

Accounting System: A fine-grained CPU resource protection mechanism for Embedded System

Midori Sugaya

Department of Computer Science
Waseda University
3-4-1 Okubo Shinjuku Tokyo, Japan
dolydcl.info.waseda.ac.jp

Shuichi Oikawa

Graduate School of Systems
Information Engineering
University of Tsukuba
1-1-1 Tennoudai Tsukuba Ibaragi, Japan
shui@cs.tsukuba.ac.jp

Tatsuo Nakajima

Department of Computer Science
Waseda University
3-4-1 Okubo Shinjuku Tokyo, Japan
tatsuo@dcl.info.waseda.ac.jp

Abstract

In ubiquitous computing environments, our daily lives will be made convenient by embedded intelligent devices. Those devices, such as car navigation systems, personal digital assistances, and cellular phones, provide various kinds of the complex services. Those devices are networked with each other and provide complicated services, through the Internet. While they provide useful services, there is an increasing possibility of security attacks, which include the unexpected execution of un-secure codes. Current information appliances have not yet fully embodied a resource protection mechanism that prevents misbehaved applications from consuming the whole CPU capacity of system resources.

In this paper, we propose Accounting System, and describe its design and implementation. The system is a resource monitoring and restriction system that has the purpose of improving the system's reliability and security. We developed the system on Linux. Our system is a very generic to offer various services, such as security improvement, overload control, and class-based accounting, that require CPU resource control.

1. Introduction

Information appliances [11] are important elements to realize the ubiquitous computing vision[6]. Nowadays, most consumer electronics appliances have computing ca-

pability in order to retrieve data from sensors, to process the data, and to control devices. The recent emergence of information appliances requires more advanced features, such as networking and GUI. Those features dramatically complicate the appliances, software systems and increase their code sizes. Networked systems need to be prepared for security attacks through the Internet. Since we expect users to be system administrators of applications, their software systems must be more robust than personal computers.

Software bugs can also cause the monopolization of CPU resources. Real-time operating systems may be vulnerable to bugs of real-time programs that easily consume the whole CPU resources. For example, a problem can happen when multimedia applications process continuous media streams. Those multimedia applications have strict timing constraints, and are given the real-time priority in general. Therefore, once they process the streams, they can easily monopolize the whole CPU resources. In such a situation, the GUI process that is usually executed on the time-sharing scheduler cannot consume necessary CPU capacity, so that a user cannot control the multimedia applications through the GUI buttons. Another example is an overload condition. If an overload condition occurs, the response time of the system becomes worse. General embedded operating systems have not yet provided a resource protection mechanism that aims to protect the CPU resources from the CPU monopolization and overload situations. Therefore, it is necessary to offer a generic mechanism to restrict the use of CPU resources to develop reliable information appliances.

In this paper, we propose Accounting System, a general-purpose resource monitoring and restriction system that prevents the excessive use of the CPU capacity of a process or a group of processes. The following two points are our design principles.

- *Simple design to be applied to various services.* Accounting System focuses on providing a simple and generic model and interface. Therefore, the system should be able to be easily applied to various kinds of services, such as security enhancement, overload monitoring, class-based accounting, and processor reservation[3].
- *Accurate resource management by using a fine-grain resolution timer.* Future information appliances require to support a fine-grained rate based execution that can be realized by the accurate resource management, for offering better response time and more stable execution.

Information appliances have become very complex. Most of them nowadays contain web browsers, Java Virtual Machines, and many other applications. Implementing these applications in a robust way needs a more powerful operating system. Linux is an open-source operating system, and supports various CPU architectures. Those features are very suitable for information appliances. Therefore, many industries consider adopting Linux for their products. There are many products that have already adopted Linux. Sony Co-coon Channel Server, Philips iPronto Remote Control, and NTT Docomo mobile phones, are the examples.

While the use of Linux is increasing, several problems are recognized. For example, Linux does not restrict the resource consumption for their processes. Thus, when malicious application programs are downloaded and executed, the programs may consume a large amount of the CPU capacity easily. To solve this problem, we propose Accounting System and implement it in the Linux kernel. The remaining of the paper is structured as follows. Section 2 describes the design and model of Accounting System. Section 3 presents the architecture and implementation of Accounting System. Section 4 shows the services that can be implemented by Accounting System. Section 5 shows the evaluations of Accounting System. Section 6 describes related work. Finally, section 7 concludes the paper.

