Planning for Coordination of Devices in Energy-Smart Environments

Ilche Georgievski
Johann Bernoulli Institute for Mathematics and Computer Science
University of Groningen
The Netherlands
initial.surname@rug.nl

Abstract

As a sustainability property, energy efficiency is of an extreme importance, especially in environments that are heavy energy consumers, such as homes and buildings. Nowadays, homes and buildings are equipped with many devices that could be exploited in order to make them smart and energy-efficient. Our vision is to bring convergence of smart environments, energy efficiency and automated planning by proposing a planning framework for energy-efficient coordination of devices in smart environments. We establish a proof of concept confirming that automated monitoring and control of devices can lead to significant savings not only on energy, but also on the amount paid for that energy. We envision use of Hierarchical Task Network (HTN) planning due to several reasons identified in our profound overview of the most popular HTN planners. We strive to answer several research questions relating to the general design of the planning framework, the use and improvement of HTN planning, the support for users to interact with the planning framework, and the evaluation of the framework in a living-lab set up at the University of Groningen.

Modern living environments, such as homes and buildings tend to be equipped with a variety of devices usually called 'smart' devices. The group of devices includes different sensors and actuators, where, both, the sensors and actuators, provide information about the environment, and only the actuators enable controlling the environment. Environments that embed such smart devices are called smart environments. However, embedding smart devices into physical environments surely does not mean having a smart environment by default. Without proper processing and computation of raw sensed data, the environment will not be able to smartly react to the contextual changes and occupant needs.

Furthermore, consider the following important problem of sustainability. Buildings account around 40% of energy consumption in European Union and up to 50% in the United Kingdom and Switzerland, being the largest CO_2 producers (EU 2010). Moreover, the energy consumption of typical industrial and commercial buildings adds up to around 30% of the total operational costs. Thus, addressing the problem and making smarter use of energy in buildings will fundamentally contribute to energy and cost savings.

As sensors and actuators provide only naive control of the environment, current buildings are not optimised with respect to energy consumption. The obvious research gap can be filled by new paradigms, approaches and frameworks that will create intelligent adaptations of the environment while increasing occupant comfort and keeping the environment in the most energy and cost efficient state.

The necessity of information processing and computation in terms of searching, reasoning and learning in order to create sophisticated adaptations of the environment brings us in the area of Artificial Intelligence (AI). Various approaches have been explored, such as learning, affective computing, temporal reasoning, fuzzy logic, agent-based systems, event-condition-action rules, and automated planning (Sadri 2011).

We find a strong motivation to develop a framework that enables smart coordination and utilisation of a variety of devices. The framework should guarantee many desired properties, such as performance, fault-tolerance and scalability in its objective to achieve energy savings. Our vision is to introduce novel methods that compose adaptations, orchestrate functionalities and accomplish goals of a building and one or more occupants.

Current Research

The focus of this paper is on the smart environments, particularly intelligent buildings, energy efficiency, and automated planning.

State of the Art

Smart and Energy-Efficient Environments The area of smart environments envisions sensitive, responsive, adaptive, and transparent environments. The important challenge here is to effortlessly emerge from an environment with a variety of embedded devices to an intelligent and computationally capable environment. In addition, a trendy vision is the transformation of conventional and intelligent environments into energy-efficient ones. Thus, the grand objective is to build environments that are unobtrusive and intelligent, satisfy occupants' needs and achieve energy efficiency.

Smart environments, their approaches and technologies are heavily reviewed in (Nakashima, Aghajan, and Augusto

---

2009; Cook, Augusto, and Jakkula 2009; Sadri 2011). These studies also discuss the role of AI in automated control, decision making and computational capabilities of smart environments.

Several AI techniques are employed in order to save energy and satisfy occupant needs in intelligent buildings. In (Boman et al. 1998), a multi-agent approach is used to monitor and control a building. Different agents are responsible for different aspects of the building. For example, a room agent ensures that a particular room is in the most energy-efficient situation, while satisfying occupant preferences at the same time. Among the several components that the room agent contains is a plan module. This module maintains sequences of actions stored in a plan library. When a particular rule is triggered, an associated with it sequence of actions is executed. The approach is tested in a simulated building environment, and the results show savings up to 40% (Davidsson and Boman 2000b; 2000a). In (Boman, Davidsson, and Younes 1999), an improvement is suggested by using automated decision making whenever an uncertain situation occurs. The agents use a pronouncer that gives decision support in dynamic and real-time manner by evaluating an input (either a decision tree or an influence diagram) and returning the best action. Another multi-agent based approach (Qiao, Liu, and Guy 2006) supports user preference learning and personalised control and feedback for which no experimental validation is provided. (Lin et al. 2010) propose a multi-agent system to automatically control an intelligent building by using information fusion. The efficiency of their system is not validated as well. To conclude, from a technical point of view, most of the studies use multi-agent approach, and from efficiency point of view, most of the studies do not present actual energy savings and do not use real-world environments.

