Electric Loads as Real-time tasks: an application of Real-Time Physical Systems

Marco L. Della Vedova and Tullio Facchinetti
Department of Computer Engineering and Systems Science
University of Pavia
via Ferrata, 1 – 27100 Pavia, Italy
Email: {marco.dellavedova, tullio.facchinetti}@unipv.it

***Abstract*—This paper describes the application of Real-Time Physical Systems (RTPS) as a novel approach to model the physical process of Cyber-Physical Systems (CPS), with specific focus on Cyber-Physical Energy Systems (CPES). The proposed approach is based on the real-time scheduling theory which is nowadays developed to manage concurrent computing tasks on processing platforms. Therefore, the physical process is modeled in terms of real-time parameters and timing constraints, so that real-time scheduling algorithms can be applied to manage the timely allocation of resources. The advantage is to leverage the strong mathematical background of real-time systems in order to achieve predictability and timing correctness on the physical process behind the considered CPS.**

The paper provides an introduction to the possible application of RTPS to energy systems. The analogy between real-time computing systems and energy systems is presented; moreover, the relationship between RTPS and related research fields is traced. Finally, the introduced techniques are proposed to optimize the peak load of power consumption in electric power systems. This method is suitable for systems spanning from small networks to smart grids.

I. Introduction

In the last decade, the field of embedded systems has been one of the most active research area for both academia and industry. Typical involved research areas related to embedded systems, and distributed embedded systems in particular, include communication systems, process control, and real-time computing. The complexity of modern distributed embedded systems, and the need of improved modeling, monitoring and control techniques of the underlying physical processes, have motivated the growing trend to bring embedded computing closer and closer to the physical process to be monitored and operated. This effort led to the brand new research field of Cyber-Physical Systems (CPS). The innovative idea behind the novel approach presented in this paper arises from an unconventional application of the real-time scheduling theory to model and control a physical process in CPS.

The real-time scheduling theory has been traditionally developed to manage the execution of processing tasks on processors under timing constraints. Several aspects have been addressed and many relevant results have been carried out since the seminal work of Liu and Layland [1], including different task models (sporadic, aperiodic, etc.), optimal scheduling algorithms, necessary and sufficient tests to guarantee the desired timely system behavior, scheduling techniques for both uniprocessor and multiprocessor platforms, just to mention a few of them.

However, in more general terms, real-time scheduling can be seen as the discipline of allocating resources over time to a set of time-consuming tasks, so that given timing constraints are satisfied. In this more general formulation, resources may not necessarily be processors or computing devices. In fact, real-time scheduling techniques have been already applied to systems other than computing systems. For example, in communication systems, real-time modeling techniques and scheduling algorithms are used to manage sets of messages over a communication channel in a timely manner [2]. In this case, an analogy holds between computing tasks and messages, as well as between processors and communication channels. The meaning of “available bandwidth” changes depending on the particular context, referring to the channel capacity in communication systems, and to processor’s processing time in computing systems. Moreover, timing constraints are enforced on execution times in one case, and (typically) on message’s end-to-end latency in the other. In other words, a real-time task must be guaranteed to terminate its execution before its deadline, while a message must be delivered to the receiver within the given time limit. This analogy allows extending to communication networks many results that have been originally developed for real-time computing systems, and vice-versa.

Some electric devices can be modeled as periodically activated tasks, with a bound on the total time that a load can remain active — thus consuming power — in each period. This bound recalls the worst case execution time (WCET) of a real-time task in computing systems. As for computing tasks, all the time properties of electric loads (periods, deadlines and activation time) must be selected according to their application requirements. Section VI provides some examples of timing constraints related to specific electrical loads. Based on the system model, a priority-based scheduling algorithm can be applied to selectively activate/deactivate each device. The goal is to meet the timing constraints of every load, while guaranteeing an upper bound on the total instantaneous power consumed by the concurrent activation of electric components.

In the real-time systems literature, there is active research on power-aware scheduling strategies to save energy while achieving timing constraints. Such scheduling policies aim at reducing the energy consumption using special features of modern electronic hardware, such as dynamic voltage scaling (DVS) [3]. As in those works, the model proposed in this paper associates a maximum consumed power to each electric device. However, we do not aim at directly reducing the overall energy required by the system. The objective is, instead, to determine a bound on the peak power consumption, and to predictably enforce this bound by scheduling electric devices activations in a timely manner.

