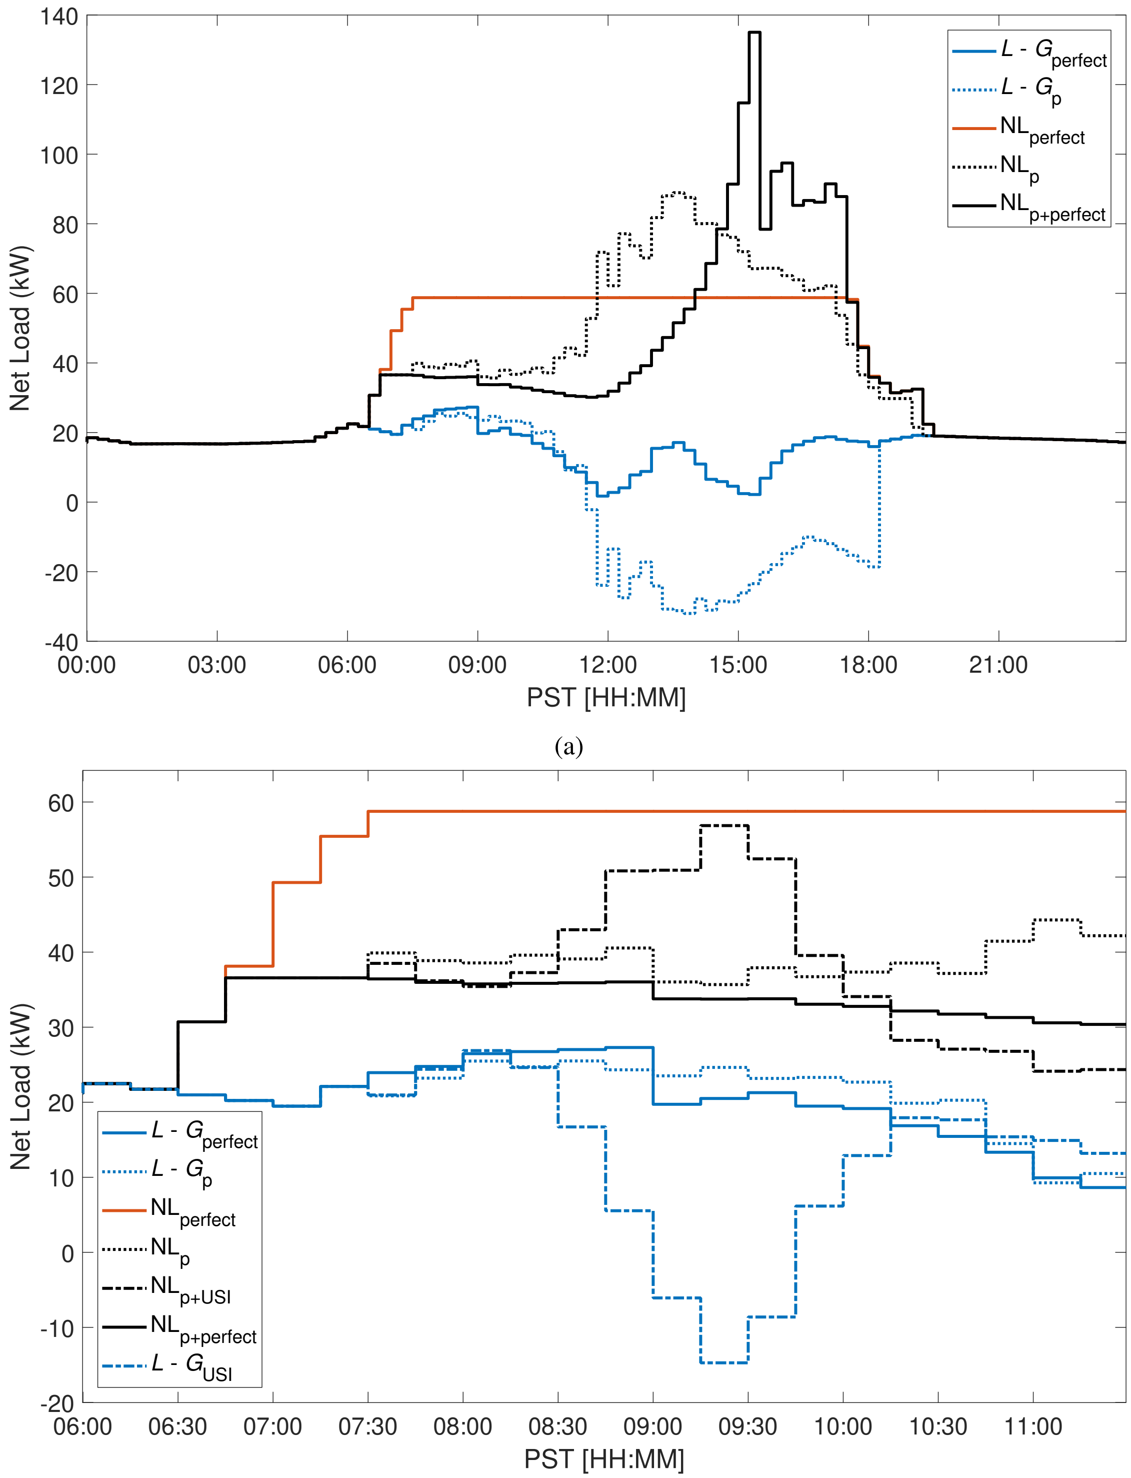

(a)

(b)

Figure 5.8: Comparison of valley filling performance using different solar forecasts (i.e. G_p, G_{p+USI}, $G_{p+perfect}$, $G_{perfect}$) on Apr 14, 2013 (a) and zooming in to the period from 06:00 to 11:00 PST (b). Line styles distinguish solar forecast scenarios, and blue and black colors differentiate net load with and without EV charging, respectively. The perfect net load forecast (load minus perfect PV forecast, solid blue) yields the ideal optimized net load (cyan).

EV charging is on an ideal trajectory; if the excellent accuracy of G_{p} had continued, the valley filling would have flattened the net load across the day. However, after 12:00 PST, G_{p} deviates from G_{perfect}. Even though the energy valley is in fact not available, the forecasts still predicts an abundant energy valley from 12:00 until 18:00 PST, thus EV charging is postponed to that time. The charge rate for each interval roughly equals the difference between f and G_{p}. Therefore, the variations in the resulting valley filling scenario NL_{p}, mirror the difference between perfect net load forecast ($L-G_{\mathrm{perfect}}$) and persistence net load forecast ($L-G_{\mathrm{p}}$). The peak error in G_{p} corresponds to the largest net load peak in NL_{p} at about 13:30 PST.

On this day, the day-ahead forecast error correction actually worsens the net load variability. Starting from 12:00 PST, when G_{p} is replaced with the USI forecast G_{USI} or the perfect forecast G_{perfect} for the next 15 min, the energy valley forecasted by G_{p} for the next 15 min is no longer available. Therefore, the optimization delays scheduling EV charging at the present time step, because it still expects an energy valley 15+ minutes from now (where G_{p} has not been updated yet). At the next time step, the situation is similar and EV charging is again rescheduled for later. Thus, non-critical EV charging is delayed until 13:30 PST when a number of early-departure EVs start to approach their charge time limit. Since those EVs must start to charge at their maximum charge capacity immediately to be fully charged (refer to Eqn. 5.11 for the constraint) by their planned departure time, the algorithm has to schedule those EVs immediately, independent of load flattening objectives, causing a steep net load rise. After 16:00 PST, another set of EVs approach their charge time limit, resulting in a second peak. Overall the resulting net load profiles $\mathrm{NL}_{\mathrm{p+USI}}$ and $\mathrm{NL}_{\mathrm{p+perfect}}$ resemble that of a maximum delay strategy, causing most EVs to concurrently charge at their full charge rate. On the contrary, operating under G_{p} only causes the algorithm to allocate the charge power more equally across the afternoon, reducing the peak in the corresponding optimized net load profile NL_{p} (dotted black). Consequently NL_{p} shows lower σ (24.6 kW) than $\mathrm{NL}_{\mathrm{p+USI}}$ and $\mathrm{NL}_{\mathrm{p+perfect}}$ (30.0 kW, 30.4 kW, Table 5.4). Lastly, the perfect forecast G_{perfect} (solid blue) yields a perfect valley filling $\mathrm{NL}_{\mathrm{p+perfect}}$ (cyan) during the EV layover time ($\sigma \cong 0$).

The subtle difference between the impacts of the three imperfect solar forecast inputs G_{p}, $G_{\mathrm{p+USI}}$ and $G_{\mathrm{p+perfect}}$ on the resulting net load profiles is further analyzed in Fig. 5.8a by zooming into the period from 06:00 to 11:00 PST. Before 07:30 PST, the lack of PV power output makes all optimized net load profiles identical. However, the algorithm still schedules to charge early-arrival EVs because of overall insufficient energy valley across the day, pushing the grid net load upward. Later on at 07:30 PST, PV starts generating power and because G_{p} and G_{USI} forecast errors differ, G_{p}, $G_{\mathrm{p+USI}}$ and $G_{\mathrm{p+perfect}}$ exhibit small optimized net load valleys. The original net load differences (all in blue) carry through to the resulting optimized net load profiles NL_{p}, $\mathrm{NL}_{\mathrm{p+USI}}$ and $\mathrm{NL}_{\mathrm{p+perfect}}$ (all in black). From 08:30 to 10:30 PST, erroneous G_{USI} causes a predicted net load valley (dashed blue), so the algorithm schedules EVs to charge, producing a net load peak (dashed black) at 09:00 PST. After 10:30 PST, the net load peak starts to diminish as G_{USI} starts to mostly follow G_{perfect} again. During the entire morning the valley filling resulting from the three imperfect solar forecasts deviates significantly from the idealized result (cyan). If the algorithm understood the true solar generation from the beginning of planning horizon, it would schedule the charging much earlier in the day and flatten the entire net load profile. This comparison further demonstrates the negative impact of forecast errors on flattening grid net load.

Furthermore, the day discussed in Figure 5.8 illustrates the relationship between day-to-day change in sky conditions (larger and more auto-correlated error in base forecast) and net load peaks (refer to Figure 5.7 and Table 5.4. On Apr 14 (the first of two consecutive overcast days of Apr 14 and Apr 15), the persistence solar forecast error, e_{p}^{G} is large (241% greater than $e_{\mathrm{p+USI}}^{G}$) because cloud conditions change from a clear day with few thin cirrus (Apr 13) to an overcast day with thick clouds. The high persistence error results in a large evening peak in optimized net load with $\sigma_{\mathrm{p}}^{\mathrm{NL}}$ of 24.6 kW. On Apr 14, correcting base forecast by better sky imagery forecasts actually worsens the net load variability ($\sigma_{\mathrm{p+USI}}^{\mathrm{NL}}$ = 30.0 kW). However, on the second overcast day (Apr 15), the persistence forecast performs better (74 percentage points), consequently both NL_{p} and $\mathrm{NL}_{\mathrm{p+USI}}$ performs better with $\sigma_{\mathrm{p}}^{\mathrm{NL}}$ and $\sigma_{\mathrm{p+USI}}^{\mathrm{NL}}$ reduced to 7.8 kW and 8.4 kW, and large net load peaks are eliminated. Similar load peaks occur on other days with changes in day-to-day sky conditions. If the present day is cloudier than the previous day, the day-ahead persistence forecast will cause the algorithm to push the EV charging peak forward (Apr 5, Apr 8, Apr 14, Apr 25). Conversely, if the present day has less clouds than the previous day, the persistence forecast will push the peak backward (Apr 9, Apr 16, Apr 26).

5.4.5 Discussion and Potential Limitations

The autocorrelation of the forecast error, i.e. persistent over- or underforecasts over a few hours, causes the largest deviations in optimized net load with erroneous forecasts from the optimized net load with perfect forecasts. Section 5.4.4 verifies that the forecast error structure is an important determinant of optimized net load variability. For example, as observed in Figure 5.8, if clear sky is predicted on a cloudy day, then the midday net load depression is over-forecast and EV charge schedule in the morning would be unnecessarily delayed. In the valley filling literature, forecast errors have mostly been modeled through statistical approaches. When forecast errors are sampled from a distribution such as in [79], then the forecast error autocorrelation is around zero making it more likely to happen that an overforecast during one interval is balanced by an underforecast in the following intervals. While forecast errors with zero autocorrelation will result in temporary small deviations in optimized net load, these idealistic forecast errors are unlikely to produce the cumulative effects that results in dramatic optimized net load deviations. This means that simple statistical models in the valley filling literature degrade the autocorrelation of forecast data, resulting in optimistic solutions.