Planning Environments exhibit smart behaviour if they are able to automatically react to environmental changes and occupant needs. Automatic adaptations are possible by incorporating sophisticated computational techniques, especially those offered by the area of AI, such as activity recognition, learning, automated planning, and temporal reasoning.

AI-inspired techniques for coordinating operations can leverage the degree of intelligence a building exhibits. Given the dynamic nature of smart environments, e.g., the constant appearance and disappearance of heterogeneous devices and their functionalities, continuous changes of the state of the devices, and the movements of occupants, the number of contextual states can be enormous. We are interested in automated planning as it provides means for achieving dynamic solutions in an atomic way. The main advantage is that solutions do not have to be hard-wired and static, but are computed on-the-fly so that they are tailored to the current state of the environment and the objectives of the occupant and the building itself. Hence, by employing automated planning, many challenges can be addresses, including the contextual awareness of the solution, handling device contingency, support of heterogeneous devices, and enabling the occupants to issue goals directly.

Automated planning has never been risen to prominence in smart environments. (Heider and Kirste 2002; 2005) propose a planner-based approach enabling the users to interact with modern environments, such as networked infotainment systems and smart homes. The justification for employing automated planning is twofold. Firstly, the system allows users to express goals (or intentions), and secondly, the system can be dynamically configurable, i.e., it reasons over available actions regardless the appearing/disappearing of devices in/from the environment. In their approach, the user intention and contextual information are translated into a particular declarative goal for which a sequence of actions is produced by the UCPOP planner (Penberthy and Weld 1992). (Ranganathan and Campbell 2004) propose a planning framework to control smart environments based on a goal specification and STRIPS-based planning. The planner takes an abstract goal specification, generates a template goal state and decides what is the ‘best’ final state. The planner, in this case Blackbox(Kautz and Selman 1999), then plans for actions. The planning framework also monitors the action executions, and, retries the actions or replans, if necessary. Actions are described in PDDL (Mcdermott et al. 1998), and can be extended with additional information about handling failures. While the approach brings users into a perspective by allowing them to specify goals through a graphical user interface, it is not experimentally validated. (Amigoni et al. 2005) propose a planning system, called Distributed HTN, based on Hierarchical Task Network (HTN) planning (Nau, Ghallab, and Traverso 2004). D-HTN plans for a centralised activity tailored on the capabilities of distributed devices. The system adopts a multiagent approach. Particularly interesting for the planning system is that for each decomposition it allows associating three values: a performance value, which indicates the expected effectiveness of the decomposition; a cost value, which indicates the resource consumption of the tasks in the decomposition; and a probability of success, which indicates the expected likeliness that no error occurs. These values provide a heuristic for a ‘good’ selection of decomposition when more available. Results from realistic experimentation demonstrate that “D-HTN is an effective approach to provide AmI systems with goal-oriented capabilities”. (Bader and Dyrba 2011) propose an approach in which the controlling of environment is accomplished through independent behaviours that produce goals. Each such behaviour defines a particular adjustment, e.g., setting up a projector and projector screen whenever a video source is connected. Goals are grouped into one objective for which a planner produces a sequence of actions. However, the creation of sequence of actions is not a pure planning process, but rather a process that maps each goal to one device state. The impact of the proposed approach is unclear as it is not experimentally validated. (Krüger et al. 2011) propose a proactive approach to control devices by inferring the user intentions. An intention is hierarchically represented by a Collaborating Task Modelling Language (CTML) model (Wurdel, Sinnig, and Forbrig 2008), while actions are represented by the PDDL model. On one hand, the CTML model is translated into a probabilistic model, and, on the other hand, a planner is used to create sequences of actions for every possible situation. These sequences are associated to states of the probabilistic model from which the most promising sequence is chosen. This approach, too, misses an experimental validation. (Kaldeli et al. 2012) propose an architecture for smart homes that incorporates automated planning in order to create adaptations at runtime. In this case, the domain is modelled as a Constraint Satisfaction Problem (CSP). The planner first prunes the actions that are irrelevant, and then the actual plan is created. The planner is triggered whenever a contextual change has happened and the planner is 'subscribed' to it. The planner has some additional features of particular importance for smart environments, such as efficient handling of variables ranging over large domains. A fully working prototype is implemented and evaluated in terms of performance and end-user satisfaction. Results show that planning is quite efficient from a technical point of view, and it satisfies the expectations and objectives from a user point of view.

