

EXHIBIT D



US005915131A

United States Patent [19]

Knight et al.

[11] **Patent Number:** **5,915,131**

[45] **Date of Patent:** **Jun. 22, 1999**

[54] **METHOD AND APPARATUS FOR HANDLING I/O REQUESTS UTILIZING SEPARATE PROGRAMMING INTERFACES TO ACCESS SEPARATE I/O SERVICES**

[75] **Inventors:** **Holly N. Knight**, La Honda; **Carl D. Sutton**, Palo Alto; **Wayne N. Meretsky**, Los Altos; **Alan B. Mimms**, San Jose, all of Calif.

[73] **Assignee:** **Apple Computer, Inc.**, Cupertino, Calif.

[21] **Appl. No.:** **08/435,677**

[22] **Filed:** **May 5, 1995**

[51] **Int. Cl.⁶** **G06F 9/40**; **G06F 13/14**

[52] **U.S. Cl.** **395/892**; **395/682**; **395/828**; **395/702**; **707/104**; **345/333**

[58] **Field of Search** **395/828**, **702**, **395/834**, **200.2**, **892**, **682**, **309**; **345/333**; **707/104**

[56] **References Cited**

U.S. PATENT DOCUMENTS

4,593,352	6/1986	Castel et al.	364/200
4,727,537	2/1988	Nichols	370/85
4,908,859	3/1990	Bennett et al.	380/10
4,982,325	1/1991	Tignor et al.	364/200
5,129,086	7/1992	Coyle, Jr. et al.	395/650
5,148,527	9/1992	Basso et al.	395/325
5,197,143	3/1993	Lary et al.	395/425
5,430,845	7/1995	Rimmer et al.	395/275
5,491,813	2/1996	Bondy et al.	395/500
5,513,365	4/1996	Cook et al.	395/800
5,535,416	7/1996	Fecney et al.	395/834
5,537,466	7/1996	Taylor et al.	379/201

5,553,245	9/1996	Su et al.	395/284
5,572,675	11/1996	Bergler	395/200.2

OTHER PUBLICATIONS

Forin, A., et al. entitled "An I/O System for Mach 3.0," Proceedings of the Usenix Mach Symposium 20-22, Nov. 1991, Monterey, CA, US, 20-22 Nov. 1991, pp. 163-176. Steve Lemon and Kennan Rossi, entitled "An Object Oriented Device Driver Model," Digest of Papers Comcon '95, Technologies for the Information Superhighway 5-9, Mar. 1995, San Francisco, CA, USA pp. 360-366.

Glenn Andert, entitled "Object Frameworks in the Taligent OS," Intellectual Leverage: Digest of Papers of the Spring Computer SOCI International Conference (Comcon), San Francisco, Feb. 28-Mar. 4, 1994, Feb. 24, 1994, Institute of Electrical and Electronics Engineers, pp. 112-121.

Hu, 'Interconnecting electronic mail networks: Gateways and translation strategies are proposed for backbone networks to interchange incompatible electronic documents on multivendor networks', Data Communications, p. 128, vol. 17, No. 10, Sep. 1988

Knibbe, 'IETF's Resource Reservation Protocol to facilitate mixed voice, data, and video nets', Network World, p. 51, Apr. 24, 1995.

