

# Management Information Base for Power Control in Computer Networks

**Roger Al-Alam Krolow   André Peres   Ingrid Jansch-Pôrto**

{roger, peres, ingrid}@inf.ufrgs.br

Federal University of Rio Grande do Sul  
Programa de Pós-Graduação em Computação (PPGC)  
Post office box 15064  
91501-970 Porto Alegre RS

## Abstract

The management of the power consumption is an important subject, which is getting attention of several segments in computers industry. Nowadays, the management is accomplished independently, in each machine, without taking into account the network requirements and the other components. This paper presents a Management Information Base (MIB) used to control the energy consumption of a computational equipment network, through the SNMP protocol (Simple Network Management Protocol). In critical situations, as when external power fails, an Uninterruptible Power Supply (UPS) is used, and the rationalization of power consumption allows a better use of the UPS batteries. During normal operation – when there is external power supply, with adequate quality–, the rationalization of the consumption results in the economy of financial resources.

The implemented MIB allows to control several equipment, as computers, hubs or switches, and it makes possible the use of the existent personal computers management standards (APM and ACPI) that allow the reduction of energy consumption without turning off the equipment. In the following, initially we introduce some considerations about the SNMP protocol and about the power management standards for personal computers, APM and ACPI. In the sequence, the MIB for power management is presented and analyzed. Some additional considerations finish the paper.

**Keywords:** MIB, SNMP, power management, protocols, network management, fault tolerance.

## 1 Introduction

Power consumption is a subject that is attracting attention in many segments of our society, as the natural resources available for power generation are limited and should be preserved. Moreover, the rationalization has essential importance in systems that operate with limited power resources, as notebooks and computer systems which are temporarily being supplied by batteries backup.

In the scenario defined for this work, a computer network is supplied by an UPS, connected to an external source of energy. When detecting external power failure (referred in this paper as power fail), the UPS provides uninterruptible energy for the system. However, depending on particular characteristics implemented by manufacturers, these equipment offer different levels of support

concerning the operational system. Some types of UPS keep energy supply until total discharge of batteries, followed by the power interruption, without previous notice to the equipment (load). This procedure is harmful, specially when these equipment are computers. Other UPS provide warnings about the remaining operation time or about the batteries imminent discharge; this allows the user to safely provide store data or the automatic shutdown of the computer, before the battery is fully discharged.

However, the operation time of the batteries can be maximized if some measures for reduction of power consumption are taken as soon as power fail is detected. Servers should be kept energized, but workstations can be forced to low consumption states or even can be turned off, after saving the necessary data and the operational context. A hierarchy of priorities can be defined for the different groups of equipment types or user classes; therefore, the shutdown can be done in a gradual and orderly way, as required in some environments. This strategy can be refined to the point of just accomplishing the immediate shutdown of the temporarily inactive stations. It may provide the shutdown of the remaining ones only if the duration of the power fail is prolonged, when the alternative power supply (battery or battery backup) of the system runs the risk of being finished. It is important to remember that the monitors are responsible for approximately half of the computer systems power consumption [NOR97]; these devices can be placed in low consumption states, without interrupting the operation of the station. Additionally, other equipment such as sound cards and printers may also be considered concerning their power consumption characteristics.

With the goal of reducing power consumption in periods of idleness of personal computers (PCs), some standards have been developed since the beginning of the 90's. The strategies which are frequently adopted are designated Advanced Power Management, or simply APM, and Advanced Configuration and Power Interface Specification, or ACPI. They specify hardware and software characteristics and also the respective interfaces, allowing the operating system (OS) to control the consumption in the hardware components.

