generating a game on the gaming machine may be tested where a maximum resource utilization allowed for the one or more ECI processes is simulated while the game is being executed. The game may be tested under a variety of operational conditions, such as when it is using a maximum number of CPU cycles or graphic processor cycles, to ensure that the generated game is adequate at the maximum resource utilization condition allowed for the one or more ECI processes. After the testing, it may be concluded that the game performance will be adequate for any combination of one or more ECI processes using up to the maximum allowable resources for the ECIs. Thus, new ECI processes may be developed after the game is released without having to test the performance of the game in combination with each new ECI.

[0204] In addition, each ECI process may be tested to determine whether they perform adequately under various resource conditions up to the maximum resources allowed for a single ECI on a gaming device. This process may allow ECI developers to develop and test ECIs and associated content that are appropriate for different resource ranges up to the maximum allowed resources without needing to test them in combination with each possible game. Further, the developer may develop multiple ECIs and associated content to perform a particular function using different amount of resources with the knowledge that each ECI will perform adequately after testing. For example, a first ECI may use vector graphics to provide an animation, which requires less memory and allows for a faster download time, as compared to a second ECI that uses pre-rendered bitmaps to provide the animation where the function of the first and second ECI are the same.

[0205] As described above, in regards to virtualization, the present invention is not limited to resource partitioning. Other examples of virtualization that may be employed in embodiments of the present invention are described as follows. Via Intel's Virtualization Technology (or the corresponding AMD technology), these microprocessor vendors have introduced features in their micro-architectures that may improve the processor's ability to run multiple operating systems and applications as independent virtual machines. Using this virtualization technology, one computer system can appear to be multiple "virtual" systems. Thus, in various embodiments, a gaming environment utilizing virtual gaming machines where the operating systems may vary from virtual gaming machine to virtual gaming machine may be employed. In a particular embodiment, a virtual gaming machine may use a core of a multi-core processor.

[0206] A virtual gaming machine may use a virtual machine monitor (VMM) A virtual machine monitor may be a host program that allows a single computer to support multiple, identical execution environments. All the users may see their systems as self-contained computers isolated from other users, even though every user is served by the same machine. In this context, a virtual machine may be an operating system (OS) that may be managed by an underlying control program.

[0207] Low interrupt latency, direct access to specialized I/O, and the assurance that a VMM won't "time slice away" the determinism and priority of real-time tasks may be important for a real-time virtual gaming machine used in a gaming environment. In one embodiment of the present invention, the combination of multi-core CPUs and Intel VT or a related technology may be used to build a real-time hypervisor based on dynamic virtualization.

[0208] A real-time hypervisor may be a VMM that uses hardware virtualization technology to isolate and simultaneously host general-purpose operating systems and real-time operating systems. Unlike a static virtualization, the dynamic virtualization implemented by a real-time hypervisor may use an "early start" technique, to take control of the hardware platform. Thus, operating systems may only be allowed to "boot" only after the real-time hypervisor has constructed a virtual machine for them. The guest operating system may be associated with a particular game provided by a software provider. Thus, in the present invention, a gaming platform may support games provided by multiple software vendors where different games may be compatible with different operating systems.

[0209] In the processors that include Intel VT an overarching operating-mode has been added, called VMX root, where a hypervisor executes with final control of the CPU hardware. A hypervisor that uses Intel VT may intercept key supervisor-mode operations executed by any software operating outside of VMX root without requiring a prior knowledge of the guest OS binaries or internals. Using this Intel VT hardware assist for virtualization, one may build a hypervisor VMM that hosts protected-mode operating systems executing in ring 0 without giving up control of key CPU resources. Also, Intel VT provides a way for the VMM to implement virtual interrupts.

[0210] In the present invention, static and dynamic virtualization may be used. Nevertheless, two advantages to building a multi-OS real-time system by using dynamic virtualization rather than static virtualization may be: first, a wide range of operating systems, both general-purpose and real-time, may be supported and, second, the boot sequence for each guest OS may be under the control of the hypervisor. The second advantage means it may possible, in embodiments of the present invention, to restart one guest OS while other guest operating systems continue to run without interruption.

[0211] TenAsys provides an example of a hypervisor that may be used in embodiments of the present invention. The hypervisor may be capable of supporting the demands of a Real-time operating system (RTOS) while simultaneously hosting a general-purpose operating system (GPOS), like Windows or Linux. The hypervisor may enhance real-time application responsiveness and reliability in a "multi-OS, single-platform" environment, by providing control over interrupt latency and partitioning of I/O resources between multiple guest operating systems.

[0212] In various embodiments, the hypervisor may be used to distinguish between resources that may be multiplexed by the VMM and those that are exclusive to a virtual machine. For example, when user interface I/O is not associated with time-critical events, input devices like the keyboard, mouse, console, disk, and an enterprise Ethernet interface may be multiplexed and shared between all virtual machines. However, hardware that is specific to a real-time control application, such as a video capture card, fieldbus interface, or an Ethernet NIC designated for communication with real-time I/O devices, may not be multiplexed between virtual machines. Using the hypervisor, specialized real-time I/O may be dedicated to its real-time virtual machine, so the RTOS and application using that I/O can maintain real-time determinism and control.

[0213] In one embodiment of a VMM some or all of the memory in each virtual machine may be swapped to disk, in order to more efficiently allocate limited physical RAM