The proposed corrective approach and the use of the realistic forecast data reveal that the deviation of practical valley filling performance from the ideal valley filling is tied to three conditions: 1) EV charging flexibility. Because of the constraint to satisfy EV energy needs, EVs with short layover periods and large energy needs limit scheduling flexibility. EV charging inflexibility prevents the load flattening objectives from being fully met even with perfect forecasts. 2) Accuracy of day-ahead forecasts. Large changes in cloud conditions on two consecutive days worsen day-ahead persistence forecasts, leading to large peaks in optimized net load. 3) Accuracy of the short-term forecast. Reducing forecast error by short-term forecast generally yields better valley filling performance (Figure 5.6). For a more accurate valley filling problem, the following improvements should be considered.

1) Accelerated adoption of EVs is needed to provide adequate amounts of controllable loads. In this work, the load and PV generation data had to be scaled down significantly for 31 EVs to just fully balance the energy valley. More financial incentives are necessary to encourage adoption of the smart charging standards.

2) The load demand timeseries was simulated through an energy simulation tool because feeder load data are typically not published by distribution system operators. As the simulated data may lack real power variability, the value of smart charging may be higher than shown in this paper.

3) Integrating economic incentives and EV battery aging models may make the study more comprehensive, but Wang et al. [106] demonstrated that peak load shaving at typical charge powers will not noticeably accelerate battery degradation in comparison to the degradation incurred from driving and calendar aging.

5.5 Conclusions

Electric utilities are experiencing unprecedented growth in the adoption of grid-connected solar PV. Solar forecasts are essential to the integration of PV and balancing supply and demand. This paper successfully demonstrates a corrective approach to mitigate negative impacts of day-ahead forecast errors using a sky imager based forecasting technique. The proposed methodology was tested against 24-hour persistence solar forecasts with one month of PV generation and substation load data. We showcase smart EV scheduling as a promising mechanism for absorbing the net load depression created by high PV penetration. This study differs from the literature regarding the forecast error modeling as we use real data to preserve autocorrelation characteristics of forecast errors and avoid statistical approximations. Our primary findings are:

1) Realistic forecast errors prevent the optimal charge strategy from flattening the net load.
2) More accurate short-term forecasts input to the corrective receding horizon optimization reduce net load variability.

Correcting the day-ahead persistence forecasts by sky imager forecasts for 15 min horizons reduces net load variability on 20 out of 30 days, and presents comparable results on the remaining days which are mostly clear. On average over the month, the standard deviation of the net load profile is reduced by 24.7%. On the 15 partly cloudy days, the short-term forecast correction reduces net load variability by 32.7%. On clear days, the proposed correction scenario optimizes net load variability to be below 3.0 kW given a scaled energy valley of 50 kW, which is only 1.5 kW worse than for a perfect forecast correction.

For future work, we will use real EV charge events and EV charging forecasts. More accurate day-ahead solar forecasts generated by numerical weather prediction and machine learning techniques will also be pursued. Because sub-optimal charging strategies worsen the net load variability, understanding possible variations associated with forecast errors around a single deterministic strategy will be of great benefit to grid operators. Finally, a stochastic optimization framework that integrates probabilistic PV forecast and EV availability and demand forecasts allows a more robust valley filling.

5.6 Appendix

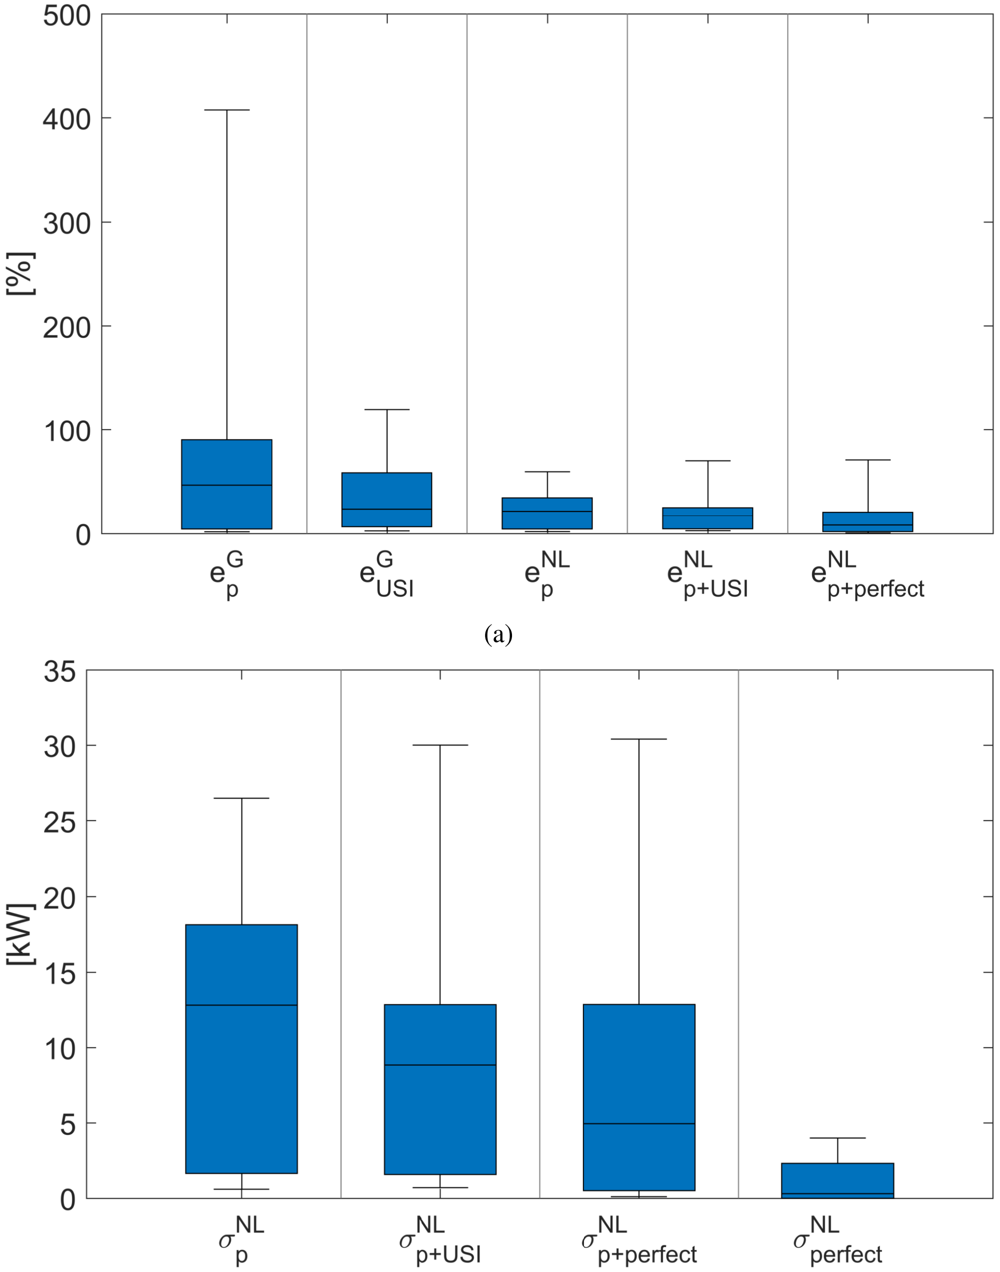

(a)

(b)

Figure 5.9: A visual representation in form of boxplot of the performance metrics in Table 5.4.

Table 5.3: Daily EV charging events database. Refer to Section 5.3 for a detailed definition of the variable names.

No.	EV Model	Arrival Time [PST]	Departure Time [PST]	Charge Capacity [kW]	Initial SOC [%]	Energy Demand [kWh]	Battery Capacity [kWh]
1	Tesla Model S	8:00	17:40	10	10	76.5	85
2	Toyota Prius	6:10	14:35	3.3	5	3.8	4
3	Chevy Volt	7:00	16:00	3.3	55	7.7	17
4	Toyota Prius	6:10	15:40	3.3	80	0.8	4
5	Nissan Leaf	6:30	16:45	6.6	95	1.2	24
6	Chevy Volt	8:30	19:15	3.3	15	14.5	17
7	Chevy Volt	8:00	17:10	3.3	60	6.8	17
8	Nissan Leaf	7:10	15:10	6.6	50	12	24
9	Toyota Prius	9:10	17:10	3.3	40	2.4	4
10	Toyota Prius	6:10	14:35	3.3	20	3.2	4
11	Chevy Volt	7:30	17:30	3.3	80	3.4	17
12	Toyota Prius	7:20	15:45	3.3	35	2.6	4
13	Toyota Prius	8:20	18:15	3.3	55	1.8	4
14	Nissan Leaf	8:20	16:40	6.6	20	19.2	24
15	Tesla Model S	6:50	17:20	10	65	29.8	85
16	Chevy Volt	7:30	16:10	3.3	30	11.9	17
17	Ford Fusion Energi	7:30	15:40	3.3	70	2.1	7
18	Toyota Prius	8:00	16:25	3.3	70	1.2	4
19	Tesla Model S	8:10	17:30	10	75	21.3	85
20	Toyota Prius	8:10	17:10	3.3	50	2	4
21	Nissan Leaf	7:00	15:40	6.6	10	21.6	24
22	Chevy Volt	8:00	18:00	3.3	30	11.9	17
23	Tesla Model S	8:00	17:15	10	90	8.5	85
24	Toyota Prius	6:40	15:50	3.3	20	3.2	4
25	Nissan Leaf	6:30	16:55	6.6	85	3.6	24
26	Tesla Model S	7:30	16:00	10	60	34	85
27	Tesla Model S	9:10	18:55	10	10	76.5	85
28	Ford Fusion Energi	7:10	16:50	3.3	50	3.5	7
29	Nissan Leaf	7:50	16:35	6.6	15	20.4	24
30	Chevy Volt	6:50	16:00	3.3	80	3.4	17
31	School Bus	11:00	16:15	70	50	57.5	115

Table 5.4: Day-by-day solar forecasts accuracy and optimized net load variability under different forecast error scenarios.

Date MM/DD/YY	e^{G}_{p} [%]	e^{G}_{USI} [%]	e^{NL}_{p} [%]	$e^{\text{NL}}_{\text{p+USI}}$ [%]	$e^{\text{NL}}_{\text{p+perfect}}$ [%]	$\sigma^{\text{NL}}_{\text{p}}$ [kW]	$\sigma^{\text{NL}}_{\text{p+USI}}$ [kW]	$\sigma^{\text{NL}}_{\text{p+perfect}}$ [kW]	$\sigma^{\text{NL}}_{\text{perfect}}$ [kW]
04/1/2013	103.8	86.8	49	25	9.6	18.1	14.3	5.2	0.8
04/2/2013	63.6	30.6	46.4	25.9	18.8	19.8	10.7	6.9	0.3
04/3/2013	24.6	16.8	22.1	9	3.5	9.5	3.8	1.3	0
04/4/2013	29.9	42.6	17.9	20	9.9	8.6	9.8	4.7	0
04/5/2013	189.1	96.9	16.9	20	20.7	15.4	15.1	15.7	3.8
04/6/2013	87.1	45.5	34	27.9	27.3	17.4	14.3	13.7	1.2
04/7/2013	73.9	54.5	23.4	16	13.5	15.3	10.4	8.9	2.7
04/8/2013	210.9	99.9	23.6	20.4	21.3	20.9	17.6	18.5	3.9
04/9/2013	125.5	4	58.2	70	70.9	18.1	21.3	21.2	0.3
04/10/2013	4.4	4.9	4.4	3.9	1.4	1.7	1.5	0.5	0.4
04/11/2013	10.9	12.6	9.2	6.5	6.7	3.1	2.1	2	0.1
04/12/2013	34.2	38.1	31.6	22.6	4	12.8	9.2	1.3	0.4
04/13/2013	89	58.1	40.8	17.9	11.3	25	11.5	9.3	3
04/14/2013	407.4	119.5	20.9	33	31.1	24.6	30	30.4	4
04/15/2013	74.1	111.1	6.5	8.2	1.8	7.8	8.4	5.8	3.5
04/16/2013	127.4	62.8	54.8	44.4	45.9	26.5	21.7	22.4	1.9
04/17/2013	62.7	6	59.6	29.3	33.7	18.4	9.2	10.5	0
04/18/2013	4	4.4	5.6	3.9	1.2	1.9	1.1	0.3	0
04/19/2013	3.4	3.8	3.7	4.1	2.9	1.2	1.1	0.7	0
04/20/2013	2.9	2.5	2.3	2.6	1.9	0.6	0.7	0.5	0
04/21/2013	1.8	7.8	1.9	5.9	0.6	0.6	2	0.2	0
04/22/2013	1.6	5	2.1	3.9	0.5	0.7	1.3	0.1	0
04/23/2013	5.6	10.6	6.5	7.4	1.5	2.2	2.6	0.4	0
04/24/2013	44.8	39.2	28.5	17.3	14.1	12.8	7.9	6.4	2.3
04/25/2013	124.4	75.8	33.9	19	19.2	22.6	12.7	12.8	3.6
04/26/2013	78.3	7.6	50.6	31.7	39.7	17	10.5	13.3	0.1
04/27/2013	3.1	7	3.7	4.6	0.9	1.3	1.6	0.3	0.1
04/28/2013	2.7	6.5	2.7	4.3	1.3	1	1.6	0.4	0.2
04/29/2013	4.7	6.5	3.7	3.5	2	1.4	1.1	0.8	0.4
04/30/2013	48.7	47.7	29.4	17.5	5.4	14.2	8.5	2.2	0

5.7 Acknowledgment

This Chapter, in full, is a reprint of the material as it appears in "Wang, G. C., Ratnam, E., Haghi, H. V. & Kleissl, J. (2019): Corrective receding horizon EV charge scheduling using short-term solar forecasting. *Renewable energy*, 130, 1146-1158." The dissertation author was the primary investigator and author of this paper.