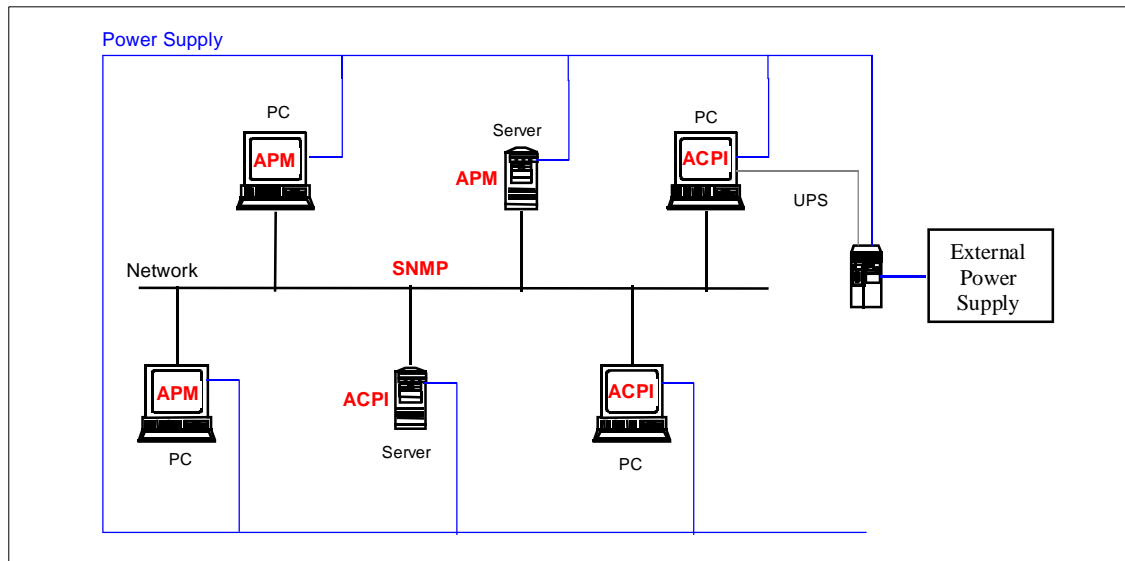
Based in these points, we propose a system of control energy, manage in conjunction of an UPS, to be used in a computer network: reduced consumption states are forced when occurs a power failure of the external supply. The result is the minimum impact in service availability and the maximization of the battery capacity. The use of native management support (APM, ACPI) joined with the SNMP protocol, which is used to obtain the information of the external power conditions and also to control equipment, is proposed to accomplish this task, as shown in figure 1.

Concerning this paper, our particular objective is to analyze the necessary information to accomplish the control of the consumption of the different kind of equipment attached to the network, through remote operations sent by a SNMP manager. From this analysis, the structure of the management information base MIB is obtained, which contains the necessary information to the control of the equipment, through an SNMP agent installed in each node of the network.

The text of the article is organized in the following way: in section 2 we present a short description of SNMP and the MIB structure. In section 3, the main standards of power management in personal computers are presented and, in section 4, these aspects are related to the control through the SNMP. The implemented MIB is presented in the section 5, including some comments on the practical results. Finally, our conclusions from this work are gathered and some perspectives are discussed.

## **2 The SNMP**

The structure of the SNMP protocol [ROS96, STA96] is based on the use of managers – programs located in a management station that is responsible for controlling and monitoring the several agents of the network – and agents – programs that contain a management information base, and are located in each node of the network to be managed.



**Figure 1 – Power Management using SNMP**

Sometimes, an agent has to accomplish the information control of a remote node that is unable to communicate through SNMP. This agent, named proxy, is responsible for translating the information between the device and the management application allowing to control and to monitor also this kind of equipment.

During the execution of its activities, the SNMP manager accomplishes basically two operations: get – reading of parameter values in agent nodes, that allows to monitor the characteristics of the node; and set – writing of parameter values in agent nodes, that results in their change. Another operation – get-next – makes easy the reading of the next value (in the sequence) of the information base.

The only way of communication started by the agent is the trap, that informs to the manager the occurrence of some special condition. This asynchronous communication is used by the agent to obtain the manager's attention. A trap does not require an answer; its arrival is neither ensured. When a trap is received, the manager has to activate the necessary interactions to the treatment of the managed node, identifying the nature and the extension of the problem.

It is important to point out that, in the SNMP protocol, during the exchanges of information with the agent, the manager does not stop its operations to wait for the answers. The SNMP has a time-out control, executed along with other management functions, that allows the detection of communication flaws among the nodes. The detection mechanism in the manager is based on the absence of answers to the messages sent.

## **2.1 Management Information**

The information kept in the agent nodes is organized in management information bases, generally referred by the acronym MIB. The MIB is defined in a structure designated Structure of Management Information, SMI [PER97].

The MIB contains the definition of the objects, that correspond to the characteristics of each resource that may interest to the management. In this way, the MIB does not contain values of instances; when the manager requests an information, the agent must accomplish the request to the node (through system calls) to transmit the corresponding value. This operation is extremely important, because the manager will always receive the updated value as the answer to a requisition.

