

Operational Characteristics of Residential Air Conditioners with Temporally Granular Remote Thermographic Imaging

Pandarasamy Arjunan

Berkeley Education Alliance for Research in Singapore Limited. samy@bears-berkeley.sg

Clayton Miller, Filip Biljecki

National University of Singapore, Singapore {clayton, filip}@nus.edu.sg

Gregory Dobler, Kyungmin Lee

University of Delaware, USA {gdobler, kmlee}@udel.edu

Kameshwar Poolla

University of California, Berkeley, USA poolla@berkeley.edu

ABSTRACT

Collecting accurate information about the operational characteristics of buildings is essential for many performance analyses including energy auditing and benchmarking. While most of the existing solutions involve manual inspection of building premises, surveys, or *in-situ* sensor deployment, these techniques are intrusive and/or time-intensive, limiting their applicability at larger, city-level spatial scales. Complementing these solutions, we present a wide-area infrared thermography (IRT) based system for measuring a building's operational characteristics remotely in a non-intrusive and scalable way, and we describe the image capture, processing pipeline, and preliminary results from an initial deployment of this system focusing on the operational characteristics of air-conditioning units. The data collection consisted of infrared imaging of a residential building for 50 days at continuous 10-second intervals. The infrared brightness variations of exterior split air-conditioners fixtures were extracted for each image, and operational attributes were then extracted from the resultant time series. Using state-based change point detection methods to determine times at which air-conditioners are operational, our preliminary analysis focuses on phenomenological patterns of activity with two main findings. First, we demonstrate our ability to determine the operational characteristics of air-conditioners and in particular the fact that they exhibit two distinct modes: continuous operation and cycling behavior. Second, we find that the fraction of "on" time spent in the cycling mode is characteristically longer for lower external temperatures, consistent with the hypothesis that the cycling mode represents a reaching of set temperature. Finally, we outline future research directions and challenges in leveraging IRT for behavioral studies of cooling system use and proactive assessment of air-conditioners towards the development of scalable virtual energy auditing.

CCS CONCEPTS

• Computer systems organization → Sensor networks; • Applied computing → Physics.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

BuildSys '21, November 17–18, 2021, Coimbra, Portugal

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-9114-6/21/11...\$15.00

<https://doi.org/10.1145/3486611.3486659>




Figure 1: *Left:* A visible wavelength image of the exterior view of the building under study in Singapore. *Right:* The infrared image of the same building with annotated AC units.

KEYWORDS

Infrared thermography, Residential Air Conditioners, Building characterization, Energy efficiency, Energy auditing, User behavior

ACM Reference Format:

Pandarasamy Arjunan, Gregory Dobler, Kyungmin Lee, Clayton Miller, Filip Biljecki, and Kameshwar Poolla. 2021. Operational Characteristics of Residential Air Conditioners with Temporally Granular Remote Thermographic Imaging. In *ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys '21)*, November 17–18, 2021, Coimbra, Portugal. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3486611.3486659>

1 INTRODUCTION

Building energy benchmarking involves quantifying a building's performance relative to its peers based on various metrics including total consumption and efficiency, and it requires the collection of fine-grained information about various operational attributes, such as air-conditioner (AC) usage and efficiency, of a number of buildings to form the comparison [2]. Building performance analysis is a related concept that refers to the quantification of the overall performance of a building during its operation. There are various methods available for both building energy characterization and performance analysis depending on the extent of such analyses, with one commonly used technique being to diagnose specific faults resulting in energy inefficient systems and operations [7].

Infrared thermography (IRT) has been widely used in various building energy diagnostics applications [8]. Broadly, there are three major application categories: detecting thermal leakages, tracking indoor mobility and assessing occupants' thermal comfort, and monitoring equipment health, degradation, and energy usage. Many of these applications involve manual inspection and analysis using handheld infrared cameras. Several studies have focused on air-conditioning systems in particular because they account for more