(Marquardt and Uhrmacher 2008) discuss the use of automated planning for creating solutions in smart environments, and the behaviour of four planners, namely UCPOP (Penberthy and Weld 1992), SGP (Weld, Anderson, and Smith 1998), Blackbox (Kautz and Selman 1999), and a simple progression planner. They define several metrics that could have impact on the planning runtime in smart environments, including the number of devices, number of services, number of services per device, and the number of pre- and post-conditions per service. Several experiments are conducted in a simulated setting. The results show that "the number of services influences the runtime of the planning process, but, in comparison, the contribution of the pre- and post-conditions has much stronger impact". Among the planners, Blackbox behaves the best and scales better than the other planners. All planners tend to have unusual long runtimes when no plan exists.

Our Vision

Smart and Energy-Efficient Environments Basically, one important question we need to answer is: *What is the design of a framework that allows devices to be seamlessly and intelligently coordinated, while minimising energy consumption and satisfying occupant preferences inside a building?* Moreover, following the trend in the energy market, the framework should be able to cope with the evolution of the Smart Grid, i.e., to consider the information on the price of energy proposed by different providers and the maximum amount of energy available at that price.

To that end, we present a prototype framework that coordinates building offices to save energy and overall energy price costs assuming the availability of the Smart Grid (Georgievski et al. 2012). Our approach is able to monitor the energy consumption of devices, monitor the energy production of small-scale generating units, acquire dynamically the prices of energy from different providers and closing contracts for short term time intervals, take into account policies for devices so to conform to occupant preferences, and automatically control in an optimal way the energy consumption of devices following the policies. Our initial research shows that intelligent and automatic control of devices is able to reduce the overall energy consumption (up to 15%) and, coupled with dynamic pricing from the Smart Grid, is able to provide considerable financial savings (up to 35%).

The prototype framework is deployed in our own offices at the University of Groningen. The offices are located on the fifth floor of a more than 10000 m^2 recently erected building[1]. The living lab will be extended continuously with new devices and will serve as a test-bed for our planning framework.

Planning We envision a planning system that satisfies several requirements. In fact, an intelligent building must support the occupants by providing a functionality to express and reason over *preferences*. For example, a building occupant may want higher temperature than the default one in the office. We argue that the domain of intelligent buildings shows an inclination towards hierarchical *representation*. We know in advance many ways how the building should be adjusted to a particular setting. For example, the knowledge of adjusting a room to a meeting setting may include turning all the lights off, setting the window blinds down, setting the projector screen to an appropriate position, and turning the projector on. Moreover, by providing extensive knowledge into hierarchies, we can simplify the description of operators, which, more or less, can map directly to the functionalities the devices provide. Given the dynamic environment, this could be beneficial in automating the way of adding a new device into the environment. A critical aspect is the *expressiveness* of the knowledge representation. Either of both ways, the hierarchical or operator representation only, must be powerful enough to express operational semantics of different devices and the semantics of different situations that could arise in intelligent and energy-efficient environments. Heider and Kirste (Heider and Kirste 2002; 2005) identify several requirements regarding the expressiveness of intelligent environments when AI planning used. The domain of intelligent buildings has a tendency towards *temporal* behaviour. Starting from temporal preferences expressed by occupants, e.g., I never want window blinds down, or building itself, e.g., a microwave must be never turned off, to properties with absolute time quantities, e.g., a boiler should be turned off at 23:00 every working day and turned on at 7:00 next working day. In addition, the ability of the planning system to generate *partially ordered plans* is highly needed. In order to deal with unexpected situations and failures that may occur during execution of some action, the planning system should *monitor* the action execution, and perform *re-planning*, if necessary. *Scalability* is an important aspect of complex environments. One system is said to be scalable, if it is capable to cope and perform under a growing amount of load. In an intelligent environment, the scalability can refer to the capability of the planning system to maintain or increase the level of performance when new resources, such as sensors or actuators are added to the environment. It is difficult to define dimensions that could measure the size of problem, but, of course, scalability is