2. Accounting System Design

2.1. Design Issues

We designed Accounting System to take into the account the following three issues.

Simplicity: The system should be simple and generic to

be used in a variety of OS services such as security enhancement, class-based accounting, overload monitoring, and processor reservation.

Accuracy : The system should monitor the CPU capacity of each process very accurately for making the execution of application more stable. A fine-grained timer is used to realize the accurate monitoring.

Portability : The system should be implemented in a variety of operating systems. The system confines the interface to a few hooks in the host kernel.

2.2. Model of Accounting System

In this section, we describe the basic concept and model of Accounting System.

2.2.1 Accounting System and Accounting Object

Accounting System is a system to offer the functions that monitor and restrict the execution time of a process or a group of processes according to user policies. We have developed it as a part of the kernel services. Accounting System provides an abstraction to manipulate the CPU capacity of each process. We named the abstraction an *accounting object*. It is a basic abstraction in Accounting System. An accounting object represents a capacity of the CPU resource on a single processor. The CPU time of each process bound to an accounting object is accumulated in it. Through an accounting object, a user can control the CPU resources. When a user process binds itself to an accounting object through the system call, the process is given the restricted execution time specified by its accounting object.

As shown in Figure 1, application A is bound to an accounting object that has 20% restriction of the CPU capacity. On the other hand, application B and C are bound to the other accounting object that has the 40% restriction. The application, bound to the specific accounting object, is not permitted to use more than the CPU capacity specified in the accounting object.




Figure 1. Accounting Object and Process




Figure 2. Time Management by Accounting Object

2.2.2 Time Management Model

To realize the restriction, an accounting object has two parameters C and T , where T represents a period that is a constant period to control the object, and C is the maximum time to be able to execute processes within T . The process bound to the accounting object cannot consume the CPU time more than C within T . When processes consume the entire CPU time within each period T , the process is blocked until the next period comes. This model has the following advantages.

- The C and T parameters are defined for respective accounting objects. This allows developers who design embedded systems to adapt the various kinds of applications that have different requirements. If the period time is unified through all accounting objects, all applications should be subject to restrict the same timing constraint.
- This model is independent of scheduling policies. Even if operating systems have several scheduling policies, they can adopt our system to control and monitor the CPU resources of the processes. Therefore, our system can be applied for the various types of the applications that run on the different scheduling policies supported by operating systems.

The two parameters, maximum computation time C , and period time T , are set to each accounting object. Figure 2 illustrates an example of the binding a real-time process to the accounting object. Let us consider the case where a user creates a new accounting object whose T and C are 100ms and 50ms respectively. If a process bound to the accounting object is started at the time 0, it monitors the execution time of the process to the extent time of C . If the execution time C of the process reaches to its maximum of 50ms, the process is blocked until the next period comes. At the end of T , which equals to the start time of the next period, Accounting System wakes up the process. Such actions are repeated until a process completes its execution. By this way, Accounting System monitors and controls the execution time of the processes by using accounting objects.

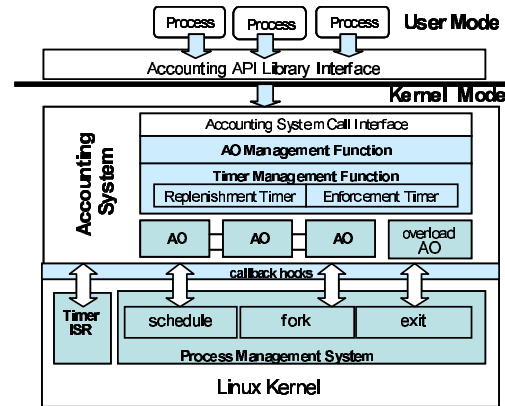


Figure 3. Accounting System Architecture

3. Implementation

3.1. System Architecture

Figure 3 depicts the overall architecture of Accounting System. The services are available by enabling the configuration of the Accounting System support at the kernel compiling time.

We consider that the architecture is simple in two points. The first, developers who design embedded systems can easily understand the abstraction as accounting object. By using it, they can easily decide the CPU utilization of the process and the group of processes by setting the two parameters, period and execution time. The second, Accounting System architecture is separated from the host kernel. Only four hooks are inserted in the host kernel. Therefore, any scheduling policies implemented in the host kernel can be used with Accounting System.