Some recent works are addressing the real-time issues related with a technique to improve the efficiency of batteries charge/discharge, for electric vehicles [4]. However, this work is limited to batteries, while our approach is oriented at establishing a general framework for managing energy systems in a real-time manner. In [5], the authors aim at finding optimal schedules for CHP (Combined Heat and Power) systems. The approach is based on global optimization through an integer linear programming formulation of the problem. However, this method is strictly limited to offline optimization, while our technique can be applied online. Moreover, we introduce the novelty of modeling electric loads using real-time parameters, to allow the use of real-time techniques for scheduling the activation of loads. Finally, in [6] the authors describe a cyber-physical energy system as a set of components modeled as dynamical systems. While the modeling of electrical devices is more advanced than the one proposed in this paper (refined modeling is subject of future research in our framework), the goal is not related with achieving peak load constraints and, again, no real-time issues are considered.

Real-Time Physical Systems (RTPS) represent a new class of systems introduced in [7]. The key idea is that a state variable changes over the time depending on the activation status of a *real-time resource*. A real-time resource is a resource that is modeled using real-time parameters and thus can be scheduled over the time using a real-time scheduling algorithm. In RTPS, real-time parameters are put into relationship with the physical state variable, and constraints on this latter are translated into constraints on real-time parameters.

Paper is organized as follows. Relations between the new class of Real-Time Physical Systems and other research fields are discussed in Section II. Section III presents a formal definition of RTPS. The application of RTPS in the field of electric systems is stated in Section IV. The analogy between energy systems and computing systems is drawn in Section V. Section VI lists some examples of electric devices and applications that are suitable for being managed by the proposed techniques, while in Section VII statistical results are provided to show the benefits of the proposed solution. Finally, conclusions are stated in Section VIII.

II. A BRIDGE AMONG RESEARCH AREAS

This section presents an analysis and outlines the relationships among different research areas which are relevant for the topic discussed in this paper. Real-Time Physical Systems can be located at the intersection of 3 main research areas: Real-Time Systems, Cyber-Physical Systems, and Hybrid Systems.

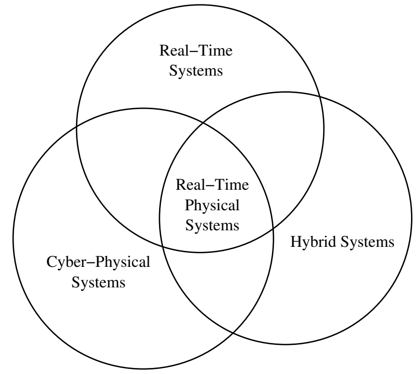

Fig. 1. Collocation of RTPS among research fields.

These relationships are depicted in Figure 1.

Real-time systems are a peculiar type of computing systems in which the correctness of computation does not only depend from the logical/numerical value of the result, but also from the time instant in which the result is made available. In other words, a logically/numerically correct result is not useful — and may lead to catastrophic consequences in some cases — when it is obtained too late. Real-time systems deal with the scheduling of processing tasks on one or many processors under timing constraints. In the field of real-time processing systems, some special cases of RTPS are already studied. Such special cases are thermal-aware and power-aware systems. The research objective is to guarantee timing constraints of computing tasks while reducing temperature and power consumption, respectively. Both problems are emerging in large computing systems as data-centers, while the latter represents a key challenge for battery-powered embedded devices. RTPS can be seen as a generalization of aforementioned real-time systems.

Cyber-Physical Systems (CPS) represent an emerging technology that aims to integrate embedded processing devices to monitor and control physical processes. Cyber-physical systems are intended to address critical applications operating in dynamic and uncertain environments, made by a high number of devices and characterized by complex relationships among components. Several factors can affect system operations, as hardware and software failures, and partial knowledge of the system operating state. Example applications include: automotive, avionics and medical systems; critical infrastructure management, as electric power and water resources; traffic control and safety; advanced robotics for manufacturing or telemedicine (see [8] for details on some specific applications). The study of CPS itself is expected to be a strongly multidisciplinary effort. RTPS are a powerful modeling approach to complex CPS.

Hybrid systems are dynamic systems presenting a two-fold behavior: a continuous and a discrete one. Their peculiarity is to model dynamic systems where the state variable may change in a continuous and discrete manner over the time. A classical example of hybrid system is a bouncing ball, in which the ball speed exhibits continuous behavior between consecutive bounces, while there is a discrete change on the speed vector in correspondence of bounce events.