On the other hand, it allows the agent accomplish system calls just when some information is requested by the management application. In figure 2, a description of standard MIB is presented.

The particular implementations used by enterprises, or the experimental ones, are accomplished through the definition of new objects added to a standard MIB, by one of the three following possibilities: a) expanding the definition of the MIB-2 subtree by defining a new subtree; b) an experimental MIB can be built for a private application. Once finished this implementation, the defined objects should be subsequently approved by the Internet Architecture Board – IAB [POS96] (IAB supplies a registration for each new implementation) and then be moved to the mgmt subtree; c) adding an extension in the private subtree.

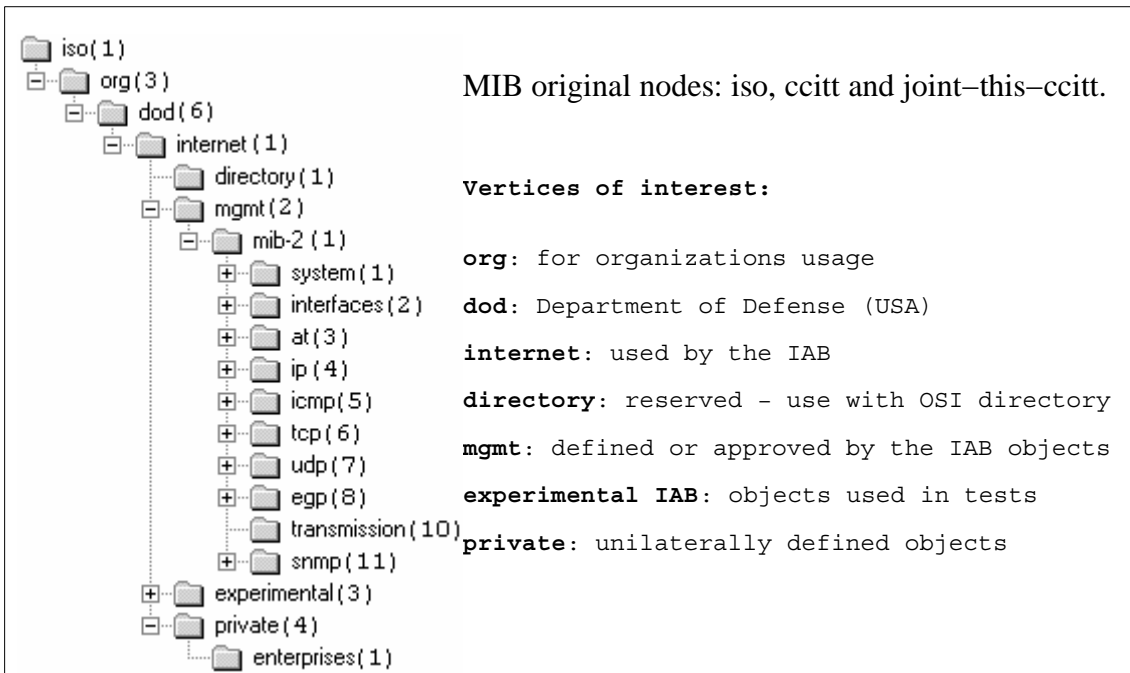


Figure 2 – MIB

### 3 The Main Standards for Power Management

Two power management standards are largely adopted by the personal computers (PC) industry; the characteristics there defined are available in the majority of the PC fabricated now.

Advanced Power Management [INT96], known by the acronym APM, was one of the first efforts of the industry of PCs for the standardization of power management. It consists of the interface specification between the layers of software – BIOS (Basic Input Output System) and of the OS – that controls the use of energy in function of the activity level of the system.

Advanced Configuration and Power Interface [INT96a], briefly referred as ACPI, describes the interface between the hardware and software for management of PCs. ACPI has been proposed to substitute APM and has taken the consumption management from BIOS to the SO. The specification covers several domains, but our description in this paper is restricted to the mechanisms for power management.

