Lawrence Berkeley National Laboratory

Recent Work

Title

Open Computing Infrastructure for Sharing Data Analytics to Support Building Energy Simulations

Permalink

<https://escholarship.org/uc/item/20v6f3rt>

Journal

Journal of Computing in Civil Engineering, 33(6)

ISSN

0887-3801

Authors

Karaguzel, OT
Elshambakey, M
Zhu, Y
et al.

Publication Date

2019-11-01

DOI

10.1061/(ASCE)CP.1943-5487.0000857

Peer reviewed

OPEN COMPUTING INFRASTRUCTURE FOR SHARING DATA ANALYTICS TO SUPPORT BUILDING ENERGY SIMULATIONS

Omer T. Karaguzel¹; Mohammed Elshambakey²; Yimin Zhu^{3*}; Tianzhen Hong⁴; William J. Tolone⁵;
Sreyasee Das Bhattacharjee⁶; Isaac Cho⁷; Wenwen Dou⁸; Haopeng Wang⁹; Siliang Lu¹⁰; and Yong Tao¹¹

Abstract: Building energy simulation plays an increasingly important role in building design and operation. In this paper, we present an open computing infrastructure, Virtual Information Fabric Infrastructure (VIFI), that allows building designers and engineers to enhance their simulations by combining empirical data with diagnostic or prognostic models. Based on the idea of dynamic data-driven application systems (DDDAS), the VIFI infrastructure complements conventional data-centric sharing strategies and addresses key data sharing concerns such as the privacy of building occupants. To demonstrate the potential of the VIFI infrastructure, we simulate an empirically-derived lighting schedule in the U.S. Department of Energy’s small office building reference model. We use the case study simulation to explore the possibility and potential of integrating data-centric and analytic-centric sharing strategies; the method of combining empirical data with simulations; the creation, sharing, and execution of analytics using VIFI; and the impact of incorporating empirical data on energy simulations. While the case study reveals clear advantages of the VIFI data infrastructure, research questions remain surrounding the motivation and benefits for sharing data, the metadata that are required to support the composition of analytics, and the performance metrics that could be used in assessing the applications of VIFI.

Keywords: Data sharing; Building energy, Simulation, Data standards, EnergyPlus, Operational schedules, Energy simulations.

¹ Assistant Professor, School of Architecture, Carnegie Mellon University, 5000 Forbes Avenue, MMCH #415, Pittsburgh, PA, 15213, Email: okaraguz@andrew.cmu.edu

² Post-Doctoral Research Associate, College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, Email: melshamb@uncc.edu

³ Professor, Department of Construction Management, Louisiana State University, Baton Rouge, LA 70803 70803 (Corresponding Author). Email: yiminzhu@lsu.edu, Corresponding Author

⁴ Staff Scientist, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, Email: thong@lbl.gov

⁵ Professor, College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, Email: william.tolone@uncc.edu

⁶ Post-Doctoral Research Associate, College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, Email: sdasbhat@uncc.edu

⁷ Research Assistant Professor, College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, Email: ichol@uncc.edu

⁸ Assistant Professor, College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte, NC 28223, Email: wdoul@uncc.edu

⁹ Ph.D. student, School of Architecture, Carnegie Mellon University, 5000 Forbes Avenue, MMCH #415, Pittsburgh, PA, 15213, Email: wanghp18@gmail.com

¹⁰ Ph.D. student, School of Architecture, Carnegie Mellon University, 5000 Forbes Avenue, MMCH #415, Pittsburgh, PA, 15213, Email: siliang1@andrew.cmu.edu

¹¹ Professor, 3301 College Avenue, Nova Southeastern University, Ft. Lauderdale, FL 33314, Email: ytao@nova.edu

INTRODUCTION

Buildings account for more than one-third of the primary energy consumption in the world; therefore, reducing energy use and greenhouse-gas emissions in the building sector is a key strategy for achieving

global energy and environmental goals. Building performance simulation has played an increasingly important role in designing, operating, and retrofitting buildings to predict energy efficiency and utility costs (Roth 2017). As more energy policies and regulations push new buildings or building retrofits towards achieve low- or zero-net energy (ZNE) goals, building energy simulations can be used as tools for evaluating and comparing design alternatives that employ integrated building technologies and control strategies. ZNE buildings also require on-site renewable energy generation, and balancing demand-side building energy efficiency and on-site generation with the supply-side renewable power requires an optimization based on various performance metrics, including energy use, energy cost, return on investment (ROI), or greenhouse gas (GHG) emissions. Moreover, energy models developed in the building design phase can be adapted for use later in the building operation phase to support performance diagnostics and operation improvements, along with guiding building retrofit strategies.

Building performance simulations rely on diagnostic and/or prognostic models to understand and predict building performance. Often, these models require significant empirical data inputs on building operations, physical and environmental conditions, or occupants and their behaviors (D'Oca and Hong 2015, Kwak et al. 2015). In addition, model calibration with empirical data is critical to the quality of building performance simulations (Coakley et al. 2014). However, gaining access to empirical data, particularly monthly or time-interval energy data, can be challenging due to the data privacy concerns of building owners. Thus, on one hand the incorporation of empirical data into building performance simulations improves simulation quality; on the other hand, it presents a data privacy challenge.

In this paper, we introduce an open computing infrastructure, or virtual information fabric infrastructure (VIFI), that integrates predictive models with empirical data to enhance building energy simulations. VIFI shares analytics on raw data, complementing conventional data sharing strategies that require exposing the raw data directly. The framework is designed based on the concept of dynamic data driven application systems (DDDAS) (Darema 2005). In essence, DDDAS include two parts: computation and measurement. Computation refers to predictive models or simulations for a particular purpose such as building energy simulations, while measurement refers to instrumentation for data

collection such as building sensors or data archives. The computation piece takes the advantage of measurement data to dynamically calibrate simulations or predictions, and provides feedback to the measurement piece to inform future data collection efforts. The feedback loop between computation and measurement allows continuous improvement in simulations and predictions.