A branch of hybrid systems research field regards the so called "switched systems". A switched hybrid system is a continuous-time system with isolated discrete switching events. Those events can be classified as: (i) state-dependent or time-dependent, and (ii) autonomous or controlled. RTPS are a special case of time-dependent controlled switched hybrid system [9] in which the switching signal is determined by a pattern of activation/deactivation generated by a real-time scheduling algorithm [10].

III. Real-Time Physical Systems

A Real-Time Physical System (RTPS) is composed by a set of resources that operate on a system characterized by some physical quantities of interest. A resource can be turned on and off by the RTPS controller, called *scheduler*.

As in CPS, RTPS are composed by two tightly interacting components: the computational part, characterized by a discrete timing, and the physical part, which has a continuous behavior. However, with respect to CPS, RTPS have the peculiarity that the "cyber" component is modeled with techniques borrowed from real-time computing systems.

The "physical" component of a RTPS is modeled by a switched dynamic system, defined by the following equation and quantities:

$$\frac{dx(t)}{dt} = f_{s(t)}\left(x(t)\right)$$

where:

- x is the vector of state variables, which are the physical quantities of interest;
- $s \in \{0,1\}^m$ is the *operation mode* of the m resources; each resource can be active (i.e. $s_i = 1$) or inactive (i.e. $s_i = 0$);
- $s(t) = [s_1(t) \ldots s_m(t)]$ is called *switching signal* or *schedule*;
- f_s is a set of 2^m vector fields representing the dynamics of the system.

The "cyber" component of a RTPS concerns the scheduling of resources or, in other words, the generation of the switching signal. The key idea behind RTPS is to use a traditional real-time scheduling algorithm for this purpose. One inherent benefit of this approach is to take advantage from the strong mathematical background which characterizes the analysis of real-time systems. Therefore, powerful analysis techniques developed over more than three decades of research on real-time systems will be leveraged to characterize timing and physical properties of the system. Among other advantages, this approach allows to deal with large and complex systems, having several types of constraints.

Figure 2 shows an example of a RTPS in which the physical value $x(t)$ has an exponential behavior. It decreases when the real-time resource is scheduled for execution, while it increases otherwise. Notice that the physical value behavior

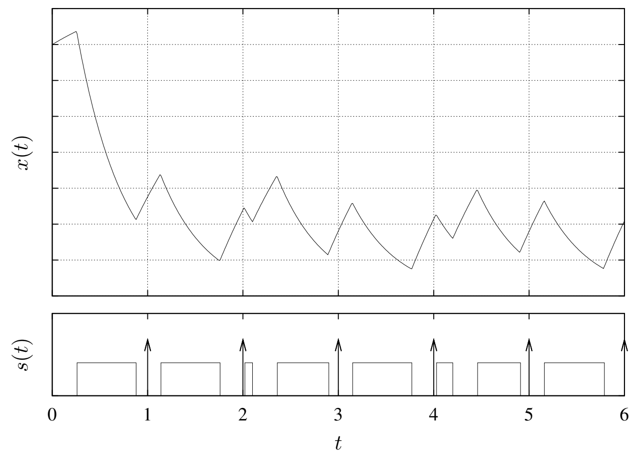

Fig. 2. Example of real-time physical system in which the physical value decreases exponentially when the resource is scheduled for execution, and it increases otherwise.

is influenced only by the switching events between execution ($s(t) = 1$) and idle time ($s(t) = 0$).

The task of the RTPS designer is therefore to choose the scheduling algorithm and to assign to each resource a set of real-time parameters that guarantees the achievement of some user requirements. A typical example of user requirement is that the physical quantities of interest must be bounded within desired working ranges.

A. Timing constraints in RTPS

To further illustrate the concept behind RTPS, and how the classical real-time scheduling theory applies to those systems, let us consider the example depicted in Figure 3. The figure shows 3 computing processes (or tasks) that are scheduled to meet individual timing constraints. Every task τ_i is described by the tuple T_i, D_i, C_i, where T_i is the activation period, D_i is the relative deadline, and C_i represents the duration (or Worst Case Execution Time, WCET). The task τ_i becomes ready for the execution at every kT_i activation time ($k = 0, 1, 2, \ldots$), and must execute for at most C_i time units in the time frame $[kT_i, kT_i + D_i]$. Since, in this example, tasks must be executed on a uniprocessor machine, the scheduling algorithm must ensure that no more than one task should be active at each time, beside the guarantee on aforementioned timing requirements.