Chapter 6

Stochastic Optimization of EV Scheduling Incorporating Real-Time Charge Events

6.1 Introduction

High photovoltaic (PV) penetration is a primary determinant of net load variability on distribution feeders, potentially causing reliability issues and reverse power flow [107]. Addressing net load variations through energy storage systems (ESSs) and predictive charging strategies has been comprehensively studied over the past decade [108, 109]. Electric vehicles (EVs) are becoming a popular distributed ESS to help mitigate net load fluctuations in the power grid due to the low additional investment cost. But the rapid adoption of EVs can negatively influence the grid network [12]. For example, uncoordinated EV charging behavior will create load peaks on the aggregated load curve, leading to power quality degradation [13] and operational cost increase [14]. Furthermore, sub-optimal charging scheduling causes additional net load peaks or worsens power grid operation [110].

Researchers have examined EV charging scheduling problems considering uncertainties of EV behavior with the objectives of frequency regulation [111], power quality impact [112], and cost minimization [113]. The uncertainties in EV availabilities are typically modeled with general assumptions using a Poisson distribution [114] or a normal distribution [115]. These simplistic distributions are inadequate to capture the stochastic nature of EV availability. Other more sophisticated approaches, including Markov chain model [116], Kernel-based method [117, 118], and Monte Carlo simulations are also frequently utilized to model

EV availability, charging behavior, locations, and profiles based on empirical distributions [119, 120].

Stochastic programming is a common approach for modeling decision making problems that involve uncertainty [121]. It requires scenarios that can be constructed through the aforementioned models or probability distributions to represent realizations of the uncertain data. The expected value of the objective function can be replaced with a statistical sample average (SAA, [122]) of the scenarios and the solution to this objective function converges to the true solution when the number of scenarios tends to infinity. This approach has proven to be effective when dealing with uncertain parameters [123]. For example, a two-stage stochastic framework was developed by [124] to address the stochastic nature of PV generation, loads, EV availability and energy demand in a cost saving problem. However, the SAA approach only yields a single averaged problem statement, which can result in suboptimal results due to the lack of representation from the possible variations away from the averaged solution. In [125], empirical characteristics of EV availability and energy demand were extracted to generate scenarios that were independently solved in a deterministic fashion. The resulting EV charging solutions were transformed into a weighted quantile distribution to provide potential net load variations. On the other hand, if the computational burden is not a concern, then an alternative approach is to solve a multitude of scenarios and assume the average of resulting solutions as the optimal solution for all scenarios.

The aforementioned probabilistic models and simple distributions barely fully characterize the stochastic nature of EV behavior and solar energy, so deviations are always expected in the resulting scenarios from the real EV charge events and actual PV productions. Moreover, in grid resource planning problems, scenario-based approaches typically use day-ahead forecasts to determine the scheduling strategies, but such solutions do not guarantee the satisfaction of real energy dispatch needs despite the consideration of scenarios since these scenarios were generated based on the information at the moment of scenario generation. Further, a simple average over all independently solved solutions cannot be justified as a true optimal solution because this optimal solution will not satisfy the objective function in every scenario, while each scenario is expected to be one of legitimate descriptions of the problem.

We presented in [110] that a corrective approach is able to remedy the impact of forecast errors and support grid operation through a deterministic problem setup. In this paper, we extend the work in [110] to address the above challenges. Specifically, we advance the prior deterministic optimization framework by considering the uncertainties in PV forecasts and EV charge events in form of scenarios, but then correcting the scenarios with real EV charge events and PV measurements in a predictive control based stochastic optimization framework. The common good EV charging strategy out of all scenarios is determined, and executed to solve the grid planning problem of interest. The proposed approach satisfies the real charge needs while is friendly to the operational use. Herein, our objectives are as follows:

1) Eliminate the averaging ambiguity by solving a common good solution out of the stochastic objective function using all scenarios at once, yielding an optimal solution with greater confidence.
2) Incorporate real-time grid data updates into the proposed stochastic optimization framework, demonstrating the ability to mitigate the impact of solar forecast errors and uncertain grid data on real world grid planning problems.
3) The charging solutions derived from a typical scenario based approach using day-ahead forecasts and perfect forecasts are presented as upper-bound and lower-bound benchmarks, respectively.

The remainder of this paper is organized as follows. Section 6.2 introduces the available EV and PV data available for our analysis, and the methodology to model the uncertainties in PV generation and EV availability. Section 6.3 defines the problem statement and objective function, introduces the methodology to produce PV forecast scenarios and EV scenarios, and proposes a corrective stochastic optimization framework implemented under a predictive control based scheme. Section 6.4 provides a detailed analysis on a sample day, and statistical results from 19 days of validation of our corrective approach. Lastly, Section 6.5 provides conclusions and future work.

6.2 Data Processing

6.2.1 EV Database

Due to the lack of real charging data, many studies constructed EV databases from travel surveys, transportation statistics and/or other assumptions. Real-world data avoids assumptions regarding the stochastic nature of vehicle driving patterns and improves the accuracy of simulated EV availability. The EV database collected at the University of California, San Diego (UCSD) campus is used in this work. There are 124 Level II charging stations managed by ChargePoint around the campus. The high concentration of EV charging creates an ideal testbed for our analysis. Data records are comprised of connection time, disconnect time, EV ID, energy consumption and charge location. In total, more than 51k charge events were retrieved from this database for the time period March 2016 to November 2018. To make the database more suitable for the workplace EV charging problem, we removed weekend and overnight charge events. Also, the charging sessions that lasted less than 15 min were eliminated as energy consumption was minimal and charging flexibility was absent in those cases. Finally, more than 38k charge events remain, spanning 972 days.

While the charge capacity of the charging station is known (i.e., level II standard), the actual charging capacity is mostly limited by the EV's onboard charger, which is not included in the dataset. Given the energy consumption and the layover period recorded in each charge event, the average charge rate can be derived per the procedure in [125]. The most commonly occurring charge rate can be assumed as the charge capacities of the EVs. The remaining sporadic charge events are either outliers or events that include idle connection presenting charging flexibility. In all 38k charge events, the primary charging capacities were found to be either 3.3 kW or 6.6 kW.

6.2.2 Empirical Modeling

Probability Distribution of EV Arrivals

The probability of EV arrivals in the time window of interest on any given workday is approximated by an empirical modeling approach. Since the impact of seasonal effects is small in southern California, it is not considered in this model. Consequently, the charge events of all 972 days are aggregated to 15 minute intervals. At time window i, the probability for v EV arrivals can be computed as follows:

$$p_v(i) = \frac{m_v(i)}{N_D}. \quad i \in [1, \dots H] \tag{6.1}$$

In this work $H = 96$ is the maximum planning horizon given time interval $\Delta t = 15$ min, while $p_v(t)$ and $m_v(i)$ represent the probability and number of occurrences of v EV arrivals during time window i, respectively. The total number of days N_D is 972. Consequently, the probability of 0 EV arrivals in each time interval can be derived:

$$p_0(i) = 1 - \sum_v p_v(i). \quad i \in [1, \dots H] \tag{6.2}$$

For better illustration, Fig. 6.1 illustrates the distribution of EV arrivals in the quarter-hourly interval. Note that 0 EV arrival is always the most likely case due to the short time window and limited number of EVs.

The probability distributions in Fig. 6.1 suggests that fewest EVs arrives in the late night and most arrives in the early morning.

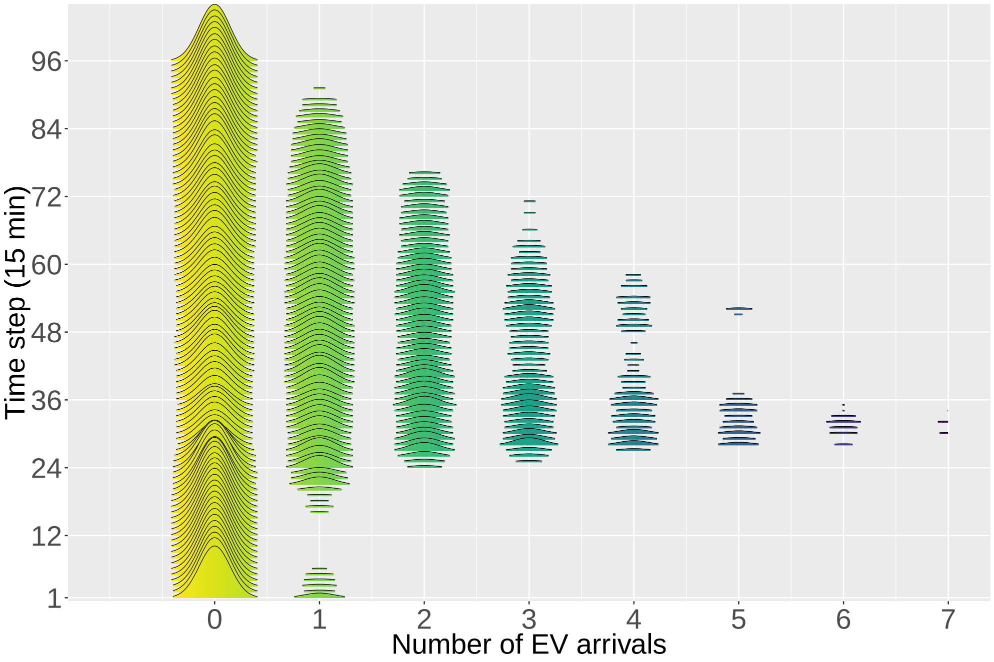

Figure 6.1: Ridgeline plot of probability distribution for the occurrence of EV arrivals in 15 min interval. The probability is calculated by Eqns. (6.1) and (6.2).

Deriving Probabilistic EV Departure Times and Energy Demand

EV departure time and energy demand are functions of the EV owners work schedule so similar work schedules can be grouped by using clustering methods [126], and *k*-means clustering is used in this work. Given the number of cluster centroids *k*, the observations (charge events) are iteratively partitioned to different cluster centroids until the distances between the points and their respective centroids are minimized. Mathematically:

$$\min \sum_{I=1}^{k} \sum_{X \in c_I} ||X - \mu_I||^2, \tag{6.3}$$

where $||\dots||$ indicates the Euclidean distance, X represents all observations, and c_I are the observations assigned to cluster I, whose centroid is represented by μ_I. While there are a number of cluster separation measures, we use the elbow method [127] to determine the optimal number of partitions k. For our analysis, a range of 1-20 clusters was initially considered, and 11 clusters were found to be optimal. We then partitioned the charge events considering $k = 11$ using Eqn. (6.3). Figure 6.2 illustrates the optimal partitions adopted in this paper.