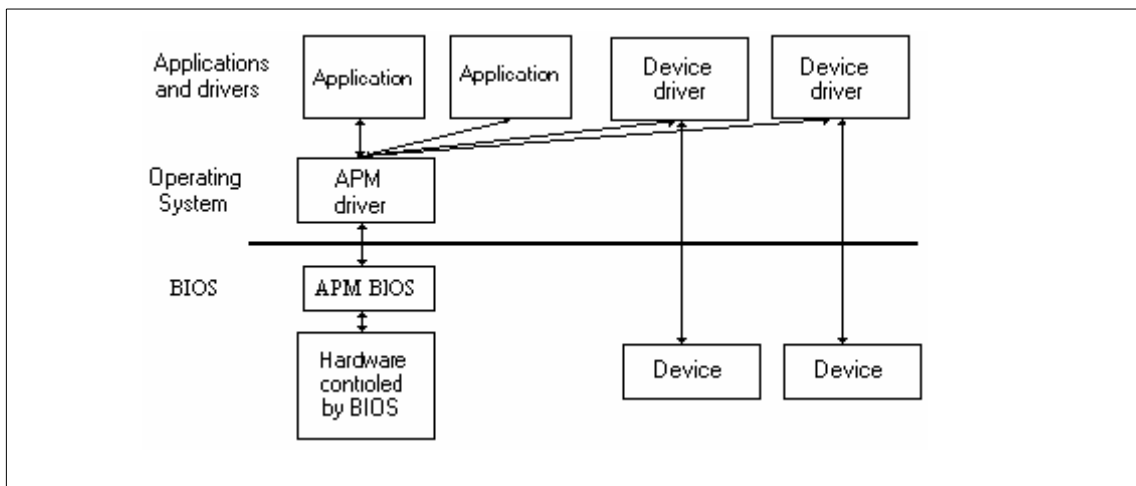
#### 3.1 The Advanced Power Management Specification

An APM system has the components illustrated in figure 3, where three different software layers are identified: BIOS; OS driver; and applications, that should cooperate to reduce the energy

consumption.

The BIOS takes care of administering the management, accomplishing it by itself when an APM driver is not available in the operating system. Each BIOS is implemented with particular properties for each platform by the machine manufacturer.

When the OS driver is loaded, it establishes a connection with the BIOS which starts to communicate the events to the OS and to execute commands sent by him. The OS driver has three main functions: to accomplish the interface between applications and BIOS; to coordinate the management strategy in multitask environments; and to identify management opportunities that are not obvious for BIOS and applications.



**Figure 3 – APM System Architecture**

After having established the connection, the OS driver requests information to BIOS in regular intervals. When BIOS reports that the system should enter in low consumption mode, the OS driver sends the information to the applications and to the device drivers to execute the appropriate procedures for saving the machine context. Just after this procedure, the BIOS is liberated to continue the operation.

The applications can contribute with the management informing to the OS about the necessary resources for the processing. In that layer, there are also drivers that realize the power management for the devices that are not directly controlled by the BIOS. Applications that participate of the power management are registered in the OS driver and notified of the management events.

The APM standard defines five power states:

- **full on:** system works normally and the devices stay energized; consumption is not reduced because APM is not active;
- **enabled:** system works normally but the processor clock is reduced or cut, when possible, causing reduction in energy consumption. Devices also can be managed from this state;
- **standby:** system stops processing and the devices are placed in states of minimum consumption. It is established after short inactivity period and reaches the whole system. The time of return to the enabled state is small;
- **suspend:** similar to standby, but the reduction in consumption is larger. This state is established when the system is maintained idle for long time. The time to return to the enabled state is great;