Two aspects of VIFI are particularly significant. First, the VIFI computing infrastructure offers a new data sharing alternative in cases where a conventional, data-centric sharing strategy is not applicable. For example, advancements in building sensor and control technologies has significantly improved the ability of building owners and managers to collect building performance data, potentially offering an unprecedented opportunity to improve building energy simulations as they apply to the design of new buildings, building retrofits, or operational strategies. Yet, access to these big data by building designers and engineers is highly constrained due to data privacy issues. In particular, building owners or facility managers may not feel comfortable sharing the data with others because such raw data often contain sensitive occupant information. In addition, the raw data may not be in the correct form for others to use. Nevertheless, data owners may support the general idea of collaborating with others to improve the quality of building energy simulations. In such scenarios, the VIFI data sharing infrastructure can alleviate the data owner's concerns about privacy and support their ability to collaborate with others without exposing their raw data.

Consequently, the VIFI data sharing infrastructure potentially allows building designers and engineers to access a larger pool of building data than would be possible using a conventional direct sharing strategy alone, which leads to another significant contribution of the infrastructure: VIFI allows building designers and engineers to enhance their simulations by combining high-quality empirical data with diagnostic and/or prognostic models for simulations. Energy simulations often have many input parameters that are strongly tied to building design or operational conditions. Such relevance significantly affects the quality of simulations. Improved relevance can be achieved by offering building designers and engineers access to a large set of relevant building data, which the new strategy can offer. Therefore, the infrastructure has the potential to change how building performance simulations are done in the future.

In the following sections, we discuss existing studies that are relevant to building energy data integration and sharing, outline the proposed VIFI infrastructure for building energy simulations, present a case study to demonstrate the implementation and use of the infrastructure, and summarize key conclusions and areas for future research.

RELATED STUDIES

The need to incorporate empirical data in building energy simulations has been extensively explored in the existing literature. Building energy simulations typically require several input variables, many of which are related to human behavior, such as an occupancy or lighting schedule (Crawley et al. 2001). Traditionally, building energy simulations are performed in a closed manner with all inputs fixed or assumed to be constant during a simulation (e.g., Zhu (2006), Pan et al. (2007)). To improve the accuracy of simulations, dynamic approaches that incorporate real-time variable inputs have recently gained attention amongst researchers. For example, (Pang et al. 2016) presented a method that integrates real-time data such as weather, plug load power, lighting power, occupancy, and room temperature settings with whole building energy simulation. Similarly, (Kwak et al. 2015) discussed real-time simulation through co-simulation to optimize building operations. However, most such applications have a focus on building operations, not building design.

While simulations using real-time data during the design stage are rare, the idea of enhancing design-stage simulations through co-simulation with other programs has been previously explored. For example, Feng et al. (2015) and later Chen et al. (2018) developed an occupancy simulation software tool to work with building energy simulation through co-simulation. Software platforms have also been developed for deploying co-simulations, such as the tight-coupling approach using the building controls virtual test bed (BCVTB) (Wetter 2011) or the loose-coupling approach based on the occupant behavior XML (obXML) (Hong et al. 2016). Additionally, OpenStudio SDK provides an ecosystem of software tools that allows a comprehensive suite of modeled variables such as lighting and occupancy to be analyzed in whole building energy simulations through an open application programming interface

(Guglielmetti et al. 2011). However, the deployment of the OpenStudio SDK requires energy models and associated data to be stored on an OpenStudio server for model execution and calibration. Thus, the potential of the proposed VIFI strategy, i.e., moving analytics to data and thus alleviating data privacy concerns, is not already explored in the OpenStudio SDK. Clearly, whole building energy simulation trends demonstrate the need for simulations that address finer component levels. However, even co-simulations cannot fully address the issue of a simulation's relevance to a particular design because the parameters of co-simulations may be tuned to data from dissimilar buildings. Thus, connections must be established between buildings with high relevance to a particular design and design-stage simulations in order to further enhance the robustness of such simulations.

Dynamic data driven application systems (DDDAS) offer a method to establish the aforementioned connections between buildings and simulations. Since their inception, DDDAS have been applied in scenarios as diverse as emergency and disaster management (Madey et al. 2007), supply chain systems (Celik et al. 2010), and threat management for urban water distribution systems (Wang et al. 2014). Potential DDDAS applications to building performance simulations have been identified as well. For example, Spitler (2006) discussed the use of model-based control systems as an application of DDDAS and the need to explore methods of data-gathering processes under a DDDAS framework. Bouktif and Ahmed (2015) discussed the application of DDDAS to the energy consumption of a residential system, where the application includes real-time metering and sensor data and incorporates human decisions into energy analytics to enhance predictions. In current building performance applications, DDDAS fits well for applications such as real-time model-based simulations for optimizing building operations (Pang et al. 2016).

While these previous applications of DDDAS in buildings concern daily operations, applications of DDDAS to building designs are not yet reported in the existing literature. Application of DDDAS during the design of new buildings or building retrofits is different from its application to operations, in particular because the input data for design simulations often cannot be collected *in-situ*. In addition, simulations during the design stage often cannot be calibrated using empirical data from buildings that are

exact proxies for the building under design (Lomas et al. 1997). Thus, the best strategy is to offer maximum access to building datasets, such that building designers and engineers are more likely to find relevant data from existing buildings for model input or calibration. As discussed, the VIFI infrastructure provides such a possibility by offering a complementary approach to the existing data-centric sharing strategy, increasing access to shared building performance data.