Accounting System is composed of the mechanisms of Accounting Object Management Function and Timer Management Function to restrict the CPU utilization of processes according to the restrictions of respective accounting objects. Those mechanisms are decoupled from the scheduling policies of the host kernel.

3.2. Process Management and Callback Hooks

In the system, callback hooks are introduced into the process management system in the host kernel for the purpose of catching events about processes. Those hooks catch the relevant events, and deliver those events to Accounting System. Accounting System assumes that the four callback hooks as described the following are inserted in the host kernel.

- *schedule_hook* : It catches scheduling events on context switching among processes. The hook delivers the scheduling events to Accounting System that changes the state of an accounting object.
- *fork_hook* : It catches events when a new process is created. These events are required to start the accounting for the newly created processes. If a parent process has already been bound to an accounting object, its child processes inherited the same accounting restriction.
- *exit_hook* : It catches events on process termination. The events notify the timing to unbind the termination of a process from its accounting object.
- *kernel_callback_hook* : It catches events when a system call returns from the kernel mode to the user mode. At that time, the system checks the state of the accounting object for the blocking of the process to execute the system call.

3.3. Timer Management Function

The management of timers has an important role of the system. In this system, there are two types of timers, a replenishment timer and an enforcement timer. A replenishment timer is a periodic timer for each accounting object that is constantly expired. An enforcement timer keeps to watch on the elapsed time of the currently executing process. They are implemented by using kernel software timers. In the following paragraph, we describe the details of them.

- *Replenishment timer* : It manages the periodic time of an accounting object. Each accounting object has its own replenishment timer. When an accounting object is created, the timer is also created. In the replenishment timer, first the expiration time is set by adding the periodic time to the current time. If the elapsed time reaches the expiration time, the registered timer handler is executed to set the next expiration time. These replenishments are recursively done by the replenishment timer until it is destroyed.
- *Enforcement timer* : It watches the elapsed time of the currently executing process. The timer accumulates the execution time of the process or processes, until it reaches to the maximum execution time of the accounting object. At that time, the registered timer handler is invoked to change the state of the accounting object. By using *schedule_hook*, the enforce timer catches the events about context switching among processes, and stops and starts its timer to monitor the CPU capacity of the process or processes.

3.4. Accounting Object Management Function

To manage accounting objects, the system provides the accounting object management function. It introduces the following accounting object's basic operations. To receive the accounting object services, an application firstly invokes the *create* operation to allocate the accounting object in the kernel memory. Each created accounting object has a unique object ID. To bind a process to the specified accounting object with the object ID, the application invokes the *bind* operation. When the application no longer requires the accounting services, it invokes the *destroy* operation to terminate the accounting object service and to deallocate the object memory from the kernel memory. The *set* operation changes the parameters of an accounting object. The *get* operation obtains the parameter in an accounting object.

3.5. Object State Transition

There are three states to represent an accounting object's status. *null* status represents the status that an accounting object is created, but no process is bound to the accounting object. *running* status represents the accounting object is running, and the status also represents that available time are left within the period time. *depleted* status represents the status that no more available time is left within the period time. The *depleted* status also represents that the accumulated execution time of a process is reached to the maximum execution time specified in the accounting object. The state of an accounting object is changed by various events. When a process binds to a specific accounting object, the state of an accounting object becomes *running*. The enforcement timer accumulates the execution time of the bound process to the accounting object by using *schedule_hook*. If the sum of the execution time of the bound processes equals to the maximum execution time of the accounting object, the object changes its state to *depleted*. The change is notified to Accounting System through *kernel_callback_hook*. Accounting System executes the action to control the CPU capacity of the bound processes. The action is classified into *block* and *signal*.

If the action is *block*, Accounting System puts the bound processes in the wait queue. The replenish timer checks the wait queue when the periodic timer expires. If they find the processes in the wait queue, *running* is set to the state of the accounting object, and the processes are woken up. On the other hand, if the action specified as *signal*, a signal is delivered to a process specified in the accounting object as a warning that the execution time of the accounting object exceeds the restricted capacity.