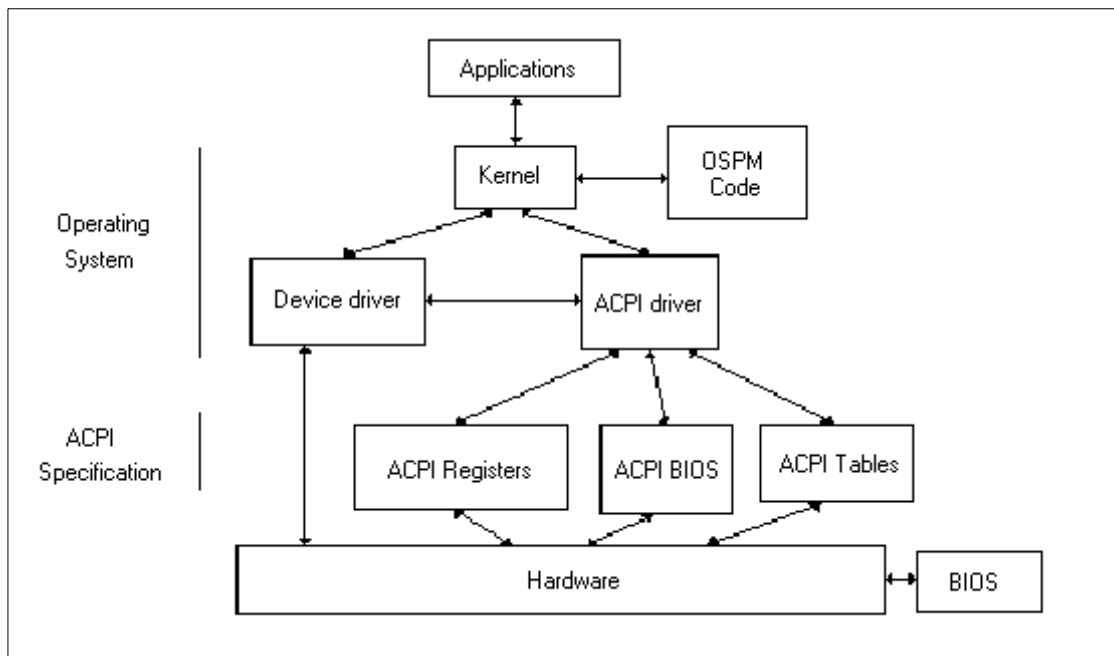
- **off**: the energy supply of the machine is interrupted. The system must be restarted to come back to the active state.

### 3.2 The Advanced Configuration and Power Interface Specification

The ACPI standard considers three basic components:

- tables: contain descriptions of the hardware and functions for energy management. The OS should possess an interpreter for those instructions;
- registers: accomplish the control of the hardware described in the tables;
- BIOS: it is the part of the firmware that implements the specification. The instructions code of BIOS makes the initialization of the machine and implements operations of hardware management.

Figure 4 shows the components of an ACPI system and illustrates the relationships among these components and others in the system.



**Figure 4 – ACPI Structure**

The ACPI standard defines four power states. When the computer is being used, it stays in the **working** state, referred by G0. When the computer is idle, it is taken to the **sleeping** state, called G1. In the following state, called **soft off** or G2, the computer consumes the smallest amount of energy. Almost all the devices are turned off, staying powered just those programmed to wake up the computer (keyboard or fax/modem card, for instance). In the **mechanical off** state, referred as G3, the machine is completely turned off.

There are similar states for devices, designated by D0, D1, D2 or D3 and for processors, identified by C0, C1, C2 or C3. The sublevels in both cases are similar to that defined for the system (G0 to G3).

While the computer is on, it is placed alternately in the states working and sleeping. In the working state, the computer has all the available functionality. Devices and processors that are not being used can be placed in low power mode (D1–D3, C1–C3); in this case, the devices can come

back to the active state in a short period of time, whose duration is variable for each device type. In the working state the computer can operate with different clock speeds to reflect the wanted or acceptable adjustment among speed, consumption, heating and noise emission.

When the computer is idle, or when the user presses the power button, the OS leads the system to a sleeping state or to the soft off state.

Soft off is the state of smaller consumption that the OS can put the equipment but, differently of a sleeping state, no context information is saved. The only powered circuits are those responsible for waking up the computer in the presence of events. The difference of that state for turned off is that the system can be wakened up by events like a phone call, a key pressed or the detection of a network packet.

### **3.2 Considerations about APM and ACPI**

The initial design of PCs (in the beginning of the 80's), has not anticipated power management properties, because there were not independent lines of supplying the devices. The APM specification, in turn, just describes software aspects, leaving the implementation of the hardware quite free. For that reason, the necessary support of hardware to the energy management was not developed. In spite of anticipating devices management, APM implementations are restricted to the management of the processor, of the hard disk and of the monitor [MIC98].

This specification leaves open the property of turning off the system by software (this is not mentioned in the standard) and it does not define wake up events. Thereby, although it has been important as the first standardized strategy of power management, APM contains restrictions to the complete power management of the machine.

ACPI is quite flexible because it extends the definition of the aspects of power management. Besides the software interface, ACPI also describes the hardware interface, in way to standardize the implementation of devices and allows the computer to be turned off by software and wakened up by events. To be conformed to this specification, the hardware architecture has to ensure individualized supply lines for each device.