In summary, existing studies concerning building energy simulation call for component level co-simulations to better address the need for different types of simulations. However, these component level simulations still rely on data that may not be representative of a particular building under design. Consequently, better access to building performance data is needed. The increasing availability of big building operation datasets coupled with the DDDAS framework offer the possibility of a new data sharing approach that allows building designers and engineers to seek those data that best match a given design or operation optimization plan and integrate these data with associated building simulations.

OPEN COMPUTING INFRASTRUCTURE FOR SHARING ANALYTICS

In this section, we present a high-level conceptual architecture for the VIFI open computing infrastructure. As the topics of simulation models and data collection from existing buildings have already been discussed extensively and are not the major focus of this work, we focus only on the application of VIFI to the integration of simulations with empirical data.

High-Level Conceptual Architecture

Fig. 1 demonstrates the high-level conceptual architecture of the VIFI infrastructure. There are three key components in the infrastructure: predictive models (such as simulations), existing buildings, and the VIFI platform. The predictive model component represents any predictive model that may be applied during building design process, such as a whole building energy simulation. The existing building component represents sources of empirical data. Data may be drawn from real-time sensor readings or from other stored logs of building operations, occupant behaviors, indoor or outdoor environmental

conditions, and other variables relevant to building research or energy applications. Finally, the VIFI platform represents a computing platform that coordinates the feedback loop between the first two components.

This computing infrastructure extends conventional application of DDDAS by inserting a generic computing component, the VIFI platform, between computation and instrumentation, the two key components of DDDAS. This extension potentially allows a wider adoption of DDDAS and better support of data sharing. In the following section, the key components of this VIFI platform are detailed.

Virtual Information Fabric Infrastructure (VIFI) Platform

VIFI is an open-source platform that can run on any operating system. While the final VIFI version will be conveyed on an open-source vault, we use the Amazon Elastic Compute Cloud (Amazon EC2) (Amazon Inc. 2018) for the initial implementation, testing, and assessment of the platform. In essence, VIFI brings analytics to locations that contain a large amount of data. VIFI allows users to access data already claimed by others; thus, users of VIFI have the capacity to perform a variety of analyses – integrating data from multiple locations, for example. VIFI permits virtual sharing of various sorts of testbeds, such as labs that create trial data, without requiring that raw datasets be moved that may be too large or sensitive to share outside the control of the primary data owners. The current implementation of VIFI consists of the following main components (Fig. 2 and Fig. 3):

- **Portable Analytic Container (PAC):** A PAC is a lightweight virtual machine, called a container, that hosts software (e.g., EnergyPlus (EnergyPlus 2018)), libraries (e.g., Eppy Python module to communicate with Energyplus (Santosh 2018)), and any other tools or operating systems required by end users to analyze data. A PAC can receive and execute end user analysis programs if the required programs are not already contained in the PAC. Leveraging container technology (e.g., Docker (Miell and Sayers 2016)), a PAC is portable and can migrate and execute on heterogeneous host platforms. A PAC facilitates reusability by hosting and utilizing different analytical libraries and programs pulled from shared

repositories (e.g., Docker hub (DockerHub 2018)). Container technology enables the movement of analytics rather than the movement of data, thus alleviating problems related to the transfer of big data – e.g., download times and/or security requirements. Each PAC is a lightweight virtual machine implemented as a Docker image (DockerHub 2018) and containing the required libraries and support programs that are required to execute the desired analysis. PACs offer a number of affordances for distributed analytics: i) they can be easily transmitted over the network due to their limited size; and ii) they simplify analytics development for inexperienced clients. The VIFI infrastructure, as explained by Elshambakey et al. (2018), is also scalable, i.e., enabling the integration of various VIFI nodes (e.g., a BIM server). The ability for VIFI workflows to access fixed VIFI nodes allows VIFI to cooperate with non-open-source resources, assuming that a VIFI user has the proper credentials.

- **Registry Service:** Distinctive PACs are stored, searched, utilized and shared through a Registry Service. In the current usage, a Docker hub (DockerHub 2018) is utilized to implement the Registry Service. Future VIFI upgrades will incorporate expansions of appropriated registry repositories to advance download and transfer time of PACs.
- **Orchestrator:** The Orchestrator provides automatic correspondence between distinctive workflow segments and the coordination of infrastructure at various locales. Every segment can be as basic as a solitary procedure running on a host, or the host itself. In the current VIFI implementation, NIFI (The APACHE software foundation 2018) functions as the orchestrator, coordinating the execution of workflow segments across the VIFI infrastructure. NIFI is an open-source tool for automating and managing the flow of information between systems. Workflow segments can dwell on the same host or distributed hosts through NIFI site-to-site communication (The APACHE software foundation 2018). The VIFI orchestrator includes loosely coupled NIFI interaction and coordination between destinations that empower fault-tolerance and scalability with an increasing number of data-stores as well as end-users.