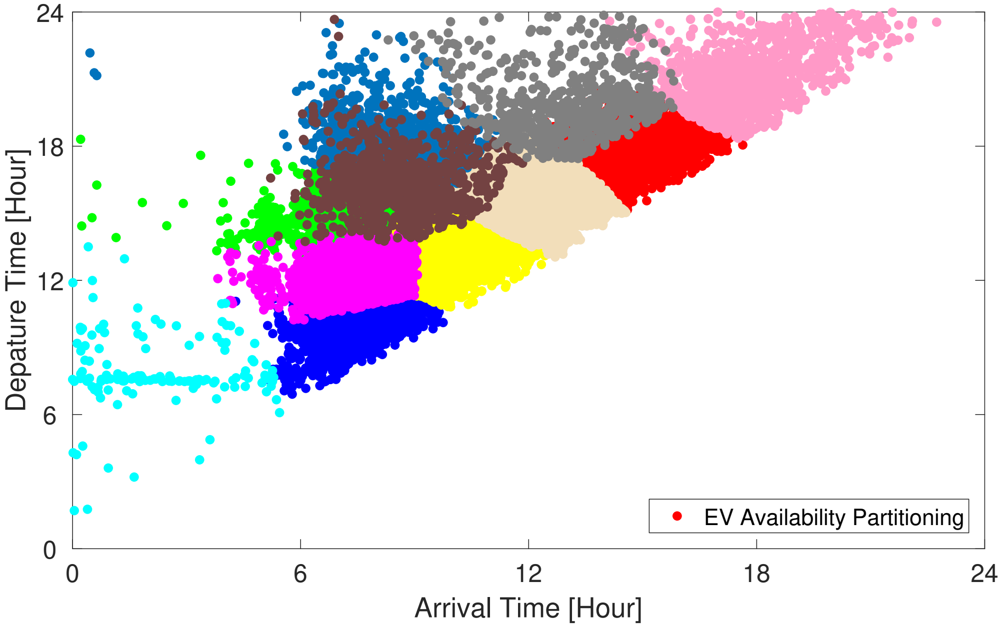

Figure 6.2: Arrival and departure times of all charge events are segmented to 11 clusters through k-means clustering. The partitions are distinguished by color.

6.2.3 PV Database

1) PV system: The PV power produced by a 3.09 MW plant and 2 MW inverter has been recorded at a 1 s sampling rate from June 28, 2012 to May 10, 2013 and has been aggregated to 15 minute averages for the purpose of this study. A significant part of January 2013 and February 2013 is missing, and these months have therefore been removed entirely from the dataset. Furthermore, the dataset contains days at which the PV power has been capped at 1.5 MW and all such days have been removed. The training and validation sets encompass data from June 28, 2012 to March 31, 2013 whereas the test set covers the full month of April of 2013.

2) NWP: To produce PV forecasts from one step ahead to 24 hours ahead, it is helpful to include Numerical Weather Prediction (NWP) forecasts. In this work, we incorporate the 1 to 24 hour ahead forecasts of the North American Mesoscale (NAM) forecast system[1], generated four times per day, at a spatial resolution of 12 km^2 at 35° latitude. However, we only use the forecasts generated at mid-night UTC.

---

[1] https://www.ncdc.noaa.gov/data-access/model-data/model-datasets/north- american-mesoscale-forecast-system-nam

Since these forecasts are in hourly resolution, we linearly interpolate the values to the 15-minute resolution required by our analysis. Furthermore, we take the average of the forecasts issued at the four nearest grid points.

3) Solar geometry: Although the NWP model forecasts the downward radiation and it is therefore known when the sun rises and sets, the low resolution of these forecasts requires us to incorporate solar geometry at a higher resolution.

6.2.4 Quantile Regression Forests (QRF)

We employ QRF to quantify the uncertainty by means of the predictive distribution. It is an extension of random forests (RFs) developed by Breiman [128]. Regression trees produce predictions with low bias but high variance, which can be improved by bootstrap aggregation where the final prediction is the result of the average of many trees [129]. In RFs, each tree is grown on a random subset of the explanatory variables. This reduces the variance and allows RFs to predict with comparable accuracy as boosted trees, but more straightforward to train and tune. The detailed implementation of QRFs can be found in [130]. We use the R package quantregForest [131] to produce the probabilistic forecasts.

6.3 Methodology

6.3.1 EV Charge Event Scenario Generation

A scenario is a specific realization of an uncertain event. In this work, an EV charge event scenario is defined as sequential EV charge events occurring in one day. We sample each EV scenario by the following steps:

(a) Set a 24-hour time window consisting of sequential time intervals $i \in [1 \dots H]$.

(b) At $i = 1$, a random number is generated to sample the number of EV arrivals from the actual population derived in Section 6.2.2.

(c) If the number of EV arrivals sampled in Step (b) is greater than zero, an EV departure time is randomly picked from the clusters which cover the present arrival time per Fig. 6.2. The energy demand associated with this select EV departure time is also picked; otherwise proceed with (e).

(d) The EVs charging capacity is randomly sampled from two equally probable charging capacities derived in Section 6.2.1.

(e) Move to the next time window by $i = i + 1$ and repeat from (b) until $i = H$.

(f) Repeat (a)-(e) for generating additional scenarios.

Note that although the battery capacity and initial state of charge (SOC) are not always available, they are not necessary in this work since we enforce complete charging per the energy requirement during the layover period. An assumption of an 85 kWh battery capacity (larger than any vehicle battery capacity in our database) is sufficient for the analysis, and the initial SOC can be derived with the battery capacity and energy demand.

6.3.2 PV Forecast Scenario Generation

PV forecast scenarios are generated from predictive distributions [132]. However, these scenarios are not innately autocorrelated, meaning that they likely exhibit significant variability, which may lead to unrealistic results. Herein, we employ a Gaussian copula that describes the dependence between the forecast horizons, which allows us to generate sequences of autocorrelated random numbers [133]. Readers are referred to [134] for the detailed implementations.

In short, we convert a sequence of autocorrelated random numbers to a scenario of 1 hour - 24 hour PV forecasts by using the integral transform. The PV forecasts are scaled down by a factor of 150 to yield a comparable energy demand to the aggregated EV charge needs in this study.

6.3.3 Scenario Corrections Using Real-time Updates

The scenarios produced in the previous subsections inherently deviate from the real data. To remedy the modeling error for better valley filling performance while meet the real-world charge needs, a scenario correction methodology is proposed. First we assume the departures and charge needs are specified by the EV owners upon their arrivals. At each time step, an EV session matching process is performed throughout the EV scenarios by validating the modeled EV arrivals occurring at the present time step against the actual EV arrivals, and the modeled departures and charge needs are corrected by the counterparts in the real EV sessions, if inconsistencies are found. The modeled EV sessions occurring in future time steps in reference to current time step are left unchanged. This matching process is iterated as time window marches forward.

On the other hand, 24 hour ahead PV forecasts in each PV scenario are continuously updated by the newly issued 1 to H step ahead PV forecasts, and the process is also iterated over time. The corrected EV and PV scenarios are used in the objective function.

6.3.4 Stochastic Optimization Framework

In this section, we advance the prior deterministic optimization algorithm in [110] to support stochastic data input. A typical stochastic framework constitutes two stages: first-stage decisions need to be taken here and now, whereas second stage decisions are taken after the uncertain events have materialized, but have to be accounted for in the decision-making stage. In this study, the common first stage decision variables are not considered, so the objective function only includes the expected value of the second stage variables, as presented in Eqn. (6.4). To allow implementing scenario correction over time, a moving time horizon mechanism is required. It improves the operation as opposed to a single optimization and control running over the entire time horizon because it allows us to introduce improved/updated predictions into the decision-making process. To proceed, we solve the following valley filling problem using all modeled EV and PV scenarios S in a 24 hour moving planning horizon:

$$\min_x \sum_s \sum_t |h_t x_{t,s} - f_s|. \quad \forall t \in T \; \forall s \in S \tag{6.4}$$

subject to:

$$x_{t,s} = L_{t,s} - G_{t,s} + \sum_v \beta_{v,t,s}. \quad \forall t \in T \; \forall s \in S \; \forall v \in V \tag{6.5}$$

$$0 \leq \beta_{v,t,s} \leq \overline{\beta}_{v,t,s} \gamma_{v,t,s}. \quad \forall t \in T \; \forall s \in S \; \forall v \in V \tag{6.6}$$

$$\sum_t \beta_{v,t,s} \gamma_{v,t,s} \Delta t = \chi_{v,t,s}^{\text{ED}}. \quad \forall t \in T \; \forall s \in S \; \forall v \in V \tag{6.7}$$

$$\gamma_{v,t,s} = \begin{cases} 1 & t_{v,s}^{\text{arr}} < t < t_{v,s}^{\text{dep}}. \\ 0 & \text{otherwise.} \end{cases} \quad \forall t \in T \; \forall s \in S \; \forall v \in V \tag{6.8}$$

$$\beta_{v,T(1),s} = \beta_v^{\text{opt}}. \quad \forall s \in S \; \forall v \in V \tag{6.9}$$

$$x_{T(1)}^{\text{opt}} = L_{T(1),\text{perfect}} - G_{T(1),\text{perfect}} + \sum_{v} \beta_{v}^{\text{opt}}. \quad \forall v \in V \tag{6.10}$$

The objective of the valley filling problem is to minimize the power difference between net load power profile $x_{t,s}$ and a preferred power profile f_s (Eqn. 6.4) subject to the power balance constraint (Eqn. 6.5), EV charge power constraint (Eqn. 6.6), the SOC constraint (Eqn. 6.7), and the EV charging availability represented by the parameters $\gamma_{v,t,s}$ (Eqn. 6.8). Specifically, in Eqn. (6.4), h_t is the matrix of the charging priority, and f_s is the vector of the preferred power of scenario s (refer to [110] for the detailed calculation). In Eqn. (6.5), $\beta_{v,t,s}$ is the charging power of the EV v in scenario s, $G_{t,s}$ and $L_{t,s}$ are the PV production forecasts and loads in scenario s, respectively. In Eqn. (6.6), the EV charging power β is constrained by the respective maximum power limit $\overline{\beta}$ and the minimum power limit 0 as Vehicle to Grid (V2G) is not considered in this work. In Eqn. (6.7), the total charged energy for an EV is limited to its energy demand χ^{ED}. The available charging period for each EV is constrained by its arrival time t^{arr} and scheduled departure time t^{dep}, as presented in Eqn. (6.8). The optimization utilizes all scenario S and solves the objective function (Eqn. 6.4) at a time by constructing a joint constraint for the first time step across all scenarios, as illustrated by Eqn. (6.9). The optimal solution that is common good to all scenarios is then solved and executed and the energy demand in each scenario is updated accordingly. The resulting net load power is determined in Eqn. (6.10). To proceed, the planning horizon T moves one-step forward and the objective function (Eqn. 6.4) is solved again by incrementing k in planning horizon T until the end of planning horizon.

6.3.5 Model Evaluation

We construct four valley filling use cases with an ideal valley filling approach and three practical valley filling approaches:

1) Optimized net load by ideal approach $\text{NL}_{\text{perfect}}$: solve Eqns. (6.4) - (6.8) once at midnight using one scenario of real PV productions and one scenario of EV availability and charge needs. If EV availability was unconstrained, the optimized charge solutions over planning horizon T are implemented in real time operation, which will yield a flat net load curve.

2) Optimized net load by day-ahead averaging approach NL_{DA}: solve Eqns. (6.4) - (6.8) once at mid-night only using the modeled scenarios of day-ahead PV forecasts and EV availability and charge needs.

The resulting day-ahead charging solutions over planning horizon T are implemented regardless of the actual EV charge needs, which is expected to yield a sub-optimal net load curve.

3) Optimized net load by corrective approach $\mathrm{NL}_{\mathrm{corrective}}$: solve Eqns. (6.4) - (6.10) recursively starting mid-night using modeled scenarios and real-time updates as presented in Section 6.3. This approach determines an optimal charging solution only for the next time step, which is expected to yield a near-optimal net load curve.

4) Dumb approach with uncontrolled charging strategy $\mathrm{NL}_{\mathrm{dumb}}$: Upon arrival, EV is scheduled to charge right away until the charge needs is satisfied, which is expected to cause giant power ramps on the net load curve.