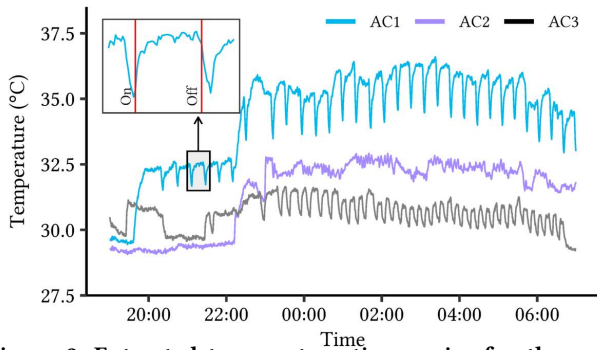


Figure 2: Extracted temperature time series for three AC units annotated in Figure 1 for one night (7pm - 7am). “Off” to “on” transitions are clearly visible as are rapid on/off “cycling” behavior. Note that AC1 and AC3 have variation in their initial “on” times and cycling behavior whereas AC2 has no cycling mode. The inset shows our annotations of starting time of “on” and “off” states for one cycle.

than 60% of the buildings’ total energy use [9], and in geographic locations with consistently warm temperatures (e.g., near the equator), space-cooling is required throughout the year [9]. Therefore, monitoring of AC usage and performance is essential to improve building energy efficiency.

Complementing these studies, in this paper, we follow the operational mode of the “Urban Observatory” (UO) facility [3] by deploying a wide-area IR camera mounted in a tall building at stand-off distance from a target building and acquiring imaging at sub-minute temporal resolution over roughly two months. Analogous to the studies on patterns of lighting variability in [4], we extract time-dependent brightnesses of the “sources” shown in Figure 1 and use characteristics of their variability to quantify patterns of operations and usage. This methodology is complimentary to conventional methods of estimates of air-conditioner end use via surveys or dense sensor deployments in that it is empirically measurement based with the potential to scale up to many buildings with a single sensor.

2 WIDE-AREA INFRARED IMAGING SYSTEM

The building under study shown in Figure 1 was located at a distance of ~ 200 meters from the IR camera and has 15 floors with approximately two apartment units on each floor. Our current observational system consists of a FLIR A300 thermal imaging camera weatherized to safeguard the camera from harsh conditions as well as a computing device for data acquisition. The camera itself supports multiple infrared resolutions up to 320×240 pixels and its thermal accuracy is $\pm 2^\circ\text{C}$ (though the conversion between IR signal and temperature by the FLIR acquisition software relies on several assumptions regarding the emissivity of the observed surfaces).

For continuous acquisition of the images at fixed intervals, we developed an image acquisition software based on the FLIR software development kit (SDK). As in the case of the UO [3], the collected thermal images are stored locally into the computing device and uploaded to a cloud storage server in real-time. We configured our data acquisition software to collect IR images every 10-seconds. Our observational campaign covered the period from 25 January to 14




Figure 3: Extracted temperature time series over 7 nights (7pm - 7am) for two ACs in our sample, one with both “continuous on” and cycling modes (top) and one without cycling (bottom). The yellowier times (higher temperature) denote periods when the AC was “on”.

March 2020 (50 days in total) with continuous observations resulting in approximately 432,000 IR images in total. In addition, we also collected the environmental data from a nearby weather station at a temporal granularity of 1-minute for the entire image collection period. While we note that this fixed vantage point provides a correspondingly fixed view of the urban scene as opposed to, for example, drone-based platforms for which multiple views may be possible, our observational methodology has the advantage of providing continuous coverage over the long temporal baselines of our study with no gaps or associated data transfer issues.

It is important to consider that proximal remote sensing techniques of urban environments at fine spatial and temporal granularity require great care to protect the privacy of the inhabitants in the scene. In this work, we follow the methodology outlined in [3] by limiting the spatial resolution of our images (the pixel resolution and distance corresponds to a pixel size of 0.27m at the surface of the building, or roughly 15 pixels for a 1m^2 patch) and not associating any of the derived variability with location in the image. Given the low resolution, single images by themselves do not contain personally identifiable information, but it is an open question regarding whether patterns of usage are personally identifiable. We discuss this possibility in Section 4, but note here that by removing the location of the sources from the analysis, we break the association between a time series and specific unit in the building.