---

[1] http://nl.wikipedia.org/wiki/Bernoulliborg

highly desirable in practical setting with a large number of facts about the environment, a large number of occupants, or a large number of operators and hierarchical descriptions. For example, the domain of GreenerBuildings project (GB 2012), which is an information and communication technologies project funded under the European Seventh Framework Programme on engineering of networked monitoring and control systems and wireless sensor networks and cooperating objects, envisions an intelligent building in which a number of devices an aims at a distributed network of hundreds of devices, thus, a very large number of actions (if we assume that each device maps to at least one action), a tens of tasks, a hundreds of occupants, and a thousands of facts. *Performance* is closely related to scalability. The ability of the planning system to generate solution fast is profound. The environment should adapt to a new state in a time-negligible manner from the moment of issuing the need for environmental change. Last but not least, the *energy awareness* of the planning framework must ensure that its plans are optimal and will transform the environment in the most energy-efficient state.

Among planning techniques that seem intrinsically interesting is Hierarchical Task Network (HTN) planning (Russell and Norvig 2003). An HTN planner is provided with a set of goal tasks that have to be repeatedly decomposed until primitive tasks are reached. Such task decomposition is based on the hierarchies called methods contained in the domain knowledge. Well-written methods can significantly reduce search space and help planner to find an efficiently executable plan. In (Georgievski and Aiello 2012), we offer an excellent consolidation point of well-established HTN approaches. We cover the state-of-the-art HTN planners and discuss their principles, planning techniques and processes, search spaces, and other peculiarities. We compare the planners based on the several criteria and we give perspectives for future research ideas where HTN planning appears to be weak or insufficiently investigated, as is the case with smart and energy-efficient environments.

We strive to answer several research questions. As it is known that HTN planning is the most used planning technique in real-world applications (Georgievski and Aiello 2012), a question comes out: *Is HTN planning adequate and efficient approach for intelligent and energy-efficient environments, such as intelligent buildings?* Given the requirements for smart and energy-efficient environments, and the general interest of automated planning community, the next question is: *How can HTN planning be extended and improved with respect to the identified requirements and the standards of the automated planning community?* This question could be divided into several sub-questions that would provide more focused directions for answers. Beside the technical challenges, the user/occupant needs with respect to the interaction with the planning framework need to be taken into account. As usually the interaction aspect is neglected, we want to give attention to it and answer the following question: *How to model a tool that will be conceptually understandable and will make the user comfortable to interact with the planning framework?* Another question relates to the practical use of the planning framework in a real working environment: *How can be the planning framework tested and evaluated in actual building settings?*

Future Steps

The static analysis that we perform and deals with study of research areas and collection of data from published literature is nearly finalised. The ongoing step is the developmental research. A first prototype of a state-based HTN planner is developed in Scala programming language. Our further work will focus on developing and describing methods and approaches that will satisfy the general design of the planning framework and the specific requirements identified for smart and energy-efficient environments. We believe that these activities will provide fundamental answers to our research questions. At the same time, we will continuously perform experimental validations in our living lab.

Acknowledgements

The work is supported by the EU FP7 Project GreenerBuildings, contract no. 258888 and the Dutch National Research Council under the NWO Smart Energy Systems program, contract no. 647.000.004.

References

Amigoni, F.; Gatti, N.; Pinciroli, C.; and Roveri, M. 2005. What planner for ambient intelligence applications? *IEEE Transactions on Systems, Man and Cybernetics, Part A* 35(1):7–21.

Bader, S., and Dyrba, M. 2011. Goalaviour-based control of heterogeneous and distributed smart environments. In *Proceedings of the 7th International Conference on Intelligent Environments*, IE '11, 142–148. Washington, DC, USA: IEEE Computer Society.

Boman, M.; Davidsson, P.; Skarmeas, N.; Clark, K.; and Gustavsson, R. 1998. Energy saving and added customer value in intelligent buildings. In *Proceedings of the 3rd International Conference on the Practical Application of lntelligent Agents and Multi-Agent Technology*, PAAM'98, 505–517.

Boman, M.; Davidsson, P.; and Younes, H. L. S. 1999. Artificial decision making under uncertainty in intelligent buildings. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 65–70.

Cook, D. J.; Augusto, J. C.; and Jakkula, V. R. 2009. Ambient intelligence: Technologies, applications, and opportunities. *Pervasive Mob. Comput.* 5(4):277–298.