The valley filling results by the three practical approaches 2)-4) are evaluated against the ideal approach 1) through normalized root mean square net load difference:

$$\varepsilon_{\mathrm{DA}} = \frac{\sqrt{\frac{1}{H}\sum_{1}^{H}\left[\mathrm{NL}_{\mathrm{DA}}(t) - \mathrm{NL}_{\mathrm{perfect}}(t)\right]^2}}{\max\left(\mathrm{NL}_{\mathrm{perfect}}(t)\right)} \times 100\%. \tag{6.11}$$

$$\varepsilon_{\mathrm{corrective}} = \frac{\sqrt{\frac{1}{H}\sum_{1}^{H}\left[\mathrm{NL}_{\mathrm{corrective}}(t) - \mathrm{NL}_{\mathrm{perfect}}(t)\right]^2}}{\max\left(\mathrm{NL}_{\mathrm{perfect}}(t)\right)} \times 100\%. \tag{6.12}$$

$$\varepsilon_{\mathrm{dumb}} = \frac{\sqrt{\frac{1}{H}\sum_{1}^{H}\left[\mathrm{NL}_{\mathrm{dumb}}(t) - \mathrm{NL}_{\mathrm{perfect}}(t)\right]^2}}{\max\left(\mathrm{NL}_{\mathrm{perfect}}(t)\right)} \times 100\%. \tag{6.13}$$

6.4 Results and Discussions

6.4.1 Aggregate Valley Filling Statistics

In this section, the same valley filling problem in [110] is solved by the proposed approaches introduced in Section 6.3.5 for a test period of 19 days. The simulations are performed under 50, 200 and 1000 scenarios, and the valley filling performances are evaluated and illustrated in Fig. 6.3. Overall, the corrective approach exhibits the best valley filling performance approach by improving the averaged nRMSD by more than 10 percentage points than the day-ahead approach and more than 40 percentage points than the dumb charging approach. The variability of net load is reduced to an nRMSD of below 10% with the proposed approach and 500 scenarios. The nRMSD stabilizes when 100 scenarios are used in both day-ahead averaging and corrective approach and the performance gain is minimal beyond using 100 scenarios. It is demonstrated that the predictive control based framework incorporating real measurements contributes more to mitigate the impact of errors in the scenarios than the scenario size and accuracy.

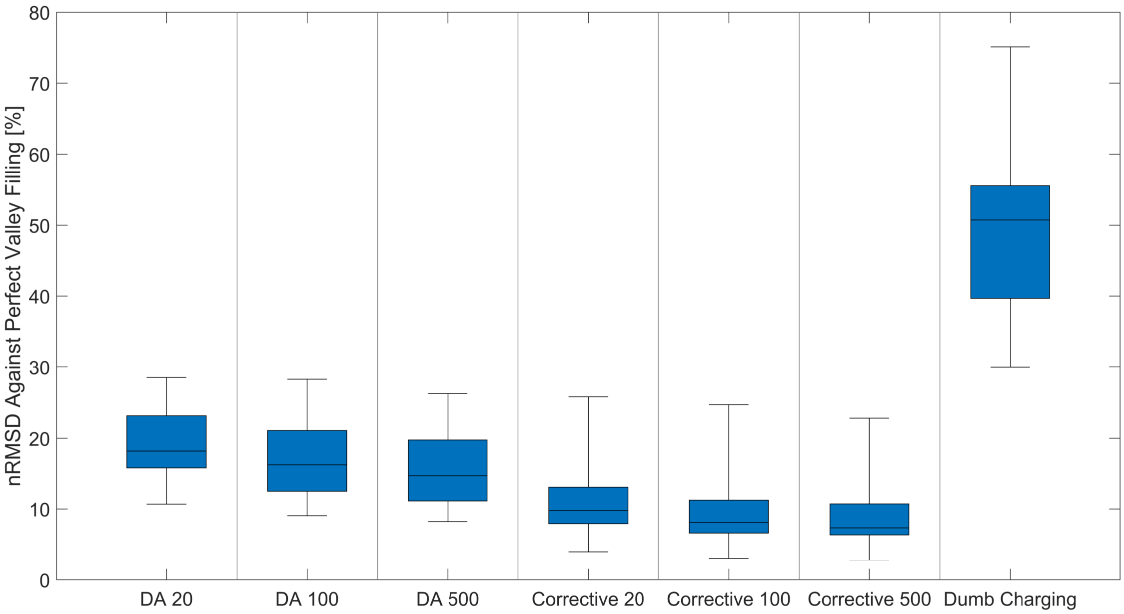

Figure 6.3: Box plot of valley filling performances between three practical approaches validated against the ideal approach over 19 days.

6.4.2 Characterizing the Impact of Modeling Errors in EV and PV Scenarios

In this section, we showcase a detailed example for Apr 21 in Figure 6.4 to further understand the deviation of the results of corrective approach (left) and day-ahead averaging approach (right) from the ideal valley filling results.

The power level derived by corrective approach rises to above the ideal net load profile early in the day but then drops at a later time and becomes fairly flat at the end of the day. We explained this phenomena in [110] that under-forecasts later in the day can cause the algorithm to schedule the EV charging earlier, resulting in a power peak. The top plot in Figure 6.5 confirms the conclusion by showing the deviation between 1000 PV forecast scenarios, the average forecast, and the real PV production. The mean of the 1000 scenarios is below the actual PV generation between 11:00 to 16:00 PDT caused by a slight cloud cover prediction in the NWP model, which in return causes the charging to be scheduled earlier. The bottom plot

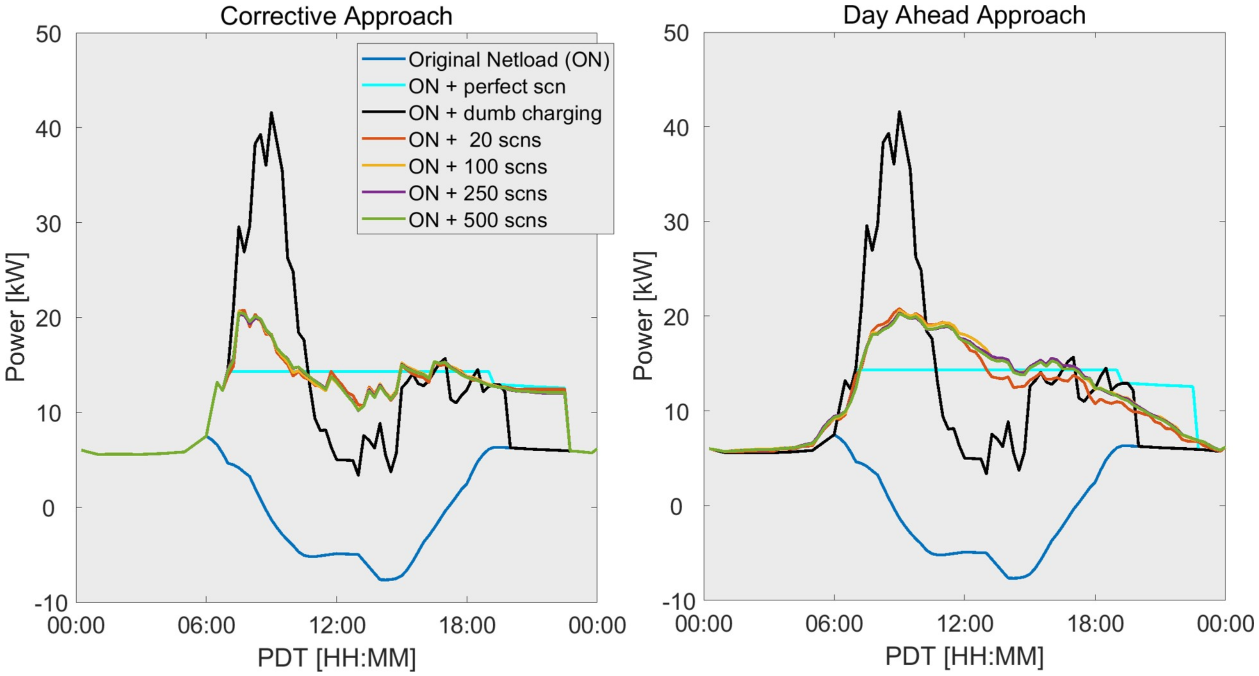

Figure 6.4: Comparison of net load valley filling performance between corrective approach (left) and day-ahead approach (right) on April 21. The cyan colored curve represents the ideal net load profile by using real PV and EV data in the day-ahead approach, the black curve illustrates the grid network quality is severely impacted by the uncontrolled EV charging, while other curves distinguish the valley filling results using different number of scenarios.

in Figure 6.5 depicts another root cause of the power peak by showing the deviation between modeled EV availability and real EV availability. From 9:00 to 12:00 PDT, the simulated EV charge events exhibit greater availability (as well as charge needs which are not displayed in the figure) than the actual, so the algorithm tends to schedule EVs during the morning to accommodate more EV arrivals later in the day. After 18:00 PDT, there is no solar production available, which can be easily predicted so PV forecast error becomes zero. Consequently, the variability in the net load profile is much smoother because the modeled EV availability and charge needs become the only error sources. The optimized net load is flattened, though a little lower than the ideal net load profile as a result of limited EV availability in the last few hours of the day.

Comparably, the right subplot in Figure 6.4 presents the net load profile after the day-ahead optimized charge solutions are exercised regardless of the changes in EV availability and PV production. The power level in the daytime is higher than the ideal power level (cyan) because the modeled EV charge needs are greater than the actual, resulting in higher power draw from the grid. After 18:00 PDT, the net load curves show a downward trend because the net load power in that period is an averaged value over the solutions of all scenarios. This demonstrates that the charging solutions that were considered optimal given the information at midnight will become sub-optimal during the actual execution, and that constitutes the value of the proposed corrective approach.

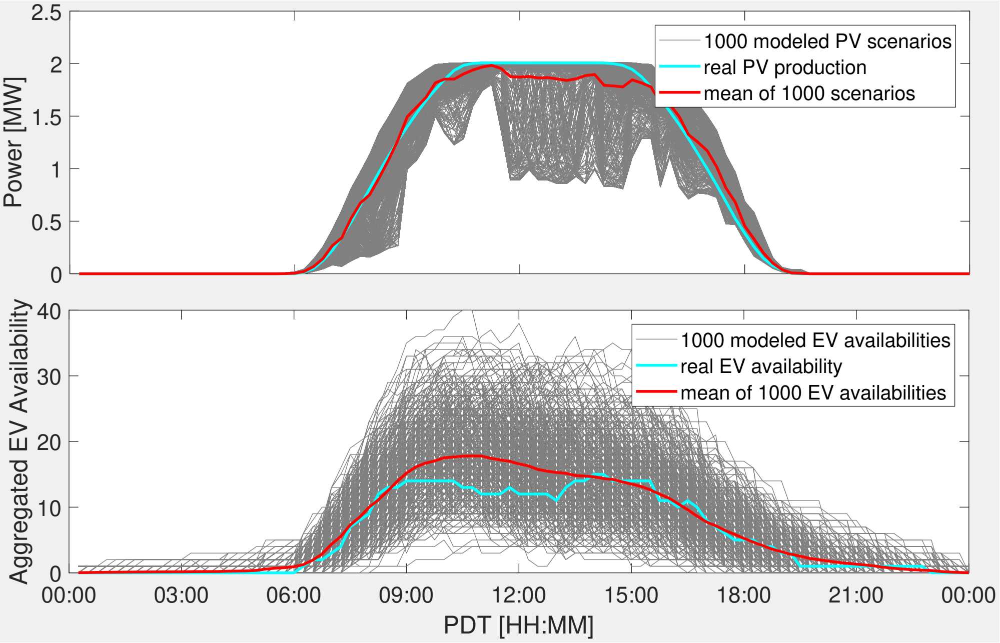

Figure 6.5: Comparison between 1000 modeled forecast scenarios (grey), averaged forecast scenario (red), and real measurement (cyan) for PV generation and EV availability on April 21.

The variability difference between the two optimized net load profiles and ideal profiles (cyan) represents the cost of perfect information on future events. In reality, it is impossible to produce day-ahead perfect forecasts, and this example demonstrates that great valley filling performance can be achieved when up-to-date information is utilized. Note that the simulated substation loads were scaled down in this study to align the orders of magnitude of the EV charge needs. While the steep net load rise early and late in the day indicates the energy valley is not sufficient for this day, the general methodology still applies.

6.4.3 Discussion

Since the day-ahead averaging approach determines an optimal solution only using the information at mid-night, the power mismatch between the day-ahead planning and real-time execution is expected. Figure 6.6 compares the satisfaction of EV energy demand with the three realistic approaches. The values indicates the charged energy in reference to the actual charge needs in percentage. For day-ahead averaging approach, the charge needs are not necessarily met because the real time EV charge events are not considered in the planning stage. For corrective and dumb charging approach, the real time EV charge events are used, so the charge needs are always fully satisfied.