The schedule in Figure 3 is one of the most common representations of the pattern generated by a real-time scheduling algorithm. On the other hand, it is straightforward to note that such schedule could represent different systems where the resource to be scheduled must be periodically activated and must stay active for a given amount of time. For example, HVAC systems (Heating, Ventilation and Air Conditioning) have an activation pattern that can be modeled in a periodic fashion. E.g., for a given environmental temperature, to maintain the room temperature within a predefined range, the air conditioning system may stay active for X minutes every

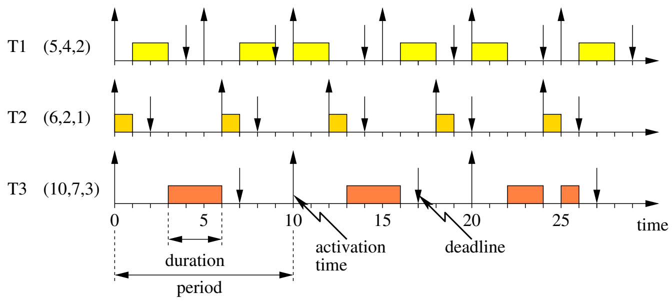

Fig. 3. A schedule of 3 real-time resources using the Earliest Deadline First (EDF) scheduling algorithm.

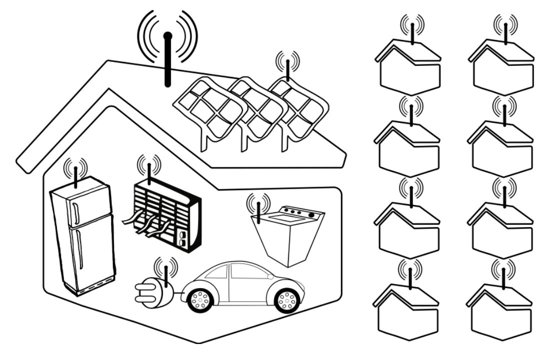

Fig. 4. Networked electric devices at building and neighborhood level. The interconnection is achieved by a wireless network.

Y minutes ($X \leq Y$), which is a typical case of real-time constraint. Therefore, the activation pattern of several HVAC could be represented by a pattern similar to that in Figure 3.

IV. CYBER-PHYSICAL ENERGY SYSTEMS

In the CPS research field, when the focus is on energy systems, the considered systems are referred as Cyber-Physical Energy Systems (CPES) [11]. Our target will be on CPES made by electrical components. In such systems, the "physical" process is made by a network of electric devices that are controlled by a complex set of interconnected embedded systems.

The current technology trend is moving towards the automatic, distributed and coordinated control of electric devices. Some examples can be found in home and factory automation systems [12], large networks of electric cars [13], and automated energy supply and distribution for town and city districts organized in the so-called *smart grids* [14].

The availability of compact and flexible embedded systems allows the effective implementation of Cyber-Physical Energy Systems. Monitoring tasks and control actions can be applied on devices composing the considered physical system. Moreover, involved embedded systems can be connected to build large distributed control networks. Figure 4 shows a set of networked electric components at building and neighborhood level. Consuming, generation (solar cells) and storage (electric cars) elements are depicted. The coordination among devices is achieved by their interconnection to the communication network. For the sake of clarity, and without loss of generality, a wireless network is used to establish the communication infrastructure. Alternative networking technologies could be mixed in a real scenario. A neighborhood made by several buildings can be managed by applying the same abstraction.

A. Managing peak load conditions

The balancing of energy utilization is fundamental for the efficient behavior of an electrical system [15], [16]. For this purpose, specific technical and economical approaches are used to control the distribution of power usage over time. One of the most widely adopted method is the *peak-load pricing*, a policy that assigns higher prices to larger peak-load demands [17].

Peak load conditions, i.e., determined by the simultaneous demand of electrical power by many users, may cause severe issues such as the disruption of power provision, leading to technical and economic issues for both suppliers and users. Moreover, during peak load conditions, the cost of energy production may unpredictably increase in a short time frame due to the impossibility of generating enough energy to satisfy the request of customers. Therefore, energy providers — that must observe contractual obligations with their customers to supply electricity at pre-defined fixed prices — may experience a relevant financial burden. On the other hand, an adequate management of peak load conditions is desirable for energy utilities [18]. An appropriate load management aiming at achieving predictable load conditions may lead to potential contractual benefits to the user. As a consequence, both energy providers and consumers are likely to be interested in load balancing and predictable energy consumption.