Computers built according to APM can be remotely managed with some restrictions. As the specification does not determine that the hardware can be turned off by software, the system can just be placed in sleeping states (standby or preferably suspend). If the shutdown is needed, the OS may be requested to shutoff the user's applications and to take the machine to a safe state for the shutdown, that should be done manually. The turn on of the machines should be made manually because this standard does not define mechanisms for the system to wake up in the presence of events. The remote manager can also act individually on the monitor and on the hard-disk.

In APM systems, additional hardware has to be provided to accomplish the total cut of the power. It is possible to create a device, controlled by the computer, between the computer and the source of energy, capable to cut the supply after the shutdown.

The computers that implement ACPI can be managed in a broader manner then the APM computers, although the procedures are of rather similar application. In those computers, the local agent can request that the computer is placed in soft off, that presents minimum energy consumption and allows the network interface to wake up the computer when receiving a packet with this order. In this way, the manager can also control the wake up of the computers. ACPI also allows the individualized control of devices.

In practice, we find in a network, computers that do not own energy management support and others with specific characteristics for the management. For the first group, it is possible to accomplish a poor utilization, with simple states: on or off. To control these equipment, it is necessary to add hardware devices that are manageable, which make possible to wake up and

shutdown them. The equipment with specific characteristics of power management should be individually studied to identify the implemented mechanisms.

## 4 Power Management using SNMP

The system, or the network in this case, must have a software manager that gathers information about the external power supply – or about the backup alternative used under power fail or emergency conditions – received by of the equipment network. A software agent has to be available in each managed equipment; it has to be able of accomplishing two basic functions:

- to be notified by the OS about the events (mainly changing) in the state of energy of the equipment and its devices, and consequently to send these information to the remote manager using SNMP traps;
- to receive commands (get, get-next, set) from the SNMP manager and to send them to the OS to change de power state of the equipment as wanted.

The implementation of the agent involves the definition of a series of system characteristics, such as the power management standard to be used, the OS that is executed and the hardware architecture of the computer (Intel, Apple, Sun workstation, among others).

The manager should be informed about the transitions in the power state of each equipment. It allows the definition of strategies in view of the reduction in power consumption. These data may be presented with graphic support, what will turn easy the user analysis. In presence of an interruption of the external energy supply, when the network directly takes profit of the batteries through an UPS, the manager should coordinate the actions that will be accomplished by the agents to take the machines to the wanted states.

## 5 Implemented MIB

The implemented MIB has been planned to cover a varied range of computer equipment. A previous study [KRO99a] has identified two approaches of power consumption management (corresponding to the APM and ACPI standards). The great majority of personal computers offered in the commercial market ensures compatibility with one of these standards. Thereby, the groups `apm` and `acpi` were created and they describe the power state of the computer (objects `systemStatus` and `procStatus`) and of each device (objects `id`, `desc` and `status` included in the tables `devACPITable` and `devAPMTable` in the MIB). The resulting structure is shown in figure 5.

The object `systemStatus` of the group `acpi` in ASN.1 (Abstract Syntax Notation One) is described in the information base as reproduced in the following:

```
systemStatus          OBJECT-TYPE
    SYNTAX              INTEGER { G0(0), G1(1), G2(2), G3(3) }
    MAX-ACCESS          read-write
    STATUS              current
    DESCRIPTION         "equipment power state"
    ::= { acpi 1 }
```

In this way, it is possible to the manager to receive information concerning the object `systemStatus`, which describes the power state of the equipment, and to change this value by notifying the agent of this operation.

Two classes of equipment do not fit in the previous categories: the equipment with own management mechanisms, and those without support to power management.