Primary Examiner—Thomas C. Lee

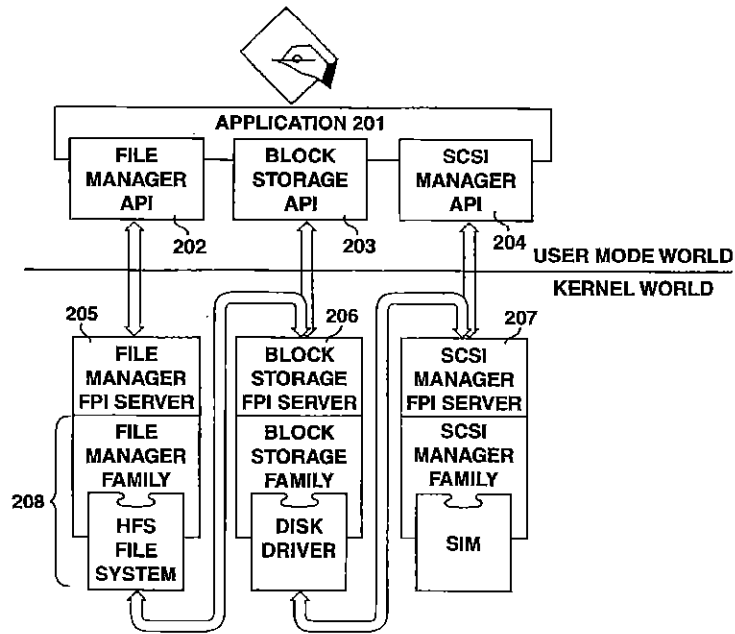
Assistant Examiner—Rehana Perveen

Attorney, Agent, or Firm—Blakely, Sokoloff, Taylor & Zafman

[57] **ABSTRACT**

A computer system handling multiple applications wherein groups of I/O services are accessible through separate application programming interfaces. Each application has multiple application programming interfaces by which to access different families of I/O services, such as I/O devices.