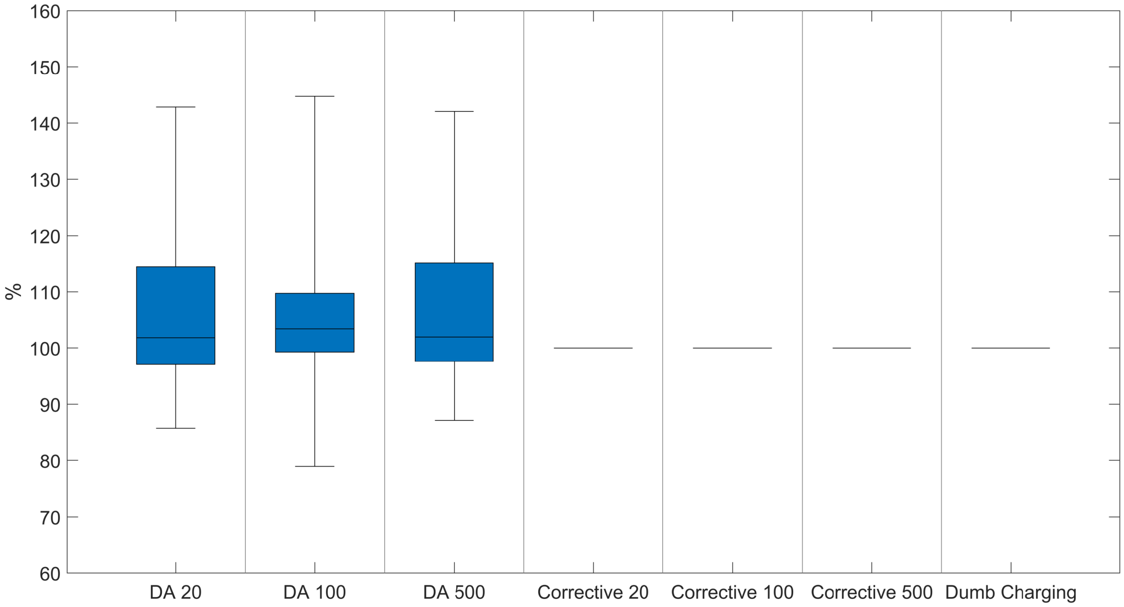

Figure 6.6: Boxplot of EV charge needs satisfaction between day-ahead averaging, corrective and dumb charging approaches over 19 days.

A potential concern of the scenario-based approaches to solving the stochastic objective function is their computational cost due to the large sample size often required to obtain an adequate representation of the original problem. In this study, we find with our corrective approach, 50 scenarios are sufficient to yield good valley filling results, which can be solved in 3 min on a common Intel i5 workstation.

6.5 Conclusions

A large penetration of EVs is expected to provide adequate amounts of controllable loads and can if charged at the workplace - help flatten the midday net load valley caused by excess solar generation. One of the key concerns in its practical implementation lies in the uncertain availability of EVs and PV power production. This paper has presented a novel predictive control based stochastic framework to mitigate the negative impacts of day-ahead PV and EV forecast errors and uncertain events using most recent PV forecasts and real time charge events. Specifically, we modeled the uncertainty in EV availability using a large database of EV charge events at UC San Diego, and built the probabilistic PV forecasts using quantile regression forests. The modeled uncertainties are converted to scenarios which are solved by the proposed framework. The common good solution out of all scenarios is solved without the ambiguous averaging operation. The proposed methodology was tested against two common approaches adopted in the grid integration studies: 1) dumb charging approach 2) day-ahead averaging approach. Our proposed method exhibits the best valley filling performance by reducing the grid net load variability to an nRMSD of less than 10%, outperforming the dumb charging approach and day-ahead averaging approach by 40 and 10 percentage points respectively. The proposed approach also proves to guarantee the real EV charge needs throughout the simulations, and it is ready for operational use. Considering that we assume the EV departure time is known upfront, future work should envelop the accommodation of uncertainty of the EV departure time. For example, the chance-constrained method [135] is a robust approach to solve optimization problems by restricting the feasible region of the solutions to a certain level. The best-case scenario and worst-case scenario can then be determined in addition to the optimal solution.

6.6 Acknowledgments

This Chapter, in full, is still in an initial draft form for future publication with Dennis van der Meer and my PhD advisor Dr. Jan Kleissl. The dissertation author was the primary investigator and author of this manuscript.

Chapter 7

Concluding Remarks

The integration of solar energy into the electricity network is becoming essential because of the continually increasing growth in PV penetration. However, stochastic atmospheric dynamics make the solar forecasts vulnerable to errors, introducing significant challenges to the grid system. The preceding chapters characterize the solar forecast errors and discuss the ways to reduce the errors. Moreover, the impact of forecast errors on the grid planning and operations are quantified and proved to be mitigated in a grid net load valley filling problem.

Firstly, since cloud base height (CBH) is an important parameter for physics based solar radiation modeling. Two methods were developed to estimate accurate local CBHs by modeling 1) Geometric relation matching solar and cloud shadow edge geometry from sky imagery with the timing of ramp events from GHI time series and 2) Correlation between ground data GHI time series and simulated GHI time series from sky images projected to different CBH. The derived CBHs are benchmarked against a collocated ceilometer and stereographically estimated CBHs from two sky imagers. Over 30 days covering all seasons, the two methods perform similarly with an nRMSD of 18.9% versus 20.8%, which is slightly worse than the stereographic method at a much lower instrumentation cost.

In keeping with that, we further enhance our CSS to measure local cloud motion vectors more reliably and frequently in favor of short-term solar forecasting. Based on the improved cloud speed measurements, we demonstrate that CBH can be derived by triangulating cloud speed and cloud pixel speed measured by a sky imager. The proposed method was validated against two-months worth of measurements from meteorological aerodrome reports (METAR) from the nearby airport and an on-site ceilometer. Typical daily RMSD is

126 m which corresponds to 16.9% of the observed CBH. Normalized RMSD remains below 30% for all days. The daily bias is usually less than 80 m which suggests that the method is robust and that most of the RMSD are driven by short-term random fluctuations in CBH. It is demonstrated that our methods provide more accurate CBH estimates in higher temporal resolutions than the local airport's measured ceiling heights, which benefits the physics-based solar forecasting.

Furthermore, having the ability to estimate local cloud speed, we develop a model to estimate the worst case scenario power ramp rate for the optimal design of a solar power plant with storage, aiming to compensate the mismatch between imperfect solar forecasts and actual solar production. The novelty of this method is that only publicly available measurements are required and high-frequency measurements are not necessary. The ability of the proposed method to bracket actual ramp rates is assessed over 10 months under different meteorological conditions, exhibiting an average compliance rate of 98.9% for a 2 min evaluation time window. The largest observed ramp of 29.7% s^{-1} is contained with the worst case estimate of 34.3% s^{-1}. The remaining 1.1% of times can be primarily ascribed to inaccurate cloud velocity measurements in conditions with sparse and/or thin clouds.

No matter how accurate a solar forecast becomes, it will never be error-free. In addition to the forecast error magnitude, we showcase the structure of forecast errors is also an important determinant in the performance of solar grid integration. We first quantify the impact of solar forecast errors on a deterministic problem of grid net load valley filling through EV charging. The case study reveals that simple statistical models for the autocorrelated forecast errors seriously underestimate its negative effects. Then we formulate a corrective optimization framework to mitigate the net load variability resulted from the autocorrelated persistence forecast errors using sky imager forecasts. The proposed methodology was tested with one month of PV generation and substation load data. Compared to using day-ahead persistence solar forecasts, the proposed corrective framework delivers a 25% reduction in the daily net load variability on a 30-day average, and a 33% reduction in the 15 partly cloudy days. It is demonstrated that correcting the autocorrelated persistence forecast errors with more accurate short-term forecasts benefits grid planning.

Finally, we consider the operational use of the proposed corrective approach by including the uncertainty in EV charge events and PV forecasts in a stochastic optimization framework. The uncertainties in PV production and EV availability are first modeled by their respective probabilistic forecasts, and then are taken care by the scenarios in the stochastic programming scheme. We propose a novel way to figure out a truly optimal solution out of the stochastic objective function by setting up a joint constraint of charge power at the next time step for all simulated scenarios. Real-time updates of forecast and real EV charge events are also artfully incorporated to ensure the real world charge needs are satisfied, which represents the key merit of the method. Two common approaches are adopted for benchmarking: dumb charging approach and day-ahead averaging approach. Overall, with 50 scenarios, we achieve the net load variability of 10%, which is half the variability of typical day-ahead averaging approach and is one-fifth as small as the variability resulted from the dumb charging. More importantly our method fully satisfies the real charge needs and supports operational use. Regarding the compute cost, 50 scenarios can be solved in 3 min and minimal performance gain is observed beyond.

In summary, while it is very challenging to achieve perfect solar forecasts, the predictive control based approach offers a nice feature to only correct the forecast errors in the present time step, and leave the rest to be corrected in the next time steps when future updates are available. Therefore, a very accurate long term forecasts are not necessary for grid integration applications, since the error will not be fully carried through to the solved solutions. That leaves room for the needs of accurate short-term solar forecasts. The deviations of our solutions from the perfect solutions represent the cost of perfect information on future events, which is known as the expected value of the perfect information. Finally, to further minimize the remaining errors, granular data need to be included in the model, and a system level design optimization might be necessary. This thesis provides a guideline that will be beneficial for researchers and engineers who are working on the operational use of imperfect solar forecasts in the field of modeling and planning of solar grid integration.

Bibliography

[1] B. Jones-Albertus, Confronting the duck curve: How to address over-generation of solar energy, US Department of Energy.

[2] H. Yang, B. Kurtz, D. Nguyen, B. Urquhart, C. W. Chow, M. Ghonima, J. Kleissl, Solar irradiance forecasting using a ground-based sky imager developed at uc san diego, Solar Energy 103 (2014) 502–524.

[3] S. Steffel, P. Caroselli, A. Dinkel, J. Liu, R. Sackey, N. Vadhar, Integrating solar generation on the electric distribution grid, IEEE Transactions on Smart Grid 3 (2) (2012) 878–886.

[4] R. Shah, N. Mithulananthan, R. Bansal, V. K. Ramachandaramurthy, A review of key power system stability challenges for large-scale pv integration, Renewable and Sustainable Energy Reviews 41 (2015) 1423–1436.

[5] F. Ueckerdt, R. Brecha, G. Luderer, Analyzing major challenges of wind and solar variability in power systems, Renewable energy 81 (2015) 1–10.

[6] R. G. Wandhare, V. Agarwal, Reactive power capacity enhancement of a pv-grid system to increase pv penetration level in smart grid scenario, IEEE Transactions on Smart Grid 5 (4) (2014) 1845–1854.

[7] S. I. Nanou, A. G. Papakonstantinou, S. A. Papathanassiou, A generic model of two-stage grid-connected pv systems with primary frequency response and inertia emulation, Electric Power Systems Research 127 (2015) 186–196.

[8] S. Eftekharnejad, G. T. Heydt, V. Vittal, Optimal generation dispatch with high penetration of photovoltaic generation, IEEE Transactions on Sustainable Energy 6 (3) (2014) 1013–1020.

[9] S. Sobri, S. Koohi-Kamali, N. A. Rahim, Solar photovoltaic generation forecasting methods: A review, Energy Conversion and Management 156 (2018) 459–497.

[10] C. Wan, J. Zhao, Y. Song, Z. Xu, J. Lin, Z. Hu, Photovoltaic and solar power forecasting for smart grid energy management, CSEE Journal of Power and Energy Systems 1 (4) (2015) 38–46.

[11] C. ISO, What the duck curve tells us about managing a green grid, Calif. ISO, Shap. a Renewed Future (2012) 1–4.

[12] F. Mwasilu, J. J. Justo, E.-K. Kim, T. D. Do, J.-W. Jung, Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration, Renewable and sustainable energy reviews 34 (2014) 501–516.

[13] R.-C. Leou, C.-L. Su, C.-N. Lu, Stochastic analyses of electric vehicle charging impacts on distribution network, IEEE Transactions on Power Systems 29 (3) (2013) 1055–1063.

[14] S. Bashash, H. K. Fathy, Cost-optimal charging of plug-in hybrid electric vehicles under time-varying electricity price signals, IEEE Transactions on Intelligent Transportation Systems 15 (5) (2014) 1958–1968.

[15] M. Lave, M. J. Reno, R. J. Broderick, Characterizing local high-frequency solar variability and its impact to distribution studies, Solar Energy 118 (2015) 327–337.

[16] P. Haessig, B. Multon, H. B. Ahmed, S. Lascaud, P. Bondon, Energy storage sizing for wind power: impact of the autocorrelation of day-ahead forecast errors, Wind Energy 18 (1) (2015) 43–57.

[17] B. Urquhart, B. Kurtz, J. Kleissl, Sky camera geometric calibration using solar observations, Atmospheric Measurement Techniques 9 (9) (2016) 4279–4294.