3.6. High Resolution Timer

For Accounting System, the high-precision clock and time stamp counter are essential to maintain the fidelity of its accuracy. To provide less than 10ms resolution for monitoring, there are some approaches. The most straightforward way is to make the clock interrupt frequency higher than 100Hz in order to issue more interrupts. This allows the kernel to provide a precise timing and to increase the system responsiveness. However, shorter intervals require CPU to spend a longer time in the kernel mode. Therefore, user programs run slower because of the system overhead to handle interrupts.

Our system solves the problem by using the high resolution timer[5]. The high-resolution timer is implemented by using the one shot mode of the ISA clock timer chip in the standard PC. Therefore, the extra timer interrupt overhead will never happen. The high-resolution timer allows us to specify microsecond granularity parameters in accounting objects.

4. Some Services Using Accounting System

In this section, we describe four services that utilize Accounting System, such as a secure resource protection, class-based accounting, overload monitoring, and processor reservation.

4.1. Avoiding to monopolize CPU resources by malicious applications

4.1.1 Secure Resource Protection

Accounting System protects the CPU resource from downloaded programs that behave maliciously. We assume that there is a manager process that starts processes to execute downloaded programs. The manager process is bound to an accounting object, of which C and T are 50ms and 100ms, respectively. When the manager process receives a request to execute a downloaded application, it creates a new process. The process is automatically bound to the accounting object to which the manager process is bound. Therefore, the newly created process cannot consume more than 50% of CPU capacity, and it cannot monopolize the CPU resource.

4.1.2 Binding Multiple processes

Accounting System provides the *bind* operation that binds a process to the accounting object. We also provide some useful interfaces to bind multiple processes to an accounting object. For example, if a malicious program consecutively invokes the *fork* system call, the whole system resources

may be consumed. In another case, it should be easy for multimedia applications, which process the video and audio data by multiple processes, to bind their processes to a specific accounting object at once.

We implement the following two methods. The first is to inherit an accounting object when creating a child process. The other is to bind an accounting object using group ID. The former function is realized by using a hook within the *fork* system call. If a parent process is bound to an accounting object, the child process who is forked from the parent inherits the same accounting object. Thus, the same restriction as its parent is applied. The latter method uses UNIX group process ID. The system binds the group of processes to a specific accounting object. It is safer than binding all processes to the same accounting object respectively. For example, when several commands connected by pipes are started, they can be easily bound to the same accounting object by using group ID.

4.1.3 Access Control

Each accounting object has an owner attribute for controlling the access to a specific accounting object. The owner attribute of the accounting object is assigned by a process that creates it. For the consideration of the security, our system allows them to be manipulated only by their owners except the privileged user. In resources reservation systems such as Linux/RK[10], there is no support of access control. Therefore, these previous system cannot be used for protecting the CPU resource from malicious programs.

4.1.4 Kernel Interface

This section shows the kernel interface offered by Accounting System. An application program needs to treat the accounting object as the *object_attributes* structure that contains basic parameters, such as C and T . We described the details of them.

- *account_create(&object_id,&object_attr)* : It creates a new accounting object. Only a privileged or owner application program sets the parameters, such as C , T , in the *object_attr* members.
- *account_destroy(&object_id)* :It destroys the accounting object specified by *object_id*. This deletes the associated timers, and frees the memory of the accounting object.
- *account_bind_pid(object_id, pid_t)* : It binds the specified accounting object through *object_id* to a process whose process ID is *pid*. An accounting object can be bound to multiple processes. If the process is terminated, the process is automatically unbound from the accounting object.

- *account_bind_pgid(object_id, pgid_t)* : It binds all processes that have the same pgid to the specified accounting object. If there is a process, which is bound to the other accounting object, this system call fails.
- *account_unbind_pid(object_id, pid_t)* : It unbinds a process from the accounting object. This is used for both of single process and a group of processes.
- *account_get(object_id, &object_attr)*: It retrieves the parameters of the specified accounting object.
- *account_set(object_id, &object_attr)* :It changes the parameters of the specified accounting object.

4.2. Class-based Accounting

In this section, we propose a class-based accounting as the one of the services that uses Accounting System. A class is defined here as a group of processes that belong to a specific scheduling class, such as the real-time scheduling class. The class-based accounting is realized by binding the processes belonging to the same scheduling class.