Given the aforementioned technical and economic issues, an efficient management of peak-load conditions has the following advantages:

1) the least efficient, i.e., the most expensive, power plants can be proactively turned off if the peak power demand is guaranteed to remain under a given threshold;
2) the electric distribution infrastructure can be tailored for lower peak loads, with less technical issues and reduced costs;
3) the profile of power usage can be smoother and flatter, allowing the final users to have better pricing conditions on the free energy market, where pricing strategies are often driven by the peak-load pricing policy [17].

Real-Time Physical Systems provide the chance to exploit real-time scheduling techniques to the management of loads in CPES. The goal is to balance the total consumed power and the peak power consumption. Real-time scheduling algorithms can be used, as mentioned in Section III, to predictably activate/deactivate electrical devices to guarantee the desired system features, in terms of timing constraints and energy consumption. The typical large size of CPES will take advantage of efficient scheduling algorithms and analysis techniques to determine the system feasibility and properties. It is worth to outline that this paper does not deal with the architecture or the engineering of a CPES. The proposed approach must be intended as a viable modeling technique for the physical energy system, allowing the development of predictable and robust control strategies based on real-time scheduling methodologies.

B. Example of application

Figure 5 shows an example of application of RTPS to the scheduling of electric loads in an apartment. The top graph represents the real measurements of the consumed power in the apartment, along a time period of 5 hours, in absence of specific control actions. It is worth to note the periodic activation of compressors 1 and 2, which cool down the fluid of two independent refrigerators. The activation of remaining loads (lighting and TV devices) depends on user actions. The peak power consumption is around 210 VA, obtained in the period from 17:40 to 18:00. The middle picture allows to figure out the contributions of each load to the total power consumption. Finally, the bottom picture shows an example of possible schedule of the same load set, generated by a coordinating management strategy, where some loads are allowed to be stopped and resumed to decrease the peak load. In particular, the second activation of compressor 1 is stopped and resumed later to allow the activation of the 100W lamp, while the second activation of compressor 2 is delayed. Such actions bring to a peak power consumption of around 155 VA, achieving a peak load reduction of 25%.

V. Electric loads as real-time tasks

The approach proposed in this paper is to exploit the periodic task model [1], widely studied in real-time systems, to represent electric loads as periodically triggered activities. An example of schedule of real-time periodic tasks is given by Figure 3. A bound is imposed on the total amount of time that a load can or should remain active in each period. As for computing tasks, all time properties of electric loads (periods, deadlines and activation time) must be selected according to their application requirements. The activation time plays the role of the WCET (Worst Case Execution Time) in real-time computing tasks. We assume that a load consumes a given amount of electric power while active, and no power when switched off.

Once loads have been modeled using real-time timing parameters, a priority-based scheduling algorithm can be applied to selectively activate/deactivate devices. For example, the Earliest Deadline First (EDF) or the Rate Monotonic (RM) policy can be used to generate the schedule [10]. The two algorithms are known to be optimal for uniprocessor platforms with full preemption, i.e., translating to the case of CPES, they can be suitably applied only when the total utilization does not exceed an upper bound that guarantees their optimality. Both aspects represent possible limitations to the proposed approach. However, they can be faced with classical techniques in the field of real-time scheduling.Multiple simultaneous activations can be manages by partitioning the set of loads as in [19], while preemptions can be limited with dedicated methods [20]. The goal is to meet the timing constraints of each load, while guaranteeing an upper bound on the total instantaneous power consumed by the concurrent activation of electric devices. Table I resumes the analogy between real-time computing and energy systems based on real-time parameters proposed in this paper.

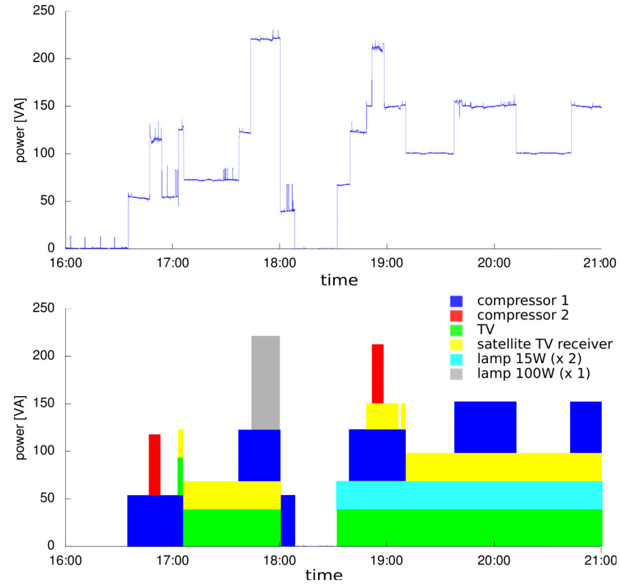

(a) Absence of load scheduling.