- **User Node:** The User Node is the passageway for a client to interface with the VIFI framework. The User Node provides a user interface (UI), communication, and basic computation capacities. The current VIFI implementation uses NIFI at every User Node to empower correspondence with other workflow parts, as well as workflow inception. Clients can plan, alter, and submit systematic contents as well as other required sources of info to VIFI through a NIFI web interface.
- **Data Site:** Data Sites are locations in the VIFI infrastructure that have varying types of heterogeneous information from different sources and models. Each VIFI-empowered information site uses NIFI and Docker Swarm (Vohra 2017). While NIFI empowers appropriated coordination and correspondence between every data site and other parts of the workflow, Docker Swarm executes clients' contents utilizing determined PACs as Docker orchestrators. Consequently, Docker Swarm conveys a group of services at every data site for parallel investigation execution without expecting clients to have any previous knowledge of the framework, stages or situations at every datum site.
- **Metadata Server:** The Metadata Server stores and lists gathered data and comments about datasets and records, which can be created by utilizing a set of enlisted extractors or physically included by clients and information proprietors. For movability, extractors are containerized as PACs. A Metadata Server supports dataset search and discovery across different sites as shown in Fig. 3.
- **Crawler:** The Crawler provides data crawling through extractors that send extracted data to the Metadata Server.
- **Watchdog:** The Watchdog updates the Metadata Server for any modifications of the registered distributed data sets.

Fig. 2 outlines the current VIFI implementation (Fig. 2(A)) and compares this implementation against conventional data-centric sharing approaches (Fig. 2(B)). Under conventional data-centric sharing approaches (Fig. 2(A)), users must first download data locally and then run required analytics. Under the

VIFI implementation (Fig.2(B)), users submit required analytics (including a Docker image and scripts) to data locations and then retrieve post-processed results. The example in Fig. 2 shows the execution of multiple requests simultaneously in both cases.

This paper focuses on the data privacy benefits of VIFI more than the benefits of VIFI in handling large datasets; nevertheless, it is noted that VIFI was previously evaluated for an earth science application where observational ($\cong 11$ GB) data and model data ($\cong 21$ GB) were analyzed and compared to evaluate the fitness of the model (Elshambakey et al. 2017). In this setting, it was observed that VIFI consumed far less processing time than a conventional data-centric approach. This is because under VIFI, Docker images need be downloaded only once and then again as updates are available. The data update rate at each data-site under the conventional approach is expected to be much higher than the update rate for Docker images; accordingly, conventional users must download updated data far more often than they would need to update Docker images under VIFI approach. An illustration of the comparative time savings from VIFI execution is found in Elshambakey et al. (2017).

CASE STUDY

We use a case study to demonstrate the application of VIFI for integration of empirical data with whole building energy simulation in the early building design phase. The case study generates dynamic lighting schedules using data processing algorithms within a Portable Analytic Container (PAC). Algorithms extract the necessary information from empirical building data, collected from the Intelligent Workspace at Carnegie Mellon University. The Intelligent Workplace at Carnegie Mellon University is a living and lived-in research laboratory dedicated to prototyping and testing of integrative high-performance building technologies such as shading integrated solar photovoltaics, automated operable windows, automated adjustable external shading devices, and daylight controlled dimmable electric lights. Building data (including occupancy presence) are collected from these systems through an extensive wired and wireless sensor network and data acquisition systems.

Measured building operation data from the Intelligent Workspace are utilized remotely by a building energy simulation tool, EnergyPlus, via the VIFI platform. The VIFI platform facilitates the operation of an analysis procedure that converts the collected empirical data into hourly building lighting schedules. Without moving any raw empirical data from its original location, VIFI executes analytics that transform the raw data collected by sensors in the Intelligent Workspace every five minutes into an hourly lighting schedule in a specific format that is readable by the remote EnergyPlus simulation engine. VIFI supports the integration of this empirically-derived lighting schedule with EnergyPlus as well as the execution of annual energy performance simulations that reflect the schedule.

Description of the Building Energy Model

We use the small-sized commercial reference office building provided by the U.S. Department of Energy Building Technologies Office (Deru et al. 2011) for our case study simulations. This model has been developed in the format of EnergyPlus version 8.7, which has been extensively utilized for computational evaluations of integrated building energy performance, mechanical and electrical systems, and thermal interactions of building spaces with increased spatial and temporal granularity (Field et al. 2010). The specific small-sized office model is selected from a comprehensive database containing 16 different hypothetical reference building definitions, which represent about 70% of the new commercial building stock in the U.S. (DOE-EERE 2018).

The DOE small office reference model is a single-story commercial office with a total conditioned and usable floor area of 511m². The building has a rectangular shape and an aspect ratio of 1.5, with its long axis on a north-south orientation (Fig. 4). Punched windows are used to represent vertical fenestration configurations, which are uniformly distributed to each cardinal orientation with a window-to-wall ratio (WWR) of 21.2% (the total window area is 59.6 m²). Thermal zone layouts are formed using the perimeter-core zoning approach, resulting in five thermal zones and an unconditioned and unoccupied attic space. The building envelope is composed of thermally massive and insulated construction assemblies that comply with the minimum requirements imposed by

ANSI/ASHRAE/IESNA Energy Standard 90.1 for non-residential building types (ASHRAE 2004). The thermo-physical properties of the main building envelope assemblies are listed in Table 1. The model is equipped with a double-pane glazing system defined in a simplified manner with U-values, solar heat gain coefficient (SHGC), and V_T inputs of 3.241, 0.385, 0.305, respectively (including the thermal and optical effects of frames and dividers). Internal loading conditions pertaining to artificial lights, electrical appliances, and occupancy are shown in Table 2 together with outdoor ventilation rates, infiltration rates, and thermostatic control limits for indoors.

The HVAC system type modeled for the building is a group of packaged single zone air conditioners (PSZ-AC) with single-speed direct expansion (DX) coils and compact gas furnaces for heating defined for each of the five thermal zones. There is no central air conditioning unit (AHU) and no economizers in this configuration. System fans are of the constant volume type without variable speed drivers (Table 3). The reference office model is simulated under the environmental boundary conditions for the location of Pittsburgh, PA, USA. A detailed statistical climate summary is shown in Table 4.