4.2.1 A Problem of Process Starvation

When a user forgets to bind a real-time process to an accounting object, the real-time process may monopolize the entire CPU resource. Even if most real-time processes are restricted, the only one real-time process can consume the rest of all the CPU resources; thus, all other time-sharing processes are starved. To avoid such a situation, the system designer should estimate the total maximum utilization of the accounting objects, which bind to the all real-time processes. We called the approach class-based accounting.

4.2.2 Class-based Accounting

The class-based accounting is realized by binding processes in the specific scheduling class to an accounting object. A class contains all processes that are scheduled by the same scheduling policy.

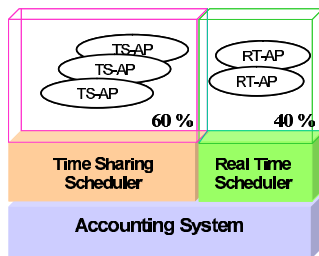


Figure 4. Class-based Accounting

The system allocates the CPU resources proportionally to each scheduling class. The total proportions of the classes are set to 100%. For example, a user can allocate 40% of the CPU capacity for the real-time scheduling class, and the remaining 60% of the CPU capacity for the time-sharing scheduling class as shown in Figure 4. The system assumes to have two classes that are necessary to allocate the CPU capacity. Each scheduling class is assigned the absolute rate of the CPU capacity. The processes bound to each class cannot excessively use the maximum utilization of the CPU resources that are set in each scheduling class. Since the real-time process cannot use more than the assigned capacity, the time-sharing process can use the remaining capacity without process starvation.

4.2.3 Implementation

The current class-based accounting allocates the CPU capacity for the real-time and time-sharing scheduling class. The accounting object for each scheduling class is created at the boot time. When it is created, all processes that belong to the same scheduling class are bound to the same accounting object. In order to set the parameter easily, class-based accounting adopts weight as its interface parameters. The weight is the parameter that represents a proportion and can easily be specified like (2,3). If weight (2,3) given, the real-time class and the time-sharing class are allocated 40% and 60% capacity of the CPU resource respectively. In class-based accounting, the accounting object of each scheduling class should use the same period time. Also, a user needs to set the resolution time that is a period for the accounting objects. For example, if the resolution is 100ms and the class proportion is 40%, the class is received 40ms within the 100ms period.

4.2.4 Kernel Interface

We design the kernel interface for class-based accounting as following.

- *weight_set(wt_ts, wt_rt, res)* : The routine sets the weight and resolution to the accounting object for respective scheduling classes. When the class accounting object is started, it monitors the execution time of the bound processes.

4.3. Overload Monitoring

4.3.1 Conceiving the overload situation

For general-purpose operating systems, controlling overload situations is very important to offer services in a stable way. In such critical situations, the computational requests from the process will exceed the time that is available in the

system. As for time-sharing processes, the average response time is increasing. We propose the monitoring function that keeps to watch the total utilization of the CPU. If the overload condition happens, the overload monitoring issues a signal to notify to the registered administrator process. If there is a notification for the administrator process, some recovery procedures can be taken to stabilize the system.

4.3.2 Design and Implementation

The overload monitoring can also be realized by Accounting System. In this system, an accounting object bound to the idle process is used to monitor the total utilization of the system resources. The idle process is executed when there is no runnable process. If the system becomes busy, the idle process consumes a very little CPU capacity. Through monitoring the CPU usage of the idle processes, the system can detect the overload situation. Our system allows an application to receive a signal when the CPU usage of the idle process becomes lower than the specified CPU.

The service is available when a system designer creates an overload accounting object, and sets the parameters to it. There is the only one overload accounting object in a system. Therefore, it has a special object_id whose number is 1. The number is reserved and never used by the other accounting objects. A system designer should explicitly set the *C* and *T* parameter to the overload accounting object. Suppose the designer sets *C* and *T* to 75ms and 100ms respectively. The overload monitoring monitors the CPU utilization not to exceed 75%. If the designer would like to check the overload condition every minute, *C* and *T* should be set to 45sec and 60sec respectively. If the total usage of CPU utilization exceeds 75%, an overload notification signal is delivered to a specific process. A system designer also should set a process as a receiver of the signal. The signal number can be freely decided by a user. In consideration of the security, only a privileged user can create and destroy the overload accounting object. The overload monitoring offers the kernel interface as following.