20 Claims, 8 Drawing Sheets






FIG. 1

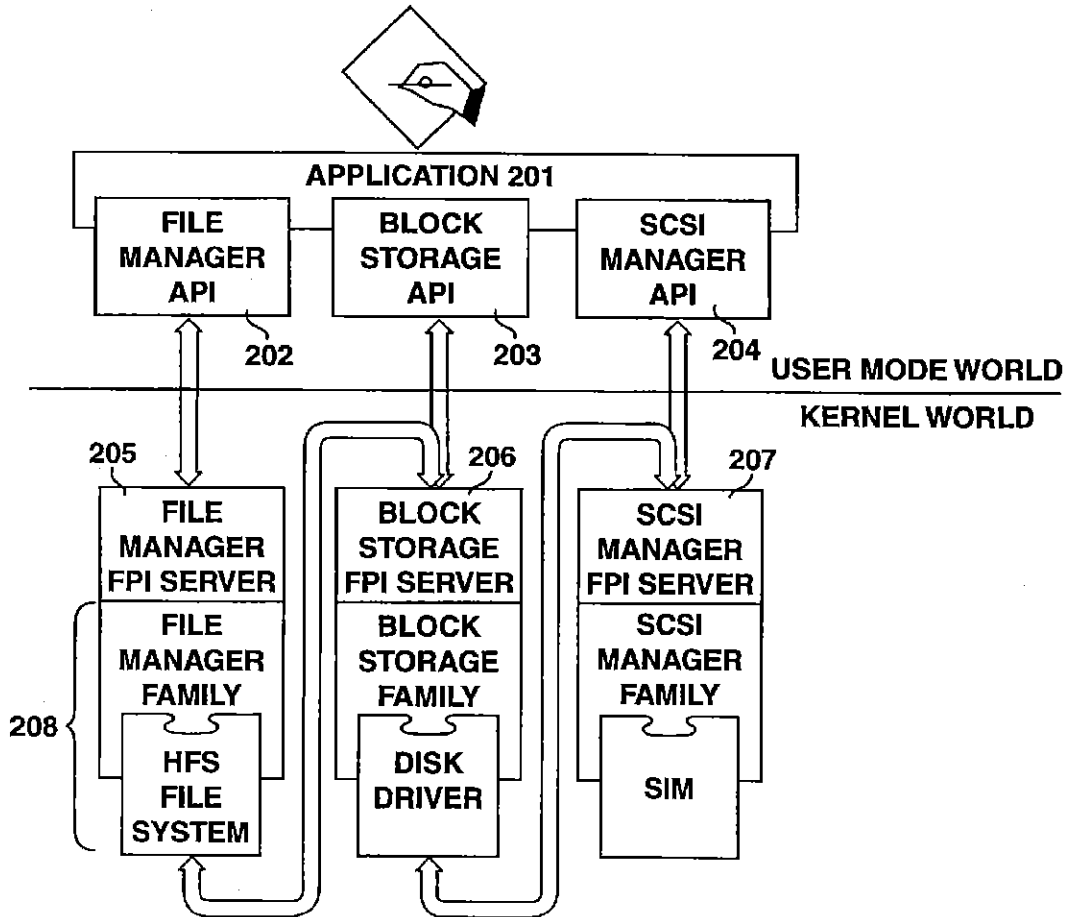