3 OBSERVATIONS AND RESULTS

Our goal in this work is to demonstrate the ability to measure time-dependent IR brightness from “sources” in the scene that vary over time, similarly to the lightscape analysis from [4], but the sources in our case being the exterior vents of the air-conditioner (AC) units. Source detection in thermal images is a challenging task due to several factors including resolution, noise floor, and weather, and so in this proof-of-concept work, we select sources via hand-selection (the use of unsupervised and supervised machine learning methods for source detection will be the subject of future work; see Section 4). For each source, we use a bounding box of 10×10 pixels and extract the maximum “temperature” within each box for each image resulting in a continuous time series of IR brightness for each hand-selected source at a cadence of 0.1 Hz.




Figure 4: On/Off state detection result for a single AC for one night (8pm - 6am) using the univariate k -Means clustering.

3.1 Air-conditioner IR time-dependence

The time series for three of the AC units are shown in Figure 2. It is clear from the figure that – as [3] found for Heating Ventilation And Cooling vents at a distance of 1 mile – the IR imaging is capable of detecting transitions in IR brightness of the AC units that relate to the duty cycle and “on”/“off” state of the unit. Split AC units consist of an indoor unit and an outdoor unit, and it is the outdoor unit that contains the compressor, condenser coil, expansion coil, and fan that increase in temperature when the unit is actively cooling in an “on” state. That state is controlled by a controller circuit that coordinates the overall operation of these units including their duty cycle based on the unit’s set-point temperature. We note that it is the activity of the controller that is manifesting as an exterior observable. The figure also qualitatively shows that there are two distinct modes of operation, a continuously “on” state, and a “cycling” state in which the unit rapidly turns on and off. We hypothesize that this phenomenology represents the unit bringing the room to temperature during the continuously on state and keeping the room at temperature during the cycling state. Figure 3 shows similar time series but over seven full nights for two AC with qualitatively different operational phenomenology: one with ubiquitous cycling behavior and another with almost no cycling behavior. Furthermore, Figures 2 and 3 demonstrate that these ACs have significant variability in their initial “on” time for each of the seven days. The AC cycling is an operational characteristic of the unit itself (See [1] for a study of the dependence of cycling frequency and outdoor temperature on AC performance) while the initial “on” time as well as the total “on” time for the AC units likely encode user behavior [6].

3.2 Air-conditioner operational state detection

In order to quantitatively assess characteristics of the AC units for correlative analysis with both timings and external variables like air temperature [5], we first extract the “on”/“off” state of a unit at a given time. To that end, we applied two unsupervised state detection algorithms:

(a) **Hidden Markov Model (HMM):** HMM are a family of probabilistic models for which the probability of an observation B at time t_i is given only by the “state” of the underlying, unobservable Markovian process to which B is related at t_i . In our case, the probability of the AC being in an on or off state at t_i is given only by the AC temperature at t_i . We assume a Gaussian distribution of the observed AC temperature for both states of our two-state HMM model. Our HMM model is applied in two steps. In the first step, a HMM is fit to the full time series generating on/off states at each

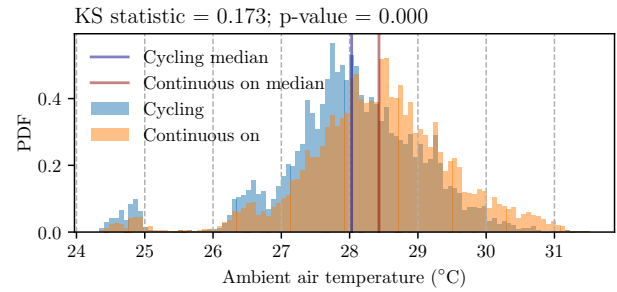


Figure 5: The distributions of air temperature at times when ACs are in a cycling mode vs a continuous on mode. Air temperatures are characteristically higher during continuous on modes, consistent with the hypothesis that the cycling mode corresponds to ACs maintaining temperature rather than cooling. The high two sample KS-statistic and low p -value indicate a very low probability that these distributions are random draws from the same parent distribution.

time. Given the difference in amplitude between fully off and off during a cycling mode, the single-pass HMM often mischaracterizes states, and so we then fit a further HMM to each sub-sequence of continuous “on” modes to extract any potential cycling behavior.