[18] J. Wang, W. B. Rossow, Determination of cloud vertical structure from upper-air observations, Journal of Applied Meteorology 34 (10) (1995) 2243–2258.

[19] J. Gaumet, J. Heinrich, M. Cluzeau, P. Pierrard, J. Prieur, Cloud-base height measurements with a single-pulse erbium-glass laser ceilometer, Journal of atmospheric and oceanic technology 15 (1) (1998) 37–45.

[20] G. Martucci, C. Milroy, C. D. ODowd, Detection of cloud-base height using jenoptik chm15k and vaisala cl31 ceilometers, Journal of Atmospheric and Oceanic Technology 27 (2) (2010) 305–318.

[21] K. Hutchison, E. Wong, S. Ou, Cloud base heights retrieved during night-time conditions with modis data, International Journal of Remote Sensing 27 (14) (2006) 2847–2862.

[22] J. A. Shaw, P. W. Nugent, Physics principles in radiometric infrared imaging of clouds in the atmosphere, European Journal of Physics 34 (6) (2013) S111.

[23] L. Liu, X.-j. Sun, X.-c. Liu, T.-c. Gao, S.-j. Zhao, Comparison of cloud base height derived from a ground-based infrared cloud measurement and two ceilometers, Advances in Meteorology 2015.

[24] A. Dessler, S. Palm, J. Spinhirne, Tropical cloud-top height distributions revealed by the ice, cloud, and land elevation satellite (icesat)/geoscience laser altimeter system (glas), Journal of Geophysical Research: Atmospheres 111 (D12).

[25] A. Prata, P. Turner, Cloud-top height determination using atsr data, Remote sensing of environment 59 (1) (1997) 1–13.

[26] N. Killius, C. Prahl, N. Hanrieder, S. Wilbert, M. Schroedter-Homscheidt, On the use of nwp for cloud base height estimation in cloud camera-based solar irradiance nowcasting, ICEM 2015, 23.-26. June 2015, Boulder, USA.

[27] M. C. Allmen, W. P. Kegelmeyer Jr, The computation of cloud-base height from paired whole-sky imaging cameras, Journal of Atmospheric and Oceanic Technology 13 (1) (1996) 97–113.

[28] E. Kassianov, C. N. Long, J. Christy, Cloud-base-height estimation from paired ground-based hemispherical observations, Journal of Applied Meteorology 44 (8) (2005) 1221–1233.

[29] D. A. Nguyen, J. Kleissl, Stereographic methods for cloud base height determination using two sky imagers, Solar Energy 107 (2014) 495–509.

[30] J. Kleissl, B. Urquhart, M. Ghonima, E. Dahlin, A. Nguyen, B. Kurtz, C. Chow, F. Mejia, University of california, san diego (ucsd) sky imager cloud position study field campaign report, Tech. rep., DOE ARM Climate Research Facility, Washington, DC (United States) (2016).

[31] G. Wang, B. Kurtz, J. Kleissl, Cloud base height from sky imager and cloud speed sensor, Solar Energy 131 (2016) 208–221.

[32] P. Kuhn, M. Wirtz, S. Wilbert, J. Bosch, G. Wang, L. Ramirez, D. Heinemann, R. Pitz-Paal, Field validation and benchmarking of a cloud shadow speed sensor, Solar Energy 173 (2018) 229–245.

[33] P. Kuhn, M. Wirtz, N. Killius, S. Wilbert, J. L. Bosch, N. Hanrieder, B. Nouri, J. Kleissl, L. Ramírez, M. Schroedter-Homscheidt, et al., Benchmarking three low-cost, low-maintenance cloud height measurement systems and ecmwf cloud heights against a ceilometer, Solar Energy 168 (2018) 140–152.

[34] V. Fung, J. Bosch, S. Roberts, J. Kleissl, Cloud shadow speed sensor, Atmospheric Measurement Techniques 7 (6) (2014) 1693–1700.

[35] C. W. Chow, B. Urquhart, M. Lave, A. Dominguez, J. Kleissl, J. Shields, B. Washom, Intra-hour forecasting with a total sky imager at the uc san diego solar energy testbed, Solar Energy 85 (11) (2011) 2881–2893.

[36] T. Schmidt, J. Kalisch, E. Lorenz, D. Heinemann, Evaluating the spatio-temporal performance of sky-imager-based solar irradiance analysis and forecasts, Atmospheric chemistry and physics 16 (5) (2016) 3399–3412.

[37] B. Urquhart, M. Ghonima, D. Nguyen, B. Kurtz, C. W. Chow, J. Kleissl, Sky imaging systems for short-term forecasting, Solar Energy Forecasting and Resource Assessment (2013) 195–232.

[38] B. Urquhart, B. Kurtz, E. Dahlin, M. Ghonima, J. Shields, J. Kleissl, Development of a sky imaging system for short-term solar power forecasting, Atmospheric Measurement Techniques 8 (2) (2015) 875.

[39] C. W. Chow, S. Belongie, J. Kleissl, Cloud motion and stability estimation for intra-hour solar forecasting, Solar Energy 115 (2015) 645–655.

[40] M. Ghonima, B. Urquhart, C. Chow, J. Shields, A. Cazorla, J. Kleissl, A method for cloud detection and opacity classification based on ground based sky imagery, Atmospheric Measurement Techniques 5 (11) (2012) 2881–2892.

[41] A. C. Cabrera, Development of a Sky Imager for Cloud Classification and Aerosol Characterization: Ph. D. Dissertation, Editorial de la Universidad de Granada, 2010.

[42] R. Marquez, C. F. Coimbra, Intra-hour dni forecasting based on cloud tracking image analysis, Solar Energy 91 (2013) 327–336.

[43] J. L. Bosch, J. Kleissl, Cloud motion vectors from a network of ground sensors in a solar power plant, Solar Energy 95 (2013) 13–20.

[44] J. L. Bosch, Y. Zheng, J. Kleissl, Deriving cloud velocity from an array of solar radiation measurements, Solar Energy 87 (2013) 196–203.

[45] P. Ineichen, R. Perez, A new airmass independent formulation for the linke turbidity coefficient, Solar Energy 73 (3) (2002) 151–157.

[46] F. A. Mejia, B. Kurtz, A. Levis, Í. de la Parra, J. Kleissl, Cloud tomography applied to sky images: A virtual testbed, Solar Energy 176 (2018) 287–300.

[47] J. Bright, C. Smith, P. Taylor, R. Crook, Stochastic generation of synthetic minutely irradiance time series derived from mean hourly weather observation data, Solar Energy 115 (2015) 229–242.

[48] T. Schmidt, J. Kalisch, E. Lorenz, D. Heinemann, Evaluating the spatio-temporal performance of sky imager based solar irradiance analysis and forecasts., Atmospheric Chemistry & Physics 15 (19).

[49] V. Gevorgian, S. Booth, Review of prepa technical requirements for interconnecting wind and solar generation, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States) (2013).

[50] Y. Ru, J. Kleissl, S. Martinez, Exact sizing of battery capacity for photovoltaic systems, European Journal of Control 20 (1) (2014) 24–37.

[51] I. De la Parra, J. Marcos, M. García, L. Marroyo, Control strategies to use the minimum energy storage requirement for pv power ramp-rate control, Solar Energy 111 (2015) 332–343.

[52] J. Marcos Álvarez, Í. d. l. Parra Laita, M. García Solano, L. Marroyo Palomo, Control strategies to smooth short-term power fluctuations in large photovoltaic plants using battery storage systems, Energies, 2014, 7 (10), 6593-6619.

[53] J. Marcos, O. Storkël, L. Marroyo, M. Garcia, E. Lorenzo, Storage requirements for pv power ramp-rate control, Solar Energy 99 (2014) 28–35.

[54] R. Tonkoski, L. A. Lopes, T. H. El-Fouly, Coordinated active power curtailment of grid connected pv inverters for overvoltage prevention, IEEE Transactions on sustainable energy 2 (2) (2010) 139–147.

[55] S. Shivashankar, S. Mekhilef, H. Mokhlis, M. Karimi, Mitigating methods of power fluctuation of photovoltaic (pv) sources–a review, Renewable and Sustainable Energy Reviews 59 (2016) 1170–1184.

[56] A. Nguyen, M. Velay, J. Schoene, V. Zheglov, B. Kurtz, K. Murray, B. Torre, J. Kleissl, High pv penetration impacts on five local distribution networks using high resolution solar resource assessment with sky imager and quasi-steady state distribution system simulations, Solar Energy 132 (2016) 221–235.

[57] P. Kuhn, B. Nouri, S. Wilbert, C. Prahl, N. Kozonek, T. Schmidt, Z. Yasser, L. Ramirez, L. Zarzalejo, A. Meyer, et al., Validation of an all-sky imager–based nowcasting system for industrial pv plants, Progress in Photovoltaics: Research and Applications 26 (8) (2018) 608–621.

[58] M. Saleh, L. Meek, M. A. Masoum, M. Abshar, Battery-less short-term smoothing of photovoltaic generation using sky camera, IEEE Transactions on Industrial Informatics 14 (2) (2017) 403–414.

[59] R. Perez, S. Kivalov, J. Schlemmer, K. Hemker Jr, D. Renné, T. E. Hoff, Validation of short and medium term operational solar radiation forecasts in the us, Solar Energy 84 (12) (2010) 2161–2172.

[60] M. Lave, J. Kleissl, E. Arias-Castro, High-frequency irradiance fluctuations and geographic smoothing, Solar Energy 86 (8) (2012) 2190–2199.

[61] R. Perez, M. David, T. E. Hoff, M. Jamaly, S. Kivalov, J. Kleissl, P. Lauret, M. Perez, et al., Spatial and temporal variability of solar energy, Foundations and Trends® in Renewable Energy 1 (1) (2016) 1–44.

[62] S. Kuszamaul, A. Ellis, J. Stein, L. Johnson, Lanai high-density irradiance sensor network for characterizing solar resource variability of mw-scale pv system, in: 2010 35th IEEE Photovoltaic Specialists Conference, IEEE, 2010, pp. 000283–000288.

[63] J. Marcos, L. Marroyo, E. Lorenzo, D. Alvira, E. Izco, From irradiance to output power fluctuations: the pv plant as a low pass filter, Progress in Photovoltaics: Research and Applications 19 (5) (2011) 505–510.

[64] M. Lave, J. Kleissl, J. S. Stein, A wavelet-based variability model (wvm) for solar pv power plants, IEEE Transactions on Sustainable Energy 4 (2) (2012) 501–509.

[65] M. Lave, J. Kleissl, Cloud speed impact on solar variability scaling–application to the wavelet variability model, Solar Energy 91 (2013) 11–21.

[66] J. Marcos, Í. de la Parra, M. García, L. Marroyo, Simulating the variability of dispersed large pv plants, Progress in Photovoltaics: Research and Applications 24 (5) (2016) 680–691.

[67] R. van Haaren, M. Morjaria, V. Fthenakis, An energy storage algorithm for ramp rate control of utility scale pv (photovoltaics) plants, Energy 91 (2015) 894–902.

[68] I. de la Parra, J. Marcos, M. García, L. Marroyo, Dealing with the implementation of ramp-rate control strategies–challenges and solutions to enable pv plants with energy storage systems to operate correctly, Solar Energy 169 (2018) 242–248.

[69] M. Lave, A. Weekley, Comparison of high-frequency solar irradiance: Ground measured vs. satellite-derived, in: 2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC), IEEE, 2016, pp. 1101–1106.

[70] G. Ari, Y. Baghzouz, Impact of high pv penetration on voltage regulation in electrical distribution systems, in: 2011 International Conference on Clean Electrical Power (ICCEP), IEEE, 2011, pp. 744–748.

[71] R. Hanna, J. Kleissl, A. Nottrott, M. Ferry, Energy dispatch schedule optimization for demand charge reduction using a photovoltaic-battery storage system with solar forecasting, Solar Energy 103 (2014) 269–287.

[72] M. Delfanti, D. Falabretti, M. Merlo, Energy storage for pv power plant dispatching, Renewable Energy 80 (2015) 61–72.

[73] S. Mazzola, C. Vergara, M. Astolfi, V. Li, I. Perez-Arriaga, E. Macchi, Assessing the value of forecast-based dispatch in the operation of off-grid rural microgrids, Renewable Energy 108 (2017) 116–125.

[74] A.-L. Klingler, L. Teichtmann, Impacts of a forecast-based operation strategy for grid-connected pv storage systems on profitability and the energy system, Solar Energy 158 (2017) 861–868.