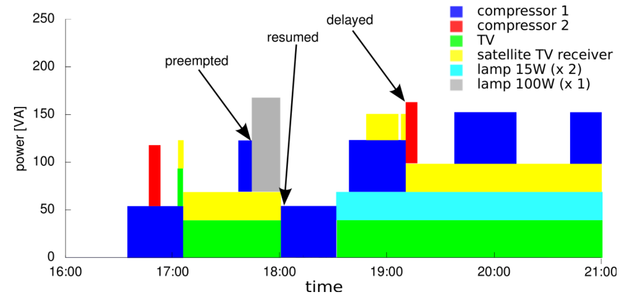

(b) Presence of load scheduling.

Fig. 5. Measurements of consumed power in an apartment. Figure (a) depicts the normal power consumption of some electric loads, both measured power and contributions of specific loads. Figure (b) shows that an adequate scheduling of load activations provides allows to achieve a peak load reduction of 25%.

TABLE I
Resume of the analogy between real-time computing systems and Cyber-Physical Energy Systems.

considered feature	Real-time Systems	Cyber-Physical Energy systems
domain	computing systems	energy systems
resource	task	load
C_i	execution time	activation time
T_i	period	period
D_i	deadline	deadline
$s(t)$	schedule	switching signal
optimization target	consumed energy	peak power

VI. Examples of load modeling

This section provides informal examples of electric devices and applications that are suitable for being integrated in a real-time management system. Their relevant characteristics are described, outlining a possible modeling of their timing properties.

A. Household appliances

Typical household devices like ovens, washing machines, dryers, dishwashers, have each a peculiar duty cycle. The tighter the timing requirements — i.e., the closer the deadline to the maximum activation time — the more constraints are imposed on the scheduling algorithm, reducing the chances of finding a lower peak load. Anyway, a certain slack is usually available in the working cycles of household appliances, and programmable devices are already used to control the activation of electric loads depending on the energy prices in the stock market.As an example, these devices are used to control washing machines in domestic environments, where postponing by a few hours the time at which the laundry is ready does not cause any problem.

B. Temperature conditioning

The target of a HVAC is to keep the room temperature within the desired range. Therefore, heating or cooling is provided depending on the actual room temperature, which is affected by the temperature of the external environment. For a given external temperature the activation pattern of a HVAC can be approximated with a periodic activity. Appropriate load parameters must be chosen considering the thermal inertia of the system. For example, the temperature of a well insulated big environment might be controlled activating a heating (or cooling) device with a period in the order of tens of minutes. Less insulated environments might instead require a more frequent activation of the device. The duration of each activation depends on the actual power of the conditioning system.

C. Lighting

Consider the corridor lights of a building, that may need to be turned on in the evening, for example at 8:30pm, and turned off in the morning at 7:00am. In this case, no service interruption can be tolerated. During the active period, the total power consumption is the sum of power consumed by each lighting device, while in the rest of the time, the power consumption is negligible.

In this simple case, the load has a period of 24h, an active time of 10:30h, and a relative deadline equal to the active time. In this way, the load is always scheduled at the beginning of the period, without allowing any interruption while the lights are switched on, as expected in this application. Electrical loads of this kind (i.e., with no activation slack) lead to an increase in the number of concurrently active loads. Since no slack is available in the activation cycle, there is no way of reducing the impact of these loads on the resulting peak load. However, it is still possible to control the activation pattern of loads having less stringent requirements. Those loads can be activated when lights are switched off.

VII. Statistical evaluation

This section provides a statistical evaluation of benefits derived by using RTPS scheduling technique against the average behavior of load activations in absence of coordination.

A set of periodically activated electric loads are modeled as a RTPS. The considered timing parameters of the i-th load are period T_i and activation time C_i. The load utilization U_i is defined as $U_i = C_i/T_i$. The utilization is a common figure in real-time systems that represents the percentage of time in which a load is active.

We study the distribution of the total power consumption during the system lifespan. For example, the maximum power consumption occurs when all the loads are active at the same time instant. This event occurs in a certain time instant with a probability that depends on the loads' utilization. Assuming an infinite system lifespan, this probability matches the percentage of time in which the event takes place.