FIG. 2

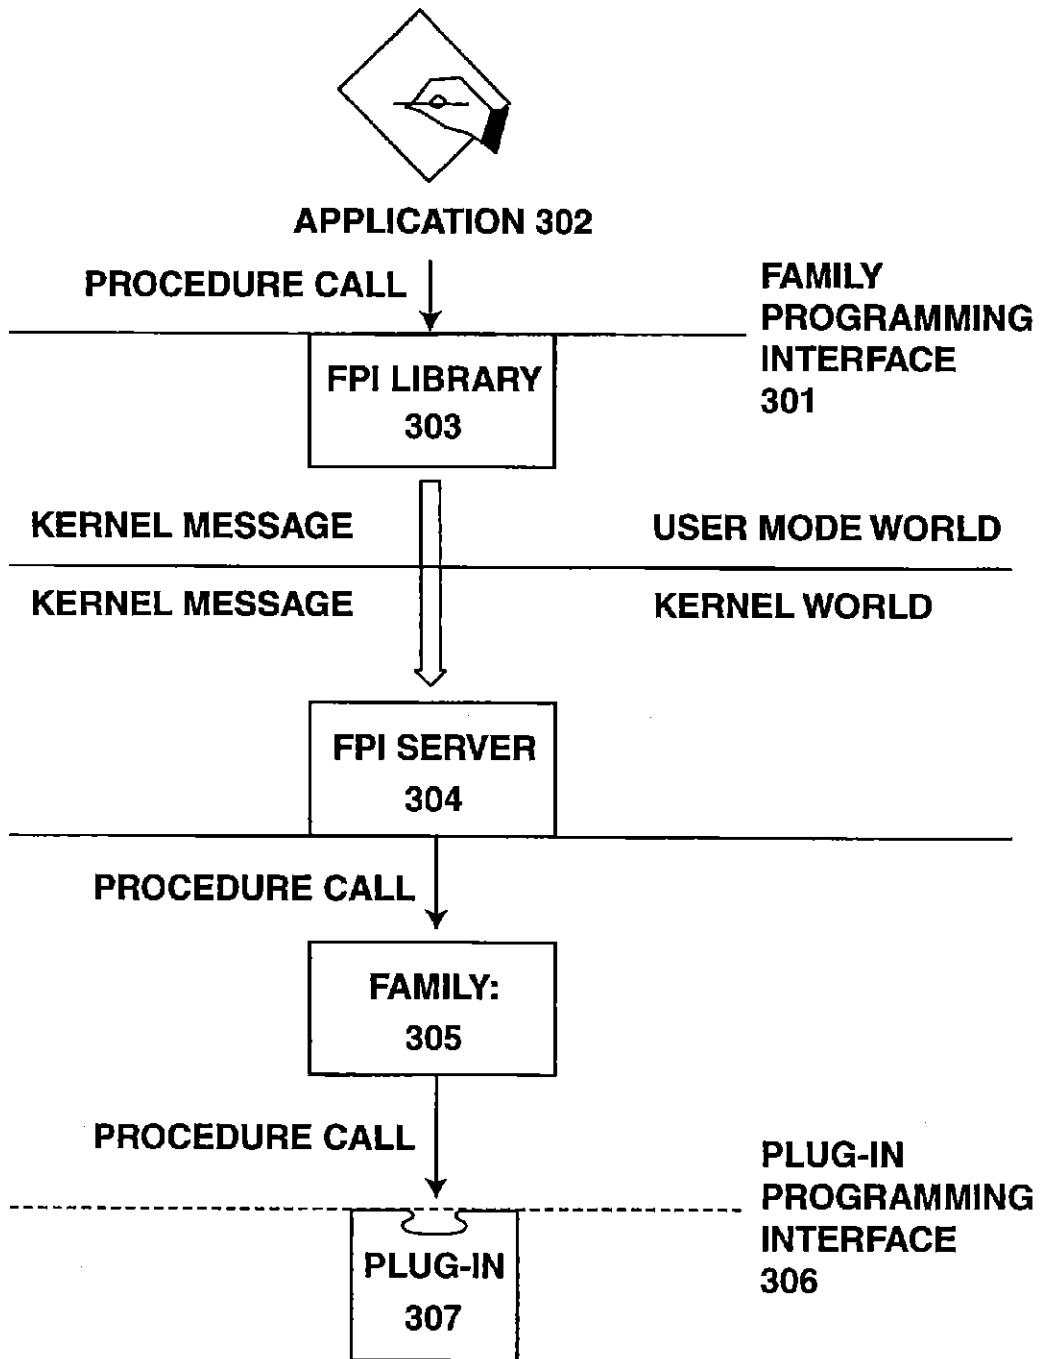


FIG. 3

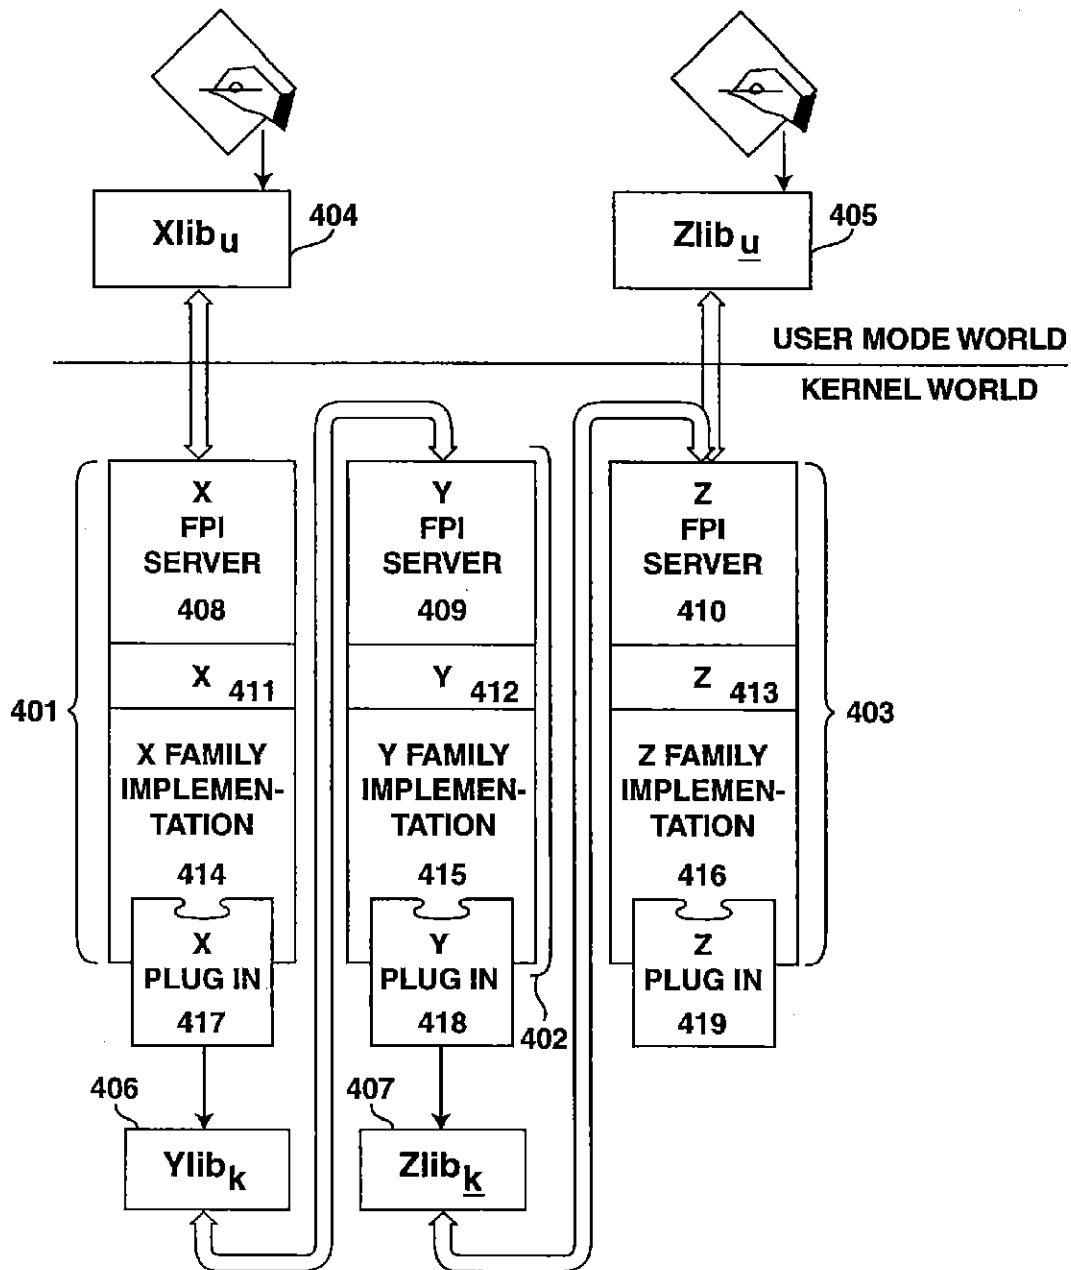


