Pervasive and Mobile Computing 26 (2016) 91–101

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Mining hierarchical relations in building management variables

Luis I. Lopera Gonzalez *, Oliver Amft

ACTLab, University of Passau, Germany

ARTICLE INFO

Article history:
Available online 28 October 2015

Keywords:
Activity
Hierarchical activity recognition
Activity discovery
Prepositional rule clustering
Unsupervised hierarchical clustering
Building automation

ABSTRACT

We present a framework to relate variables as they occur in a modern building management system (BMS) that processes data from building-installed sensor and actuators. Our group mining framework extracts a unified event time series as changes in building management variables, derives prepositional variable association rules, and extracts hierarchical variable groups from the derived rules. Variable changes typically occur by either occupants interacting with the building or as response to outdoor environmental changes. As a user enters a building, a sequence of sensor activations will create a specific temporal event pattern that is mapped into a variable hierarchy by our framework. Similarly, as outdoor lighting changes, a variable hierarchy appears that relates variables to the change. To extract variable groups, we introduce a novel hierarchical transitive clustering (HTC) algorithm that constructs a rooted variable tree and then clusters the tree to represent variable group relations. HTC is parameter-free and works unsupervised. We evaluated the group mining framework in living-lab data recorded in different office environments during 14 months. As typical for BMS operation, variables in our dataset represent measurements and control states of building-installed devices and processed context information. HTC showed a correctness of over 0.91 and an average variable coverage of 75%, this improving variable coverage by 40 p.p. compared to previous work. We successfully detected alternative hierarchies and show how variables relate across office rooms.

1. Introduction

Effective operation of modern buildings relies on interconnected ubiquitous devices. Examples of ubiquitous devices installed in many large buildings include movement detectors and light sensors, besides actuators, e.g., ceiling lights and air ventilation. The devices are represented as measurement and control variables in a building management system (BMS). The variables are then combined in rules or more advanced interpretation algorithms to infer building context states. For example, Milenkovic et al. [1] showed how several office user activities could be recognised from sensors already available in many buildings. Detected user activities are again variables. A BMS thus links measurements from sensors and operations of actuators. For example, occupancy measurements in a meeting room, indoor temperature, and outdoor weather conditions are used to control optimal operation of a heating, ventilation, and air-conditioning (HVAC) system in this meeting room.

The motivation of building owners and operators to further increase measurement and automation functions are manifold, including saving energy [2], and improve occupant comfort under energy efficiency constraints [3,4]. Buildings,

* Corresponding author.
E-mail addresses: luis.loperagonzalez@uni-passau.de (L.I. Lopera Gonzalez), oliver.amft@uni-passau.de (O. Amft).

http://dx.doi.org/10.1016/j.pmcj.2015.10.009
1574-1192/© 2016 Published by Elsevier B.V.

such as office towers may utilise about ten measurement and control variables per desk, resulting in 100 variables for a 10-occupant open office space. In a typical office building with ten floors and 500 desks, as many as 50.000 variables can be expected, not considering corridors, elevators, meeting rooms, etc. Priyadarshini et al. [5] reported an application example of wireless sensor nodes, where ubiquitous integration and rapid increase in devices and variables is well illustrated. For scalability, BMS development is thus migrating to a service-oriented architecture, as illustrated by Degeler et al. [6]. Nevertheless, a great deal of manual labour is required during commissioning to correctly relate sensors and actuators, derive context variables, and maintain consistency during a building's lifetime.

Whether newly build or refurbished, commissioning of a BMS is today performed by expert technicians that work together with device and system suppliers to correctly interlink variables. During maintenance, further effort is spent to correct variable naming errors and updating configurations upon device exchanges, failures, or upgrades. It is technically conceivable that any modern building-installed device would automatically identify itself, revealing its location and interrelation of variables with other devices. However, several constraints prohibit self-configuring features from being implemented in individual building-installed devices, including device cost, size, power consumption, and robustness for continuous operation over many years. Due to the cost of large BMS installations, sensors and actuators need to be affordable. For example, EnOcean (www.enocean.com) develops wireless self-powered sensors and actuators, including motion detectors and wall switches, which are even used in new building constructions to minimise physical network infrastructure. Therefore, methods that could mine variable relations from a central BMS based on the variables' data streams are sought. In earlier work, we showed that mining relations among BMS variables and grouping variables has many applications, including detecting missing or broken devices, detecting and correcting variable naming errors, and deriving new variables such as people count in office spaces [7].

We present a novel approach to derive groups and group hierarchy from variable association rules that are mined in event time series of variable states. Our hypothesis is that variables that relate to the same physical space in a building will provide events with temporal relation. Variables that have a causal relation, i.e., a variable changes in response to a change in another variable, should have a temporal dependency in states that could be extracted from their event time series. For example, occupants walking through a building space towards a desk will create changes in presence and activity variables and potentially in light actuator states, first for the building space, i.e. room, and then at the desk. For our approach, the variables could have different origin, including sensor measurements, actuator states, and derived states. Our approach considers state changes across all variable as one unified, temporal ordered event stream and derives variable grouping without supervision. As rules are mined, temporal relation of events provides the basis for variable groupings. Subsequently, variable hierarchy is determined to describe dependencies.

In particular, the paper provides the following contributions:

1. We present a group mining framework to extract variable association rules from the unified variable event stream and then extract variable groups from the rules. A novel parameter-free hierarchical tree clustering (HTC) algorithm is introduced to group variables into hierarchical structures. HTC works in two steps, where initially a rooted variable tree is constructed based on the extracted variable association rules and subsequently the tree is clustered and trimmed to represent variable group relations.
2. We evaluate our mining framework using actual BMS data of several months and derive variable grouping performance. The grouping performance is compared to expert generated groups and to groups derived from BMS configurations. The evaluation confirmed that HTC provides relevant groupings that are robust even if only subsets of the BMS data was considered as source.

In our previous work, we devised a weighted transitive clustering (WTC) algorithm to discover variable grouping [7]. However, the groups derived by WTC provided no hierarchy information. As a result, when an environmental variable is associated with more than one room group, the linked room groups are merged into one. Additionally, to distinguish different hierarchy levels, e.g. room and desks, WTC needs to be run with different parameter settings. In contrast, the HTC algorithm presented in this work produces hierarchical groups that describe the relationship between rooms and desks. The hierarchical groups can also describe interconnections between groups created by environmental variables. In this work, we demonstrate that HTC can outperform the WTC algorithm. HTC finds groups that associate a majority of variables and identifies alternative relations between variables. Moreover, due to the hierarchical representation, variable relations could be inspected graphically.

2. Related work

In wireless sensors networks, device localisation is an active field of research. One common approach to localisation is device-based, i.e. using sensor node communication capabilities to estimate a node's relative location with respect to its neighbours [8,9]. Aspens et al. [10] provide a detailed theory for automatically constructing graphs that represent the sensor location in a bi-dimensional or three-dimensional space. They considered the distance between neighbouring nodes known and assume fixed positions of beacon nodes. In practice, the location of beacons is manually configured in order to map a node's location to a physical location. Furthermore, beacons become critical path points in the system, as broken beacons may prevent correct node localisation.