VIFI Implementation of the Case Study

Fig. 5 shows the VIFI implementation of the case study. Although the case study was implemented on the Amazon EC2, the authors emulated a distributed environment as in actual applications, including a building information model (BIM) server, a user node, a data server, and an energy server. The BIM server contains the BIM of the case building. The user node represents the location where building designers and engineers work. The data server represents the access point to empirical data, where raw lighting schedule data are stored in this case. The energy server is where EnergyPlus is located.

The current capability of VIFI can be extended by incorporating fixed VIFI nodes (e.g., the BIM server and the EnergyPlus server) where analytics do not have to run inside containers (Elshambakey et al. 2017). The ability of VIFI workflows to access fixed VIFI nodes allows VIFI to cooperate with non-open-source resources given that a VIFI user has the proper credentials. The services provided by the BIM server, the data server(s), and the energy server are registered with VIFI so that a user can determine

what operations or analytics are available for their applications. In the case study, the following operations and analytics are registered with VIFI first in the PAC registry (Fig. 5):

1. the BIM server (fixed VIFI node): accepts requests to create an input data file (IDF) based on the BIM model for energy simulations and returns an IDF;
2. the data server: creates and returns hourly lighting schedules based on the raw data; and
3. the energy server (fixed VIFI node): accepts an IDF and returns energy simulation results.

The following is the workflow of the VIFI implementation of the case study:

1. Steps 1 and 2: a building designer or engineer contacts the BIM server (i.e., Step 1) to get an EnergyPlus Input Data File (IDF) (i.e., Step 2), which contains default schedule settings. At this point, some or all parts of the IDF need to be updated based on proper data analysis (e.g., lighting schedule or occupancy schedule). In this case study, the BIM server was accessed directly to simplify the implementation. In future implementations, the BIM server can be implemented as a fixed VIFI node that does not require a PAC to access. Sharing a BIM directly with all partners is a conventional data-centric strategy; therefore, the building designer or engineer can access BIM on a shared network. The integration of BIM with VIFI represents an extension to the conventional data sharing strategy. The new strategy (i.e., sharing analytics) and the conventional strategy (i.e., sharing data directly) are complementary to each other and can work hand-in-hand to solve practical problems.
2. Steps 3, 4, and 5: These steps create an IDF that better reflects design or operation conditions. In Step 3, the building designer or engineer searches for lighting schedules that best match the design or operation conditions that he or she expects. The metadata server of VIFI (Fig. 3) helps identify relevant data sources that are accessible through VIFI. In the case study, the authors assume that the collected data source from the Intelligent Workspace represents the best match.

Once the data source or sources are identified, the building designer or engineer sends a request to the data source(s) to generate the required schedule in Step 4. In the case of energy

simulations, the operations needed to develop better schedules are implemented as standard analytics in VIFI. The analytics are small-sized applications that run inside Docker containers (Miell and Sayers 2016) using Docker Swarm (Vohra 2017). Docker containers are produced using Docker images hosted at different VIFI registries (e.g., Docker hub (DockerHub 2018)). Docker images contain all required libraries/dependencies to run the user's analytics. Thus, end users such as a building designer or engineer do not have to worry about configuration of the required analytics.

The analytics are executed on data at the data source(s). In the case study, the raw data are collected by sensors every five minutes continuously over multiple years at the Intelligent Workspace. The raw data cannot be used by others directly for different reasons: 1) the occupancy or lighting schedule requires hourly data but the raw data are collected every five minutes; 2) the data may be in different format and schemas (e.g., csv, gbXML (GreenBuildingXML 2018), obXML (Hong et al. 2015a; Hong et al. 2015b)) from what is needed for creating the IDF; 3) the raw data may contain data that are irrelevant to the analysis; and 4) the data are too sensitive to share in their raw format as they may compromise occupant privacy. Thus, instead of downloading data from the Intelligent Workspace to the building designer's or engineer's local site, the data owner (Carnegie Mellon University (CMU) in this case) can publish the analytics that are allowed to operate on its data sources to VIFI and the building designer or engineer can find the analytics at the PAC registry. The building designer or engineer then uses the analytics to access data and obtain the desired schedule updates. Importantly, this approach does not require a user to have significant knowledge about programming, so the person can stay focused on the main simulation tasks.

Once the building designer or engineer receives the updated hourly lighting schedule, the original IDF is updated with this new information in Step 5.

3. Steps 6 and 7: Finally, the updated IDF file is used to run energy simulations using the EnergyPlus server (Cao et al. 2011) in Step 6. In Step 7, the building designer or engineer receives simulation results.

Creating, Sharing, and Executing Analytics

Currently VIFI supports two mechanisms for creating, sharing, and executing analyses. The first mechanism uses Docker containers to encapsulate the analytics. This mechanism supports the reusability of different analysis applications and building new analytics from existing ones. The second mechanism uses a fixed point VIFI node to host a specific software such as a BIM or the EnergyPlus server. Fixed point VIFI nodes are useful in cases where VIFI must interact with software that is difficult to containerize (e.g., proprietary software without available Docker images).

Fig. 6 shows the sample application of the first mechanism to create a lighting schedule for energy simulation. The authors used lighting schedules as an example in the case study to demonstrate creating, sharing, and executing analytics - lighting schedules are one of many input schedules to EnergyPlus. Analytics take the form of a dedicated Python script that generates a lighting schedule in IDF (Input Definition File) format using the empirical data collected at CMU. The input to this analytics script is a file in CSV format with two columns: the first column contains timestamps, and the second column indicates the artificial light status, either on or off as recorded by sensors. The output is a whole year *Schedule:Compact* IDF object. Between the input and the output, the analytics perform data parsing and extraction using the Python Pandas library. The script is generic and can thus be shared among different application cases to create improved lighting schedules.