For the first class – equipment with own management politics – we have created the `generic` group that, in addition to the previously defined objects for the groups `apm` and `acpi`, has a power



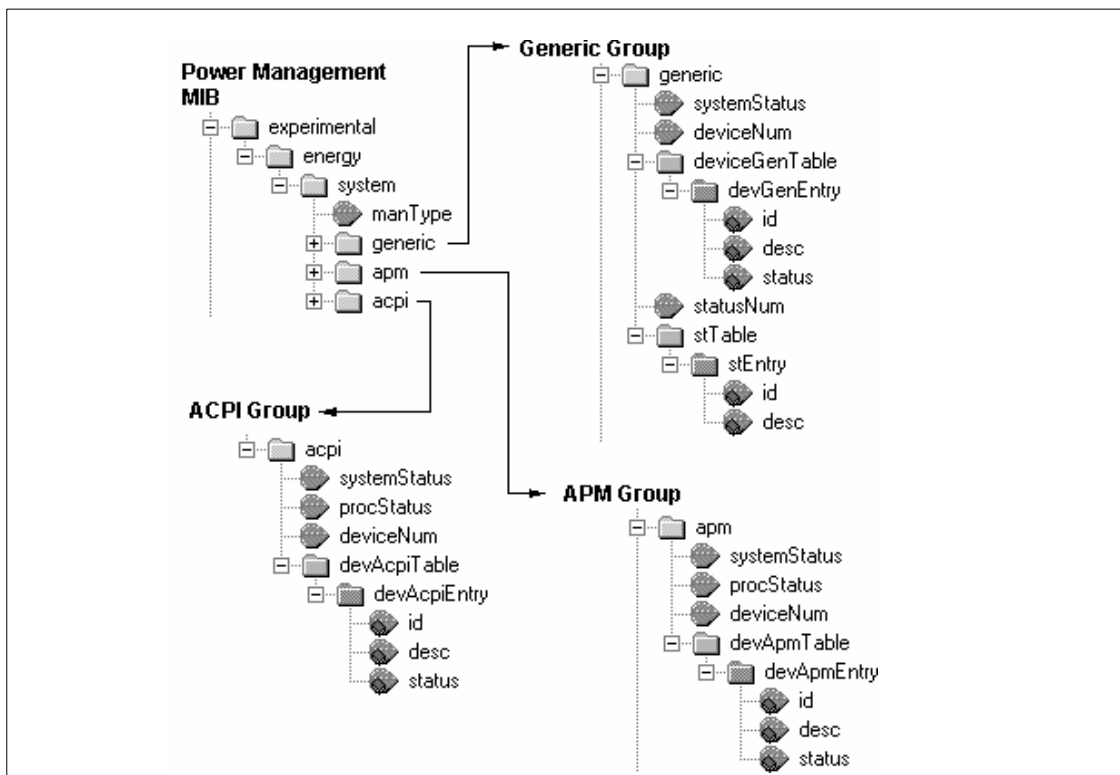
states description table (stTable).

For the second class – equipment with no support for power management – we can also use that group although with quite reduced functionality. These equipment can just indicate that they are powered on.

The information base owns an object manType to describe the type of management that is supported by the equipment. The ASN.1 description of the object is shown in the following:

```

manType
SYNTAX          OBJECT-TYPE
MAX-ACCESS     INTEGER {generic(1), apm(2), acpi(3)}
STATUS         read-only
DESCRIPTION    current
               "equipment management type"
 ::= {system 1}
  
```



**Figure 5 – The implemented MIB**


By hypothesis, an external power fail occurs in the computer network defined by figure 1. The equipment that implement ACPI and that are temporarily in low power state (without use), may be turned off immediately. The agent is responsible to inform the OS that to accomplish the safe saving of the open files). In the same way, the machines that implement APM and are in low power state can request the shutoff of the applications and of the OS (shutdown). These machines should be turned off manually.

The computers that are processing will be informed of the event and, depending on their operational priority – established individually or by classes – will start to save files and will perhaps provide the shutdown after a predefined period of time, or when the batteries reach a predefined

state. Computers that are running processes without interaction with the user can have their respective monitors turned off by the manager, resulting in smaller power consumption by the overall system.

To validate the operation of the proposed MIB, we developed an SNMP agent whose main functions are to control and to modify the equipment states. This agent has been implemented in a personal computer that uses the ACPI standard approach, under the Microsoft Windows 98 OS.

A SNMP manager, which provides the exchange of information with this agent, has been developed to monitor and to define the values modification. The manager monitors the UPS battery status through an SNMP proxy agent, as shown in figure 6.



**Figure 6 – Scenario of experiment**

The UPS informs the battery load conditions of the proxy agent (A) through the serial port. This agent implements a particular MIB, that has been proposed as a result of a study involving several UPS MIB implementations. This MIB is driven to the quality of energy management and to control UPS equipment. Additional information about this MIB may be asked directly to the authors; just a preliminary paper on it has been published [PER98]. This agent is executed on a PC computer, under Microsoft Windows NT OS; it has been implemented as an extension of the Microsoft agent [MUR98].