- *overload_create (&object_attr, pid)* : The function binds an idle process to the specific accounting object. At the same time, the user should set an administrator process ID that accepts a signal from the system, and set a signal number that is used to deliver the signal to the administrator process when the overload condition occurs.
- *overload_destroy(void)* : It destroys the overload accounting object.

4.4. Resource Reservation

On application running on our system can request the reservation of a certain amount of CPU resources. The OS-

enforced resource reservation can be introduced similar to CPU capacity reserves [3][10], which provide the framework for managing the processor capacity. Accounting System can implement the framework easily by adding a mechanism for the admission control as a user-level service.

5. Evaluation

We have evaluated Accounting System by running several benchmarks. The evaluation uses a standard PC that has Celeron 300MHz CPU and 512MB of RAM, and the results were measured using the built-in high-resolution time stamp counter.

5.1. Costs of Basic Operations

The evaluation shows the cost of the basic operations in Accounting System. Table 1 shows the average cost for invoking each kernel function.

Table 1. Costs of Basic Operations

API	Costs(usec)
create	39.3
bind	21.1
unbind	6.2
destroy	10.8
set	10.4
get	2.4




Figure 5. With/Without High Resolution Timer Support

The result shows the creation time is longer than the other system function's cost. The creation cost includes the

operation of allocating a memory region and creating dynamic timers. The bind function also includes the searching cost through all process in the system.

5.2. Accuracy of Accounting System

This section evaluates the accuracy of Accounting System. To evaluate the effectiveness of high-resolution timer support, we use the benchmark program that executes the infinite-loop. We create two accounting objects that set 20ms and 25ms of CPU time every 100ms, and bind the program to each accounting object. Figure 5 shows the result of the consumed CPU ratio with/without the support of a high-resolution timer. In this figure, the bottom two lines are the results of the high-resolution timer support. It shows that each 20% and 25% resource restriction is correctly limited.

In the contrast, without the high-resolution timer support, the program bound to the accounting object whose execution time is 25ms executes 30ms within 100ms. These are apparently due to the fact that Linux updates its internal timer-clock at the 10ms intervals. Even if a user sets a parameter whose resolution is less than 10ms, it is impossible to ensure the resolution. The results shows that the effectiveness of the accurate resource accounting.

5.3. Effectiveness of CPU Protection

This section evaluates the effectiveness of the CPU resource protection in Accounting System. First, we show the effect of the security attacks, then, show the result of the class-based accounting.

5.3.1 Attacks to create multiple processes

Figure 6 shows the consumed CPU ratio of multiple processes bound to an accounting object. To evaluate the effectiveness for protecting from attacks to create multiple processes, we run an evil program that invokes `fork()` consecutively and produces a lot of child processes that execute their infinite-loops. The response time becomes very bad by this malicious program because these processes consume the whole CPU resource. In the evaluation, if the first process is bound to the 30% restricted accounting object, the forked processes cannot use more than 30% of the CPU capacity, even if the numbers of processes are increasing. The result shows the effectiveness of Accounting System as shown in Figure 6.

The loss of the accuracy appears as the numbers of created process is increased. In the current implementation, a process that executes a system call is blocked after the system call is returned, when the execution time of the accounting object is expired. If there are multiple processes




Figure 6. The Numbers of Process and Accounting Object Ratio

that are bound to the same accounting object, and they execute system calls, the system calls can be executed before blocking themselves. This is a reason to exceed the specified execution time when the number of process bound to the same accounting object is increased.

5.3.2 Class-based Accounting

Figure 7 and Figure 8 show the results of the evaluation of class-based accounting. To evaluate the effectiveness of the class-based accounting, the benchmark runs a real-time process to execute an infinite loop with/without the class-based accounting. Figure 7 shows the result when the class-based accounting is not used. The result shows that all CPU resources are consumed by the real-time process.

Then, we evaluate the class-based accounting in which each class is set the weight(4,1) where 80% and 20% CPU capacity are allocated for real-time processes and time-sharing processes respectively.