(b) **Univariate k -Means clustering:** This algorithm clusters the given univariate data into k clusters by minimizing the weighted sum of within-cluster distances. We set $k = 2$ which selects the optimal decision boundary between on/off states (clusters) in the AC temperature time series.

We use the `depmixS4` and `Ckmeans.1d.dp` packages in R to implement these two algorithms. We apply these two models to the time series of the difference between the AC “temperature” and the ambient air temperature for each AC for each night between 8pm and 6am. We take this difference to eliminate a slowly varying drift in the IR brightness of the AC units due to ambient cooling (since such a drifting baseline is not accounted for in our two-state models). The result for one night is shown in Figure 4. The figure demonstrates that univariate k -Means clustering can effectively determine the on/off transitions in the time series. Furthermore, we have chosen five random nights for each of our 15 ACs and annotated the on/off transitions, as shown in Figure 2, for comparison with the model results. We find an accuracy of 86.9% and 86.7% on the determination of the state of the unit at a given time using a univariate k -Means clustering and HMM, respectively.

3.3 Cycling behavior and ambient temperature

The cycling behavior in Figures 2-3 is characteristic of some AC units in our study, but not all as in Figure 6 which shows the fraction of “on” time each AC unit spends in the cycling mode for each night if we define an AC as being in a cycling mode at time t if there are 2 or more on/off transitions within a 10 minute window. While some units spend most of their “on” time cycling other spend almost no time cycling. This could be due to a variety of factors including equipment differences or variable set temperatures by users. The latter explanation would only be consistent with the figure if the cycling mode indicates that the room has reached the set temperature and the AC unit is cycling to keep the room at that temperature. Were this the case, we would expect that cycling states

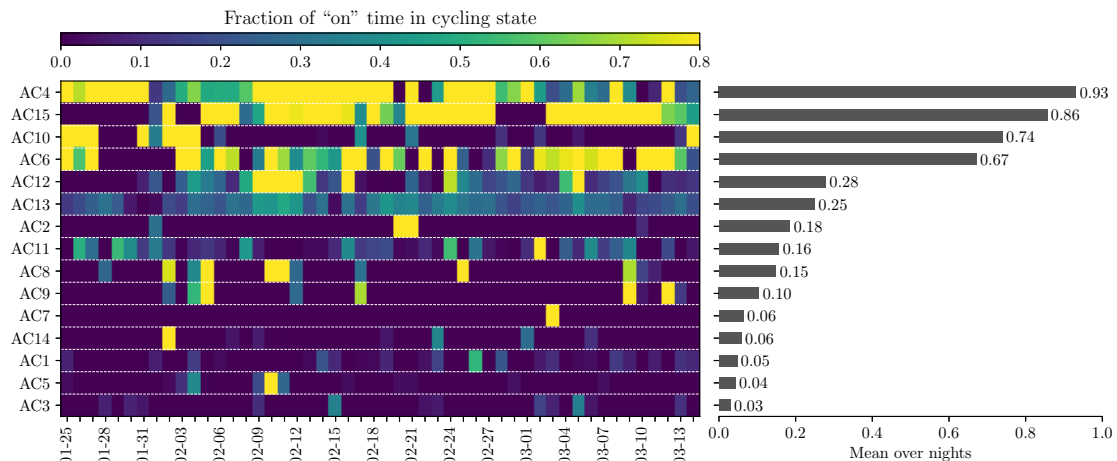


Figure 6: *Left:* the fraction of “on” time spent in the “cycling” state for each AC (Y-axis) on each night (X-axis) in our study. While some AC units spend most of the night cycling, others do not cycle at all, potentially indicating different AC types or differences in user set temperatures. *Right:* the average cycling fraction over all nights for each AC.

characteristically correspond to lower ambient air temperatures since, regardless of set temperature, ACs will take longer to reach the set temperature for higher ambient air temperature. Figure 5 shows that this is indeed the case. The distribution of ambient air temperatures at times when ACs are cycling has characteristically lower values than times when the AC is continuously on. Furthermore, a two-sample Kolmogorov-Smirnov (KS) test yields a value of 0.173 with a p -value of ~ 0 indicating a very low probability that the two distributions represent random draws from the same parent distribution. This is consistent with our hypothesis that cycling states correspond to maintenance of temperature as opposed to primarily cooling.