The SNMP manager (B) (developed for Linux OS) detects the occurrence of changes in the conditions of the power supplied by the UPS and acts in the agents of the controlled equipment. The behavior of the manager is determined by the configuration files that contain the rules which implement mechanisms corresponding to the policies for the different scenarios. This manager is capable to interact with many UPS. When detecting the occurrence of pre-defined power events, the manager sends SNMP messages to the agent of the controlled equipment, changing its state. A detailed description of the manager can be found in [KRO99b]. More details may be asked directly from the paper authors.

The power control agent (C), in accordance with the notifications of the manager, turns off its devices when requested. In our prototype, when identifying the occurrence of low battery event, the manager sends a message to the agents to turn off the computer monitors, reducing the consumption for about a half, or placing all the equipment in a low consumption state. To control the consumption of the monitor, the agent interacts directly with the BIOS through the interruption 15 service. The control of the equipment consumption is carried through the *SetSystemPowerState* Windows API System Call.

Other possibilities of energy control are being developed as, for instance, agents for equipment

with APM support and agents compatible with other OS platforms.

## 6 Conclusions

We may identify two main properties in this work: simplicity on the development of agents which are capable to interact with the OS and with the SNMP manager, without overloading the node; and also the flexibility that allows the implementation of agents for a large spectrum of equipment.

The developed MIB makes possible the construction of a centralized power management system that is able to rationalize the energy consumption. Using this power management MIB, the manager may send commands that control the power consumption in different equipment modifying the power state. The information described by the MIB also allows the creation of management applications to visualize the use and the consumption of equipment network.

Additionally, the management system may be used as much in situations of external power fail, where an UPS equipment ensures the uninterruptible power supply of the computer network, as in situations where the goal is the rationalization of energy in view of natural and financial resources economy.

An SNMP agent which is able to accomplish the operations of the information base presented in the Microsoft Windows 98 operating system is being tested. This agent accomplishes the management of machines that implement the ACPI standard.

## 7 References

- [INT96] Intel, Microsoft. Advanced Power Management (APM) BIOS Interface Specification. Version 1.2. 1996. Available from: <ftp://developer.intel.com/ial/powermgm/apmv12.pdf>.
- [INT96a] Intel, Microsoft, Toshiba. Advanced Configuration and Power Interface Specification. 1996. Available from: <http://www.teleport.com/~acpi/DOWNLOADS/ACPIsepc10.pdf>.
- [KRO99a] Krolow, R. Management of Energy in Personal Computers. PPGC–UFRGS, Technical Report, February 1999. (in Portuguese; original title: Gerência de Energia em Computadores Pessoais).
- [KRO99b] Krolow, R. Power Management System for Networked Equipment. PPGC–UFRGS, IV Semana Acadêmica, 1999. (in Portuguese; original title: Sistema de Gerência de Energia para Equipamentos em Rede).
- [MIC98] Microsoft. OnNow: Device Power Management. 1998. Available from: <http://www.microsoft.com/hwdev/desinit/DEVICEPM.HTM>.
- [MUR98] Murray, J. Windows NT SNMP. O'Reilly & Associates. California. 1998.
- [NOR97] Nordman, B. et al. User Guide to Power Management for PCs and Monitors. University of California, Berkeley, 1997. Available from: <http://eande.lbl.gov/EAP/BEA/LBLReports/39466/39466>.
- [PER97] Perkins, D.; Mcguinness, E. Understanding SNMP MIBs. Prentice Hall, 1997.
- [PER98] Peres, A. Definition and Construction of a Management Information Base for UPS Management. PPGC – UFRGS, III Semana Acadêmica, 1998. (in Portuguese, original title: Definição e Construção de uma Base de Informações de Gerência para UPS).

- [POS96] Postel, J. Internet Architecture Board. Homepage. URL: <http://www.iab.org/iab>
- [ROS96] Rose, Marshall T. The Simple Book – An Introduction to Networking Management. Revised Second Edition. Prentice–Hall Inc. Upper Saddle River, NJ. 1996.
- [STA96] Stallings, William. SNMP SNMPv2 and RMON – Practical Network Management Second Edition. Addison–Wesley, Reading, Massachusetts. 1996