The implementation of the second mechanism is straightforward. The Eppy Python module (Santosh 2018) was used to integrate non-open-source software such as EnergyPlus with VIFI. In the case study, the Linux-based EnergyPlus Server, version 8.7.0 (EnergyPlus 2018) is used. The EnergyPlus Server runs the energy simulation based on the updated IDF file and returns the energy use of the sample building.

Simulations and Results

To demonstrate the usefulness of VIFI platform in improving building energy simulations from the point of occupancy representativeness, the authors performed a series of simulation runs using comparable models. A baseline simulation model representing common energy modeling practices uses a standard lighting schedule provided by the Building Technologies Office of the U.S. Department of Energy. Other simulation models use empirically-derived lighting schedules derived based on three occupancy profiles, i.e., a professor, a graduate student, and a postdoctoral researcher. The standard lighting schedule lacks relevance to actual design or operational conditions, while empirically-derived schedules hold certain relevance to specific design or operation conditions. The main objective of this demonstration and subsequent comparative analyses is to reveal the impact of using empirically-derived schedules on the predictive capability of whole-building energy models. In turn, the impact shows the importance of using VIFI to improve building energy simulations.

Fig. 7 clearly indicates that empirically-derived electric lighting schedules pertaining to different occupancy types do not match well with standardized schedules (shown in dashed lines), which lack dynamism within a typical weekday. The results also signify the need to use design-specific or building case-specific operational schedules to enhance the predictive capability of building energy models. Meanwhile, empirically-derived lighting schedules also show some degree of variations among themselves, which is due to highly varied space use styles for the selected occupant types. Compared to the graduate student (with 1,633 hrs of occupancy), the user types of professor and post-doctoral researcher spend much more time in their offices with occupied hours of 3,094 hrs and 2,676 hrs, respectively. However, all occupancy time is shorter than the assumption of standard schedules, 5,097 hrs in a typical year.

Given fractional hourly schedules, which are drawn from the outputs from EnergyPlus' annual energy simulation runs, the authors calculated the cumulative power fraction (CPF) for electric lighting systems, defined as the sum of the number of hours in a year when systems are operating in full power, and conducted the comparisons shown in Table 5. The CPF metric together with its derivative of Full

Power Equivalent ($FPE = CPF/8,760$) metric revealed that while occupancy types of professor (Type I) and postdoc (Type III) had very close CPFs compared to the standardized schedule, there existed a considerable deviation for the graduate student (Type II) with a 45% change from the standard CPF, i.e., 2,863. This is due to the fact that Type I and Type III have similar peak power levels to those of the standard schedule while Type II shows variations in peak power, its frequency, and event timing.

It is clear that with similar CPF and FPE metrics, the annual lighting energy consumption of Type I (professor) is very close to the model alternative using a standard schedule (only -0.7% deviation observed between the two). Type III (post-doc) occupancy type has an CPF and FPE of 2,437 and 0.28 respectively for lights, which results in about 14.9% decrease in annual lighting energy consumption compared to the standardized schedule model. Simulation results for Type II (grad student) show significant departures from the standard schedule model with deviations of up to -45.7% for lighting energy consumption.

To demonstrate the dynamic effects of alternative hourly lighting schedules on cumulative energy performance, we performed annual energy simulations facilitated by the VIFI platform. Annual site energy consumption levels for the end uses of space heating, cooling and electric lighting (kWh) are extracted from annual EnergyPlus simulation runs for comparison purposes (Fig. 8). It should be noted that since the lights in a building energy model contribute to internal heat gains, deviations in the predictions of the energy models are also found in space heating and cooling energy consumptions. For example, results in Fig 8 indicate that due to reduced usage of lights in the Type II (grad student) model, annual space heating energy was increased by 16%, while annual space cooling energy was reduced by 17% compared to the standard schedule model. While such increments and decrements appear to balance each other out in the particular climate zone (Pittsburgh, PA), where heating dominates, such effects would have important implications for cooling energy savings estimates in cooling-dominated climates.

Discussion

Our case study demonstrates three key attributes of the VIFI computing infrastructure. First, the case study shows the technical feasibility and application potentials of integrating data-centric and analytic-centric sharing strategies. In the building industry, data-centric sharing has been the norm and technologies such as building information models are vehicles allowing stakeholders to directly share data created by others. The case study shows that VIFI complements the conventional approaches by utilizing both data and analytics for sharing purposes. This capability may allow building designers and engineers to access more data than what conventional methods such as building information models can provide. Second, the case study shows the VIFI platform enables the use of empirically-driven operational building schedules for energy performance simulations instead of relying on pre-determined and generic schedules, which have no apparent contextual and occupational relevance to building design and operation characteristics. Such an approach can significantly affect the predictive ability and representativeness of energy performance models and can be a critical element of model calibration and validation procedures. Third, the case study has demonstrated the technical feasibility of the creation, sharing, and execution of analytics that can be executed across distributed sites (e.g., the BIM server, the data server, and the energy server). This capability is at the core of the VIFI platform and its proposed computing infrastructure.

Furthermore, the case study demonstrates that the VIFI platform can be instrumental to the automation of specific labor-intensive energy modeling tasks (e.g., repeated development of annual hourly schedules for different building spaces with list lengths as high as 8,760 elements). Such automation not only reduces human errors but also speeds up the process of model development and interpretation of results. Additionally, the VIFI computing infrastructure can be instrumental to model-based building control applications, which carry building energy models from the design stage through to the operational stage in order to perform predictions across a specific time horizon, allowing optimal control of building mechanical and electrical systems for increased energy efficiency and occupant comfort.