[75] P. Coker, J. Barlow, T. Cockerill, D. Shipworth, Measuring significant variability characteristics: An assessment of three uk renewables, Renewable energy 53 (2013) 111–120.

[76] M. Q. Raza, M. Nadarajah, C. Ekanayake, On recent advances in pv output power forecast, Solar Energy 136 (2016) 125–144.

[77] N. Zhang, C. Kang, Q. Xia, J. Liang, Modeling conditional forecast error for wind power in generation scheduling, IEEE Transactions on Power Systems 29 (3) (2013) 1316–1324.

[78] V. Mohan, J. G. Singh, W. Ongsakul, An efficient two stage stochastic optimal energy and reserve management in a microgrid, Applied Energy 160 (2015) 28–38.

[79] J. Lujano-Rojas, G. Osório, J. Matias, J. Catalão, A heuristic methodology to economic dispatch problem incorporating renewable power forecasting error and system reliability, Renewable Energy 87 (2016) 731–743.

[80] M. Marzband, M. Javadi, J. L. Domínguez-García, M. M. Moghaddam, Non-cooperative game theory based energy management systems for energy district in the retail market considering der uncertainties, IET Generation, Transmission & Distribution 10 (12) (2016) 2999–3009.

[81] Y. Zhang, T. Zhang, R. Wang, Y. Liu, B. Guo, Optimal operation of a smart residential microgrid based on model predictive control by considering uncertainties and storage impacts, Solar Energy 122 (2015) 1052–1065.

[82] E. Scolari, F. Sossan, M. Paolone, Irradiance prediction intervals for pv stochastic generation in microgrid applications, Solar Energy 139 (2016) 116–129.

[83] W. Wu, K. Wang, B. Han, G. Li, X. Jiang, M. L. Crow, A versatile probability model of photovoltaic generation using pair copula construction, IEEE Transactions on Sustainable Energy 6 (4) (2015) 1337–1345.

[84] B. Ngoko, H. Sugihara, T. Funaki, Synthetic generation of high temporal resolution solar radiation data using markov models, Solar Energy 103 (2014) 160–170.

[85] M. Marzband, N. Parhizi, J. Adabi, Optimal energy management for stand-alone microgrids based on multi-period imperialist competition algorithm considering uncertainties: experimental validation, International transactions on electrical energy systems 26 (6) (2016) 1358–1372.

[86] A. Michiorri, J. Lugaro, N. Siebert, R. Girard, G. Kariniotakis, Storage sizing for grid connected hybrid wind and storage power plants taking into account forecast errors autocorrelation, Renewable energy 117 (2018) 380–392.

[87] E. Nuño, M. Koivisto, N. A. Cutululis, P. Sørensen, On the simulation of aggregated solar pv forecast errors, IEEE Transactions on Sustainable Energy 9 (4) (2018) 1889–1898.

[88] E. Nuno, M. Koivisto, N. Cutululis, et al., Simulation of regional day-ahead pv power forecast scenarios, in: 2017 IEEE Manchester PowerTech, IEEE, 2017, pp. 1–6.

[89] Y. Ru, J. Kleissl, S. Martinez, Storage size determination for grid-connected photovoltaic systems, IEEE Transactions on sustainable energy 4 (1) (2012) 68–81.

[90] A. Nottrott, J. Kleissl, B. Washom, Energy dispatch schedule optimization and cost benefit analysis for grid-connected, photovoltaic-battery storage systems, Renewable Energy 55 (2013) 230–240.

[91] S. Chen, H. B. Gooi, M. Wang, Sizing of energy storage for microgrids, IEEE Transactions on smart grid 3 (1) (2011) 142–151.

[92] A. H. Habib, V. R. Disfani, J. Kleissl, R. A. de Callafon, Quasi-dynamic load and battery sizing and scheduling for stand-alone solar system using mixed-integer linear programming, in: 2016 IEEE Conference on Control Applications (CCA), IEEE, 2016, pp. 1476–1481.

[93] M. Marzband, A. Sumper, J. L. Domínguez-García, R. Gumara-Ferret, Experimental validation of a real time energy management system for microgrids in islanded mode using a local day-ahead electricity market and minlp, Energy Conversion and Management 76 (2013) 314–322.

[94] E. L. Ratnam, S. R. Weller, C. M. Kellett, An optimization-based approach to scheduling residential battery storage with solar pv: Assessing customer benefit, Renewable Energy 75 (2015) 123–134.

[95] S. Habib, M. Kamran, U. Rashid, Impact analysis of vehicle-to-grid technology and charging strategies of electric vehicles on distribution networks–a review, journal of Power Sources 277 (2015) 205–214.

[96] Z. Wang, S. Wang, Grid power peak shaving and valley filling using vehicle-to-grid systems, IEEE Transactions on power delivery 28 (3) (2013) 1822–1829.

[97] P. Denholm, M. O'Connell, G. Brinkman, J. Jorgenson, Overgeneration from solar energy in california. a field guide to the duck chart, Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States) (2015).

[98] S. Sae, Electric vehicle and plug in hybrid electric vehicle conductive charge coupler (2010).

[99] S. Stüdli, E. Crisostomi, R. Middleton, R. Shorten, Optimal real-time distributed v2g and g2v management of electric vehicles, International Journal of Control 87 (6) (2014) 1153–1162.

[100] T. Wu, Q. Yang, Z. Bao, W. Yan, Coordinated energy dispatching in microgrid with wind power generation and plug-in electric vehicles, IEEE Transactions on Smart Grid 4 (3) (2013) 1453–1463.

[101] S. J. Kang, J. Park, K.-Y. Oh, J. G. Noh, H. Park, Scheduling-based real time energy flow control strategy for building energy management system, Energy and Buildings 75 (2014) 239–248.

[102] O. Sundstrom, C. Binding, Flexible charging optimization for electric vehicles considering distribution grid constraints, IEEE Transactions on Smart grid 3 (1) (2011) 26–37.

[103] J. Mattingley, Y. Wang, S. Boyd, Receding horizon control, IEEE Control Systems Magazine 31 (3) (2011) 52–65.

[104] P. Braun, L. Grüne, C. M. Kellett, S. R. Weller, K. Worthmann, Model predictive control of residential energy systems using energy storage and controllable loads, in: European Consortium for Mathematics in Industry, Springer, 2014, pp. 617–623.

[105] T. Wang, H. Kamath, S. Willard, Control and optimization of grid-tied photovoltaic storage systems using model predictive control, IEEE Transactions on Smart Grid 5 (2) (2014) 1010–1017.

[106] D. Wang, J. Coignard, T. Zeng, C. Zhang, S. Saxena, Quantifying electric vehicle battery degradation from driving vs. vehicle-to-grid services, Journal of Power Sources 332 (2016) 193–203.

[107] M. Ropp, J. Newmiller, C. Whitaker, B. Norris, Review of potential problems and utility concerns arising from high penetration levels of photovoltaics in distribution systems, in: 2008 33rd IEEE Photovoltaic Specialists Conference, IEEE, 2008, pp. 1–6.

[108] M. Y. Suberu, M. W. Mustafa, N. Bashir, Energy storage systems for renewable energy power sector integration and mitigation of intermittency, Renewable and Sustainable Energy Reviews 35 (2014) 499–514.

[109] M. Katsanevakis, R. A. Stewart, J. Lu, Aggregated applications and benefits of energy storage systems with application-specific control methods: A review, Renewable and Sustainable Energy Reviews 75 (2017) 719–741.

[110] G. C. Wang, E. Ratnam, H. V. Haghi, J. Kleissl, Corrective receding horizon ev charge scheduling using short-term solar forecasting, Renewable energy 130 (2019) 1146–1158.

[111] J. Donadee, M. D. Ilić, Stochastic optimization of grid to vehicle frequency regulation capacity bids, IEEE Transactions on Smart Grid 5 (2) (2014) 1061–1069.

[112] R.-C. Leou, J.-H. Teng, H.-J. Lu, B.-R. Lan, H.-T. Chen, T.-Y. Hsieh, C.-L. Su, Stochastic analysis of electric transportation charging impacts on power quality of distribution systems, IET Generation, Transmission & Distribution 12 (11) (2018) 2725–2734.

[113] W. Su, J. Wang, J. Roh, Stochastic energy scheduling in microgrids with intermittent renewable energy resources, IEEE Transactions on Smart grid 5 (4) (2013) 1876–1883.

[114] M. Alizadeh, A. Scaglione, J. Davies, K. S. Kurani, A scalable stochastic model for the electricity demand of electric and plug-in hybrid vehicles, IEEE Transactions on Smart Grid 5 (2) (2013) 848–860.

[115] L. Jian, Y. Zheng, X. Xiao, C. Chan, Optimal scheduling for vehicle-to-grid operation with stochastic connection of plug-in electric vehicles to smart grid, Applied Energy 146 (2015) 150–161.

[116] P. Grahn, K. Alvehag, L. Söder, Phev utilization model considering type-of-trip and recharging flexibility, IEEE Transactions on Smart Grid 5 (1) (2013) 139–148.

[117] B. Wang, Y. Wang, H. Nazaripouya, C. Qiu, C.-C. Chu, R. Gadh, Predictive scheduling framework for electric vehicles with uncertainties of user behaviors, IEEE Internet of Things Journal 4 (1) (2016) 52–63.

[118] H. V. Haghi, S. Lotfifard, Z. Qu, Multivariate predictive analytics of wind power data for robust control of energy storage, IEEE Transactions on Industrial Informatics 12 (4) (2016) 1350–1360.

[119] M. Neaimeh, R. Wardle, A. M. Jenkins, J. Yi, G. Hill, P. F. Lyons, Y. Hübner, P. T. Blythe, P. C. Taylor, A probabilistic approach to combining smart meter and electric vehicle charging data to investigate distribution network impacts, Applied Energy 157 (2015) 688–698.

[120] H. Kamankesh, V. G. Agelidis, A. Kavousi-Fard, Optimal scheduling of renewable micro-grids considering plug-in hybrid electric vehicle charging demand, Energy 100 (2016) 285–297.

[121] J. R. Birge, F. Louveaux, Introduction to stochastic programming, Springer Science & Business Media, 2011.

[122] W.-K. Mak, D. P. Morton, R. K. Wood, Monte carlo bounding techniques for determining solution quality in stochastic programs, Operations research letters 24 (1-2) (1999) 47–56.

[123] A. J. Kleywegt, A. Shapiro, T. Homem-de Mello, The sample average approximation method for stochastic discrete optimization, SIAM Journal on Optimization 12 (2) (2002) 479–502.

[124] Y. Wang, B. Wang, C.-C. Chu, H. Pota, R. Gadh, Energy management for a commercial building microgrid with stationary and mobile battery storage, Energy and Buildings 116 (2016) 141–150.

[125] G. Wang, V. Disfani, J. Kleissl, Scenario based stochastic optimization of probabilistic ev charging scheduling, in: 2018 IEEE Innovative Smart Grid Technologies-Asia (ISGT Asia), IEEE, 2018, pp. 552–557.

[126] E. Xydas, C. Marmaras, L. M. Cipcigan, N. Jenkins, S. Carroll, M. Barker, A data-driven approach for characterising the charging demand of electric vehicles: A uk case study, Applied energy 162 (2016) 763–771.

[127] R. L. Thorndike, Who belongs in the family?, Psychometrika 18 (4) (1953) 267–276.

[128] L. Breiman, Random forests, Machine learning 45 (1) (2001) 5–32.

[129] J. Friedman, T. Hastie, R. Tibshirani, The elements of statistical learning, Vol. 1, Springer series in statistics New York, 2001.

[130] N. Meinshausen, Quantile regression forests, Journal of Machine Learning Research 7 (Jun) (2006) 983–999.

[131] S. Siegert, Specsverification: forecast verification routines for ensemble forecasts of weather and climate. r package version 0.5–2. 2017.

[132] J.-F. Toubeau, J. Bottieau, F. Vallée, Z. De Grève, Deep learning-based multivariate probabilistic forecasting for short-term scheduling in power markets, IEEE Transactions on Power Systems 34 (2) (2018) 1203–1215.

[133] R. B. Nelsen, An introduction to copulas, Springer Science & Business Media, 2007.

[134] J. Munkhammar, J. Widén, An autocorrelation-based copula model for generating realistic clear-sky index time-series, Solar Energy 158 (2017) 9–19.

[135] Y. Huang, L. Wang, W. Guo, Q. Kang, Q. Wu, Chance constrained optimization in a home energy management system, IEEE Transactions on Smart Grid 9 (1) (2016) 252–260.