If load activations are independent, the occurrence of the maximum power consumption is the product of utilization of active loads. Conversely, a load activation schedule based on RTPS technique prevents the simultaneous activation of loads in the same scheduling group. In [19], authors show how to build scheduling groups based on power consumption and utilization. Loads in each group are scheduled using the EDF algorithm. The heuristic method presented in [19] will be used in this section.

Figure 6 shows the distribution of the total power consumption during the system lifespan in systems composed by three different load sets, one for each plot. Nominal parameters P and U are shown in each caption. These examples show that the RTPS scheduling method reduces the peak load while compressing the power consumption distribution around the average value.

Finally, Figure 7 shows the power distributions with and without using the scheduling techniques as a function of the number of loads. Each box-plot is the average of the power distribution of 100 randomly generated load sets. Nominal power and utilization have been generated with a uniform distribution in $[1, 10]$ and $[0.1, 0.7]$, respectively. Again, this plot confirms that RTPS technique compresses the power consumption distribution and significantly reduces the peak load.

VIII. Conclusion

This paper presented an application of Real-Time Physical Systems to the modeling of CPES. The paper foster the possibility to use real-time techniques to model energy systems in order to achieve its predictable timing behavior. The goal is to optimize a given figure; in this paper the reduction of the peak load of power consumption is addressed. This approach opens the door to further investigation of proposed methods in field such as building automation, smart grids, demand-side management and energy efficiency.

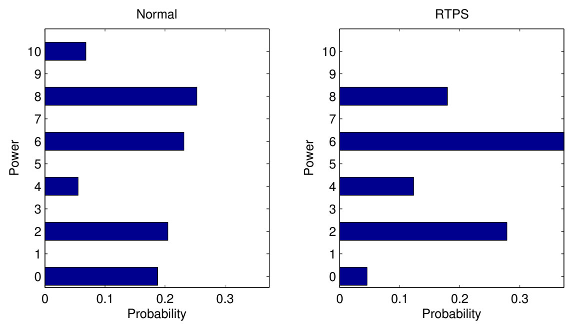

(a) Results with $n = 3$ loads. $P = (2, 2, 6)$kW, $U = (0.38, 0.32, 0.55)$.

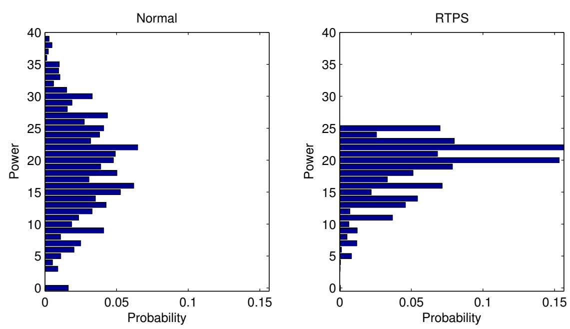

(b) Results with $n = 7$ loads. $P = (1, 3, 3, 4, 4, 5, 6, 7)$kW, $U = (.57, .38, .45, .21, .49, .14, .67, .43)$.

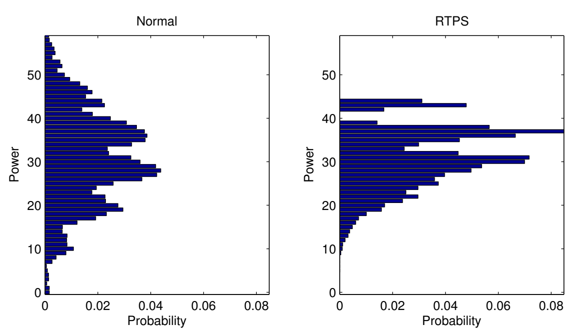

(c) Results with $n = 11$ loads. $P = (1, 2, 2, 3, 4, 4, 5, 7, 8, 8, 9)$kW, $U = (.24, .28, .56, .54, .51, .62, .30, .46, .16, .69, .44)$.

Fig. 6. Example of power distribution over time consumed by 3, 7 and 11 loads, without using RTPS technique (left) and using RTPS technique (right). The peak is reduced by 20%, 37% and 26%, respectively.