CONCLUSIONS AND FUTURE RESEARCH

This study has demonstrated the potential of the VIFI open computing infrastructure to offer a novel, analytics-based data sharing strategy that complements conventional data-centric strategies. By sharing analytics on data, the new strategy points towards a future where empirical data from existing building operations are closely connected with decision-making processes that drive building design and operation through additional cyberinfrastructure. The study has demonstrated many possibilities for using this novel infrastructure to enhance building performance simulations. For example, the study shows the potential for the VIFI platform to support both data and analytic-centric sharing. This new strategy bears significant implications for the building design and engineering community, where data fragmentation has been recognized as a major challenge to data sharing. As shown in the case study, designers and engineers can access both a building information model and empirical data in an integrated manner through VIFI. This strategy enables designers and engineers to access a larger pool of data that are not only more relevant to improving their simulations, but may also set the foundation for future applications of advanced methods such as machine learning to transfer knowledge that is gained during building operations back to new building designs.

This study only examined the functional capabilities of the open computing infrastructure, and many other aspects related to the infrastructure require further research. For example, our case study was implemented with the help of computer scientists who are technical experts in deploying and implementing applications using VIFI. Looking ahead, it is critical that the VIFI platform is useful to those in design and engineering domains who do not have the technical background to understand technical details of the proposed infrastructure such as the PAC and registry. Therefore, a user-friendly workflow process is important to develop for future VIFI implementations. In addition, while the VIFI platform does not need data owners to share data *per se*, a certain level of access to data needs to be granted to outside users in order for the VIFI analytics to work on the datasets. This requirement goes beyond the issue of addressing data privacy and security issues, bringing forth more fundamental questions about the motivation and benefits of sharing data. Rich datasets are needed for access by VIFI

analytics in order to realize the full benefits of the proposed infrastructure, and the incentives for creating and sharing such datasets must be improved. For example, “information as commodity” (Smith 1983) provides an interesting idea to drive data creators and users to share data. Such an idea can be pilot-tested in a small community, such as those working on the development and initial applications of VIFI. Experimentation with this idea may shed more light on how to increase the sharing of data, which is important to other VIFI functions. Another important issue to be explored is the quality of metadata, which plays a significant role in searching for, reusing, and executing analytics on the VIFI platform. Although metadata are not addressed in this case study, this topic has been discussed for decades in the buildings research community and as a result many metadata schemas are available, for example: gbXML, Industry Foundation Classes (IFC), obXML, and CityGML. The VIFI research community must develop a strategy for building from such existing metadata schemas and integrating new concepts as needed to support VIFI applications. Finally, there is a need to develop a set of performance metrics to measure the benefits and risks of using the VIFI infrastructure; these metrics should include technical, economic, and social measures at the individual, company, and society levels.

ACKNOWLEDGEMENTS

This paper was partially supported by the U.S. National Science Foundation - Data Infrastructure Building Blocks (DIBBS) program’s funding for the project “Virtual Information Fabric Infrastructure (VIFI) for Data-Driven Decisions from Non-Shareable Distributed Data” (NSF DIBBs Award #1640818). This work was also supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the U.S. Department of Energy.

REFERENCES

- Amazon Inc. (2018). “Amazon Elastic Compute Cloud (Amazon EC2).” <http://aws.amazon.com/ec2/#pricing> (Jun. 28, 2018).
- ASHRAE. (2004). *ANSI/ASHRAE Standard 90.1-2004: Energy standard for buildings except low-rise residential buildings*. Atlanta: American Society of Heating, Refrigeration and Air Conditioning Engineers. Atlanta, GA.
- Bouktif, S., and Ahmed, W. K. (2015). “Monitoring System for Residential Energy Consumption.” *Int. J. New Comput. Archit. their Appl.*, 5(4), 156–164.
- Cao, Y., Li, H., Dong, W., Tang, Z., and Zhang, G. (2011). “Energy-saving performance analysis of distributional building energy system.” *Asia-Pacific Power Energy Eng. Conf. APPEEC*, IEEE, Wuhan, China, 1–4.
- Celik, N., Lee, S., Vasudevan, K., and Son, Y.-J. (2010). “DDDAS-based multi-fidelity simulation framework for supply chain systems.” *IIE Trans.*, 42(5), 325–341.
- Chen, Y., Hong, T., and Luo, X. (2018). “An agent-based stochastic Occupancy Simulator.” *Build. Simul.*, 11(1), 37–49.
- Coakley, D., Raftery, P., and Keane, M. (2014). “A review of methods to match building energy simulation models to measured data.” *Renew. Sustain. Energy Rev.*, Elsevier, 37, 123–141.
- Crawley, D. B., Lawrie, L. K., Winkelmann, F. C., Buhl, W. F., Huang, Y. J., Pedersen, C. O., Strand, R. K., Liesen, R. J., Fisher, D. E., Witte, M. J., and Glazer, J. (2001). “EnergyPlus: Creating a new-generation building energy simulation program.” *Energy Build.*, 33(4), 319–331.
- D’Oca, S., and Hong, T. (2015). “Occupancy schedules learning process through a data mining framework.” *Energy Build.*, Elsevier B.V., 88, 395–408.
- Darema, F. (2005). “Dynamic Data Driven Applications Systems: New Capabilities for Application Simulations and Measurements.” *Int. Conf. Comput. Sci.*, Springer, Berlin, Heidelberg, 610–615.
- Deru, M., Field, K., Studer, D., Benne, K., Griffith, B., Torcellini, P., Liu, B., Halverson, M., Winiarski, D., Rosenberg, M., Yazdanian, M., Huang, J., and Crawley, D. (2011). “U.S. Department of Energy