The Xserver (X Window System display server) process runs on the time-sharing scheduler. On Xserver, some game programs are running. Figure 8 shows the result that the real-time processes can not exceed 80% CPU resources, where the Xserver process can use 20% CPU resource at a maximum.

6. Related Work

UNIX offers the system call `setrlimit` to control the maximum resource consumption. They can set the resource limit to a process. However, the resolution is more than one second. The resolution is too coarse for information appliances that need to process data within a short time.




Figure 7. Without Class-Based Accounting

Linux/RK[10] proposes the resource reservation that reserves resources, such as CPU, disk and network capacity for executing real-time applications. The resource reservation prevents a system from being overloaded by controlling the admission of new processes. Also, it monitors that a process does not consume more CPU capacity than the specified CPU usage. However, they did not have any supports of using this system for increasing security. Accounting System takes the advantage of the reservation system to separate resources with each process, and develops to provide more general functions to restrict the process execution.

The proportional fair scheduler[1][2] allocates CPU capacity to each process fairly. The allocation of the CPU capacity is controlled in a rigorous way. Therefore, each process can consume the same amount of CPU resource. Each process can assign *weight* to allocate the different amount of CPU resource. However, the approach is difficult to group processes to allocate CPU capacity. To solve this problem, the hierarchical proportional fair share scheduling [7][8] is proposed. In the approach the higher level multiple schedulers run on the lower level proportional fair scheduler. Each higher-level scheduler consumes the proportional ratio of CPU capacity. The proportional fair scheduler is very promising, however it offers the lack of accuracy and it offers different semantics from the original Linux scheduling semantics. Some Linux applications assume the Linux's original scheduling semantics.

A resource container [4] is a model for fine-grained resource management mechanism. The resources, such as CPU and sockets are contained in respective resource containers. However, the resource container does not offer the accurate resource accounting. Therefore, it is difficult to achieve the fairness in the short time interval.




Figure 8. With Class-Based Accounting

7. Conclusion

In this paper, we have proposed Accounting System. The system provides the simple and generic accounting model that can be applied various services that are used to increase the system security and reliability. It also supports a fine-grained resource accounting with the high-resolution timer. This makes the system more stable because the application can obtain the necessary rate precisely even in the microsecond resolution. The performance evaluation results have showed that the system effectively protects the malicious use of CPU resources.

References

- [1] A.K. Parekh and R.G.Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case. *IEEE/ACM Trans. Networking*, Vol. 1, No. 3, pp. 344-357, June 1993.
- [2] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queuing Algorithm. *Journal of Internet working Research and Experience*, pp.3-26, October 1990.
- [3] Clifford W. Mercer, Stefan Savage, and Hideyuki Tokuda. Processor capacity reserves: Operating system support for multimedia applications. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, pages 90-99, May 1994.
- [4] G. Banga, P. Druschel, and J. C. Mogul. Resource containers: A new facility for resource management in server systems. In *Proceedings of the 3rd USENIX*

Symposium on Operating Systems Design and Implementation, Feb. 1999.

- [5] <http://high-res-timers.sourceforge.net/>
- [6] Mark Weiser, The Computer for the 21st Century, Scientific American, Vol 265, No.3, 1991.
- [7] P. Goyal, X. Guo, and H. M. Vin. A hierarchical CPU scheduler for multimedia operating systems. In Proceedings of the Second Symposium on Operating Systems Design and Implementation (OSDI). USENIX, October 1996.
- [8] P. Goyal, H. M. Vin, and H. Chen. Start-time fair queuing: A scheduling algorithm for integrated services packet switching network. IEEE/ACM Transactions on Networking, Vol 5 No.5, pp690-704, October. 1997.
- [9] Sally Floyd, Jacobson, V., Link-sharing and Resource Management Models for Packet Networks. IEEE/ACM Transactions on Networking, Vol. 3 No. 4, pp. 365-386, August 1995.
- [10] Shuichi Oikawa and Ragnathan Rajkumar, Portable RK: A portable resource kernel for guaranteed and enforced timing behavior. In Proceeding of IEEE RTAS, 1999.
- [11] W. P. Sharpe, S. P. Stenton, Information Appliances, The Human-Computer Interaction Handbook : Fundamentals, Evolving Technologies and Emerging Applications - Human Factors and Ergonomics, Chapter37, Lawrence Erlbaum Assoc Inc Published 2002.