4 DISCUSSION AND CONCLUSIONS

We have presented results from a continuous infrared-imaging campaign of a building at stand-off distance with persistent observations over 50 nights at a temporal granularity of 10 seconds between images. By extracting the time-dependent brightness of exterior AC vents, which present as bright IR sources in our images, and analyzing them for on/off transitions, we find both operational and behavioral characteristics associated with AC use. In particular, we note that the application of state-based change point detection models to time series of AC IR brightness result in on/off detections consistent with dual operational modes: a continuous “on” mode followed by a cycling mode characterized by rapid on/off transitions. The distribution of these on/off is consistent with a scenario in which the continuous on mode represents a cooling stage and the cycling behavior represents a temperature maintenance stage.

We note that the implications for this work are numerous. For example, we have also shown that not all AC units display cycling behavior and the analysis of cycling frequency and characteristics may reveal faults or operational differences in AC units. Furthermore, the cycling frequency itself (and its stability) may indicate potential inefficiencies or pending breakage of AC components. Finally, the aggregate behavioral properties of AC use can yield important information for characterizing energy end-use, particularly if AC IR brightness correlates with energy consumption for a

given unit. Relatedly, while the spatial resolution of our IR imaging is certainly sufficiently low to preclude the possibility of collecting personally identifiable information (see also [3, 4]), the patterns of AC use may be characteristic of an individual user and so while we have presented completely anonymized and aggregated results here, it is an open question at what level of disaggregation behavioral patterns become unique to a certain AC unit. Each of these use cases and considerations will be the focus of future work.

ACKNOWLEDGMENTS

This work is supported by the National Research Foundation of Singapore through a grant (#1645964) for the Singapore-Berkeley Building Efficiency and Sustainability in the Tropics (SinBerBEST) program and by James S. McDonnell Foundation Complex Systems Scholar award (number:220020434), United States.

REFERENCES

- [1] Marc Bruant, Gérard Guarracino, and Pierre Michel. 2001. Design and tuning of a fuzzy controller for indoor air quality and thermal comfort management. *International journal of solar energy* 21, 2-3 (2001), 81–109.
- [2] William Chung. 2011. Review of building energy-use performance benchmarking methodologies. *Applied Energy* 88, 5 (2011), 1470–1479.
- [3] Gregory Dobler, Federica B Bianco, Mohit S Sharma, Andreas Karpf, Julien Baur, Masoud Ghandehari, Jonathan Wurtele, and Steven E Koonin. 2021. The Urban Observatory: A Multi-Modal Imaging Platform for the Study of Dynamics in Complex Urban Systems. *Remote Sensing* 13, 8 (2021), 1426.
- [4] Gregory Dobler, Masoud Ghandehari, Steven E Koonin, Rouzbeh Nazari, Aristides Patrinos, Mohit S Sharma, Arya Tafvizi, Huy T Vo, and Jonathan S Wurtele. 2015. Dynamics of the urban lightscape. *Information Systems* 54 (2015), 115–126.
- [5] E Elgendy, M Melike, and M Fatouh. 2018. Experimental assessment of a split air conditioner working with R-417A under different indoor and outdoor conditions. *International Journal of Refrigeration* 85 (2018), 268–281.
- [6] Hiromi Habara, Rakuto Yasue, and Yoshiyuki Shimoda. 2013. Survey on the occupant behavior relating to window and air conditioner operation in the residential buildings. In *13th Conference of International Building Performance Simulation Association*. 2007–2013.
- [7] Guanqing Lin, Hannah Kramer, and Jessica Granderson. 2020. Building fault detection and diagnostics: Achieved savings, and methods to evaluate algorithm performance. *Building and Environment* 168 (2020), 106505.
- [8] Elena Lucchi. 2018. Applications of the infrared thermography in the energy audit of buildings: A review. *Renewable and Sustainable Energy Reviews* 82 (2018), 3077–3090.
- [9] Bin Yang and Thomas Olofsson. 2017. A questionnaire survey on sleep environment conditioned by different cooling modes in multistorey residential buildings of Singapore. *Indoor and Built Environment* 26, 1 (2017), 21–31.