Davidsson, P., and Boman, M. 2000a. A multi-agent system for controlling intelligent buildings. In *Proceedings of the Fourth International Conference on Multiagent Systems*, 377 –378.

Davidsson, P., and Boman, M. 2000b. Saving energy and providing value added services in intelligent buildings: A mas approach. In *ASA/MA*, 166–177.

EU. 2010. Directive 2010/31/eu of the european parliament and of the council on the energy performance of buildings.

GB. 2012. Greenerbuildings project @ONLINE.

Georgievski, I., and Aiello, M. 2012. An overview of hierarchical task network planning. JBI Preprint JBI 2012-12-5, University of Groningen.

Georgievski, I.; Viktoriya, D.; Pagani, G. A.; Nguyen, T. A.; Lazovik, A.; and Aiello, M. 2012. Optimizing energy costs for offices connected to the smart grid. *IEEE Transactions on Smart Grid* 3:2273–2285.

Heider, T., and Kirste, T. 2002. Supporting goal-based interaction with dynamic intelligent environments. In *Proceedings of the 15th European Conference on Artificial Intelligence*, 596–600.

Heider, T., and Kirste, T. 2005. Smart environments and self-organizing appliance ensembles. In *Mobile Computing and Ambient Intelligence*. Springer.

Kaldeli, E.; Warriach, E. U.; Lazovik, A.; and Aiello, M. 2012. Coordinating the web of services for a smart home. *ACM Transactions on the Web*.

Kautz, H., and Selman, B. 1999. Unifying sat-based and graph-based planning. In *Proceedings of the 16th international joint conference on Artifical intelligence - Volume 1*, IJCAI'99, 318–325. Morgan Kaufmann Publishers Inc.

Krüger, F.; Ruscher, G.; Bader, S.; and Kirste, T. 2011. A context-aware proactive controller for smart environments. In *Proceedings of the 2nd Semantic Models for Adaptive Interactive Systems Workshop*, volume 747 of *CEUR Workshop Proceedings*.

Lin, Z.; Guiqing, Z.; Bin, S.; Xiuying, X.; and Qiao, Y. 2010. Building energy saving design based on multi-agent system. In *The 5th IEEE Conference on Industrial Electronics and Applications*, ICIEA'10, 840–844. IEEE.

Marquardt, F., and Uhrmacher, A. M. 2008. Evaluating ai planning for service composition in smart environments. In *Proceedings of the 7th International Conference on Mobile and Ubiquitous Multimedia*, MUM '08, 48–55. New York, NY, USA: ACM.

Mcdermott, D.; Ghallab, M.; Howe, A.; Knoblock, C.; Ram, A.; Veloso, M.; Weld, D.; and Wilkins, D. 1998. PDDL - The Planning Domain Definition Language. Technical report, CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control.

Nakashima, H.; Aghajan, H.; and Augusto, J. C. 2009. *Handbook of Ambient Intelligence and Smart Environments*. Springer Publishing Company, Incorporated, 1st edition.

Nau, D. S.; Ghallab, M.; and Traverso, P. 2004. *Automated Planning: Theory & Practice*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Penberthy, J. S., and Weld, D. S. 1992. Ucpop: A sound, complete, partial order planner for adl. In *KR*, 103–114.

Qiao, B.; Liu, K.; and Guy, C. 2006. A multi-agent system for building control. In *Proceedings of the IEEE/WIC/ACM international conference on Intelligent Agent Technology*, IAT '06, 653–659. Washington, DC, USA: IEEE Computer Society.

Ranganathan, A., and Campbell, R. H. 2004. Autonomic pervasive computing based on planning. In *Proceedings of the 1st International Conference on Autonomic Computing*, ICAC '04, 80–87. Washington, DC, USA: IEEE Computer Society.

Russell, S. J., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Pearson Education.

Sadri, F. 2011. Ambient intelligence: A survey. *ACM Computing Surveys* 43(4):36:1–36:66.

Weld, D. S.; Anderson, C. R.; and Smith, D. E. 1998. Extending graphplan to handle uncertainty and sensing actions. In *Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence*, AAAI '98/IAAI '98, 897–904. Menlo Park, CA, USA: American Association for Artificial Intelligence.

Wurdel, M.; Sinnig, D.; and Forbrig, P. 2008. Ctml: Domain and task modeling for collaborative environments. *J. UCS* 14(19):3188–3201.