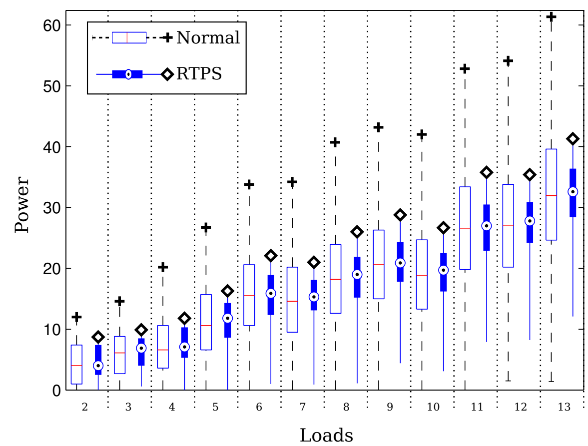

Fig. 7. Distribution of power consumption over time as a function of the number of loads. For each number of loads, the average of 10 random load sets is plotted.

REFERENCES

[1] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard real-time environment," *Journal of the Association for Computing Machinery*, vol. 20, no. 1, pp. 46–61, January 1973.

[2] P. Pedreiras, P. Gai, L. Almeida, and G. Buttazzo, "FTT-Ethernet: A flexible real-time communication protocol that supports dynamic QoS management on Ethernet-based systems," *IEEE Transactions on Industrial Informatics*, vol. 1, no. 3, August 2005.

[3] H. Aydin, R. Melhem, D. Mossé, and P. Mejía-Alvarez, "Power-aware scheduling for periodic real-time tasks," *IEEE Transactions on Computers*, vol. 53, no. 5, pp. 584–600, May 2004.

[4] H. Kim and K. Shin, "Scheduling of battery charge, discharge, and rest," in *Proceedings of the 30th IEEE Real-Time Systems Symposium (RTSS)*, December 2009.

[5] M. Bosman, V. Bakker, A. Molderink, J. Hurink, and G. Smit, "The microCHP scheduling problem," in *Proceedings of the Second Global Conference on Power Control and Optimization (PCO)*. Springer Verlag, June 2009, p. 8.

[6] M. D. Ilic, L. Xie, U. A. Khan, and J. M. F. Moura, "Modeling future cyber-physical energy systems," in *IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, July 2008, pp. 1–9.

[7] M. L. Della Vedova, M. Ruggeri, and T. Facchinetti, "On real-time physical systems," in *Proceedings of the 18th International Conference on Real-Time and Network Systems (RTNS)*, November 2010.

[8] J. A. Stankovic, I. Lee, A. Mok, and R. Rajkumar, "Opportunities and obligations for physical computing systems," *Computer*, vol. 38, no. 11, pp. 23–31, 2005.

[9] J. Lunze and F. Lamnabhi-Lagarrigue, Eds., *Handbook of Hybrid Systems Control: Handbook of Hybrid Systems Control*. Cambridge University Press, 2009.

[10] G. Buttazzo, *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Boston: Kluwer Academic Publishers, 1997.

[11] S. T. Morris, A. Srivastava, B. Reaves, S. A. K. Pavurapu, R. Vaughn, W. McGrew, and Y. Dandass, "Engineering future cyber-physical energy systems: Challenges, research needs, and roadmap," in *Proceedings of the IEEE North American Power Symposium*, October 2009.

[12] "HomePlug homepage," 2009, http://www.homeplug.org/.

[13] "Betterplace project homepage," 2009, http://www.betterplace.com/.

[14] R. Brown, "Impact of smart grid on distribution system design," in *Proceedings of the IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, July 2008, pp. 1–4.

[15] A. I. Cohen, "An optimization method for load management scheduling," *IEEE Transactions on Power Systems*, vol. 3, no. 2, May 1988.

[16] H. Hermanns and H. Wiechmann, "Future design challenges for electric energy supply," in *Proceedings of the IEEE Conference on Emerging Technologies and Factory Automation (ETFA)*, September 2009.

[17] M. A. Crew, C. S. Fernando, and P. R. Kleindorfer, "The theory of peak-load pricing: A survey," *Journal of Regulatory Economics*, vol. 8, no. 3, pp. 215–248, November 1995.

[18] D. Du Bois, "Time of use electricity billing: How puget sound energy reduced peak power demands (case study)," *Energy Priorities Magazine*, February 2006.

[19] T. Facchinetti, E. Bini, and M. Bertogna, "Reducing the peak power through real-time scheduling techniques in cyber-physical energy systems," in *Proceedings of the First International Workshop on Energy Aware Design and Analysis of Cyber Physical Systems (WEA-CPS)*, April 2010.

[20] M. Bertogna and S. Baruah, "Limited preemption edf scheduling of sporadic task systems," *IEEE Transactions on on Industrial Informatics*, vol. 6, no. 4, pp. 579–591, November 2010.