- commercial reference building models of the national building stock.” 1–118.
- DockerHub. (2018). “Docker Hub.” <<https://hub.docker.com>> (Jun. 28, 2018).
- Elshambakey, M., Khalefa, M., Tolone, W. J., Das Bhattacharjee, S., Lee, H., Cinquini, L., Schlueter, S., Cho, I., Dou, W., and Crichton, D. J. (2017). “Towards a distributed infrastructure for data-driven discoveries & analysis.” *2017 IEEE SmartWorld, Ubiquitous Intell. Comput. Adv. Trust. Comput. Scalable Comput. Commun. Cloud Big Data Comput. Internet People Smart City Innov.*, 1–8.
- EnergyPlus. (2018). “EnergyPlus Linux Installation.” <<https://energyplus.net/installation-linux>> (Jun. 28, 2018).
- Feng, X., Yan, D., and Hong, T. (2015). “Simulation of occupancy in buildings.” *Energy Build.*, Elsevier B.V., 87, 348–359.
- Field, K., Deru, M., and Studer, D. (2010). “USING DOE COMMERCIAL REFERENCE BUILDINGS FOR SIMULATION STUDIES National Renewable Energy Laboratory , Golden , CO SUMMARY OF INPUTS AND.” *Proc. SimBuild*, New York City, New York, 85–93.
- GreenBuildingXML. (2018). “Green Building XML (gbXML) Schema.” <<http://www.gbxml.org>> (Jun. 28, 2018).
- Guglielmetti, R., Macumber, D., and Long, N. (2011). “OpenStudio: an open source integrated analysis platform.” *Proc. Build. Simul. 2011*, (December), 442–449.
- Hong, T., D’Oca, S., Taylor-Lange, S. C., Turner, W. J. N., Chen, Y., and Corgnati, S. P. (2015a). “An ontology to represent energy-related occupant behavior in buildings. Part II: Implementation of the DNAS framework using an XML schema.” *Build. Environ.*, Elsevier Ltd, 94, 196–205.
- Hong, T., D’Oca, S., Turner, W. J. N., and Taylor-lange, S. C. (2015b). “An Ontology to Represent Energy- related Occupant Behavior in Buildings Part I: Introduction to the DNAs Framework.” *Build. Environ.*, 92, 764–777.
- Hong, T., Sun, H., Chen, Y., Taylor-Lange, S. C., and Yan, D. (2016). “An occupant behavior modeling tool for co-simulation.” *Energy Build.*, Elsevier B.V., 117, 272–281.
- Kwak, Y., Huh, J. H., and Jang, C. (2015). “Development of a model predictive control framework

- through real-time building energy management system data.” *Appl. Energy*, Elsevier Ltd, 155, 1–13.
- Lomas, K. J., Eppel, H., Martin, C. J., and Bloomfield, D. P. (1997). “Empirical validation of building energy simulation programs.” *Energy Build.*, 26(3), 253–275.
- Madey, G. R., Barabási, A.-L., Chawla, N. V., Gonzalez, M., Hachen, D., Lantz, B., Pawling, A., Schoenharl, T., Szabó, G., Wang, P., and Yan, P. (2007). “Enhanced Situational Awareness: Application of DDDAS Concepts to Emergency and Disaster Management.” *Int. Conf. Comput. Sci.*, Springer, Berlin, Heidelberg, 1090–1097.
- Miell, I., and Sayers, A. H. (2016). *Docker in practice*. Computing and Computers.
- Pan, Y., Huang, Z., and Wu, G. (2007). “Calibrated building energy simulation and its application in a high-rise commercial building in Shanghai.” *Energy Build.*, 39(6), 651–657.
- Pang, X., Nouidui, T. S., Wetter, M., Fuller, D., Liao, A., and Haves, P. (2016). “Building energy simulation in real time through an open standard interface.” *Energy Build.*, Elsevier B.V., 117, 282–289.
- Roth, A. (2017). “Building Energy Modeling 101: What Is It and What Is DOE’s Role?” <<https://www.energy.gov/eere/buildings/articles/building-energy-modeling-101-what-it-and-what-doe-s-role>> (Jun. 22, 2018).
- Santosh, P. (2018). “eppy 0.5.44 documentation.” <<https://pythonhosted.org/eppy/index.html#>> (Jun. 28, 2018).
- Smith, K. A. (1983). “Information as A Commodity or Public Good.” *Proc. 25th Annu. Conf. Natl. Fed. Abstr. Inf. Serv.*, Philadelphia, 27–31.
- Spitler, J. D. (2006). “Building performance simulation: the now and the not yet.” *HVAC&R Res.*, 12(3a), 711–713.
- The APACHE software foundation. (2018). “APACHE nifi.” <<https://nifi.apache.org/index.html>> (Jun. 28, 2018).
- Vohra, D. (2017). *Docker Management Design Patterns: Swarm Mode on Amazon Web Services*. Apress.
- Wang, L., Chen, D., Liu, W., Ma, Y., Wu, Y., and Deng, Z. (2014). “DDDAS-based parallel simulation

- of threat management for urban water distribution systems.” *Comput. Sci. Eng.*, 16(1), 8–17.
- Wetter, M. (2011). “Co-simulation of building energy and control systems with the building controls virtual test bed.” *J. Build. Perform. Simul.*, 4(3), 185–203.
- Zhu, Y. (2006). “Applying computer-based simulation to energy auditing: A case study.” *Energy Build.*, 38(5), 421–428.