FIG. 4

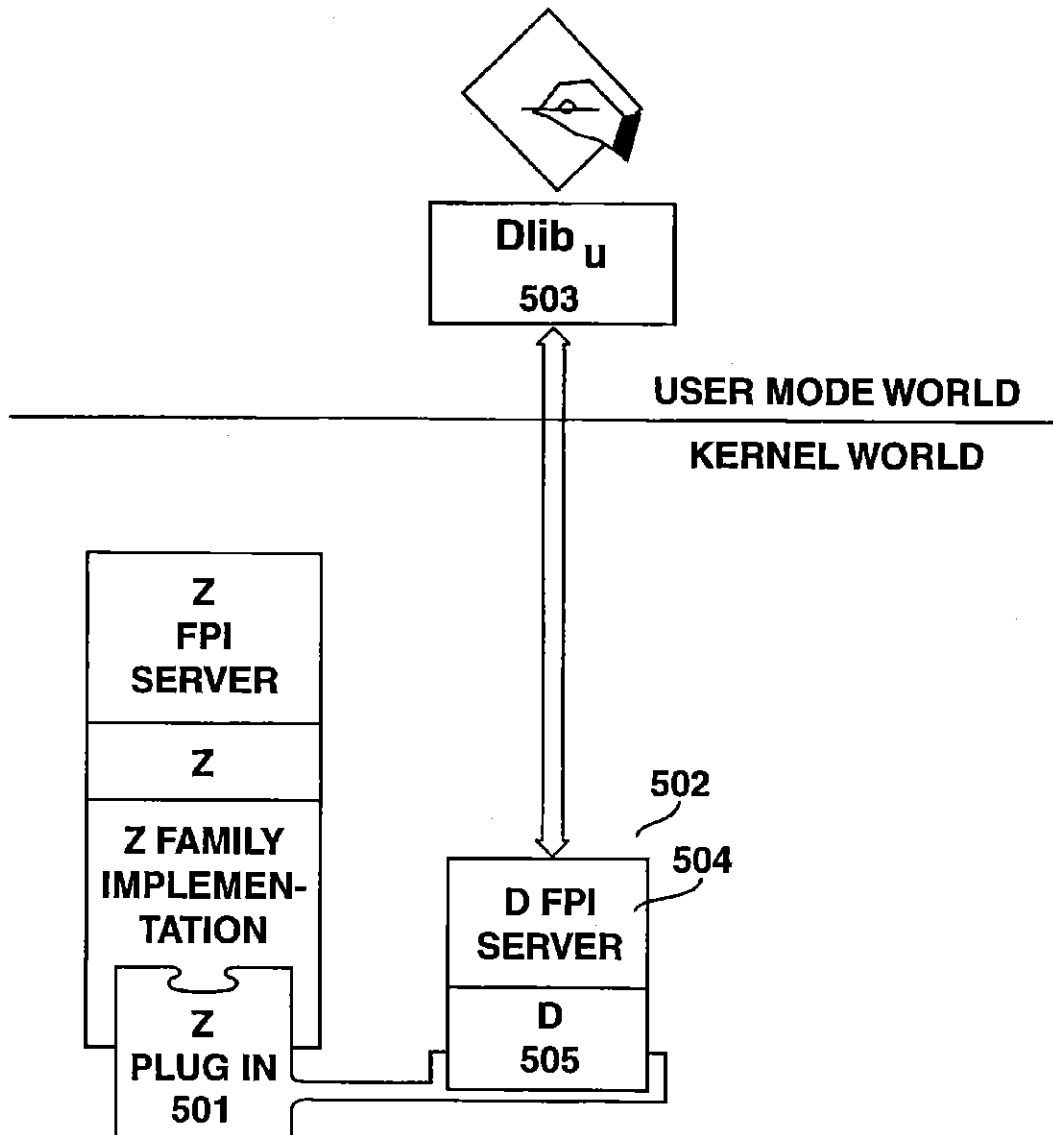


FIG. 5




FIG. 6

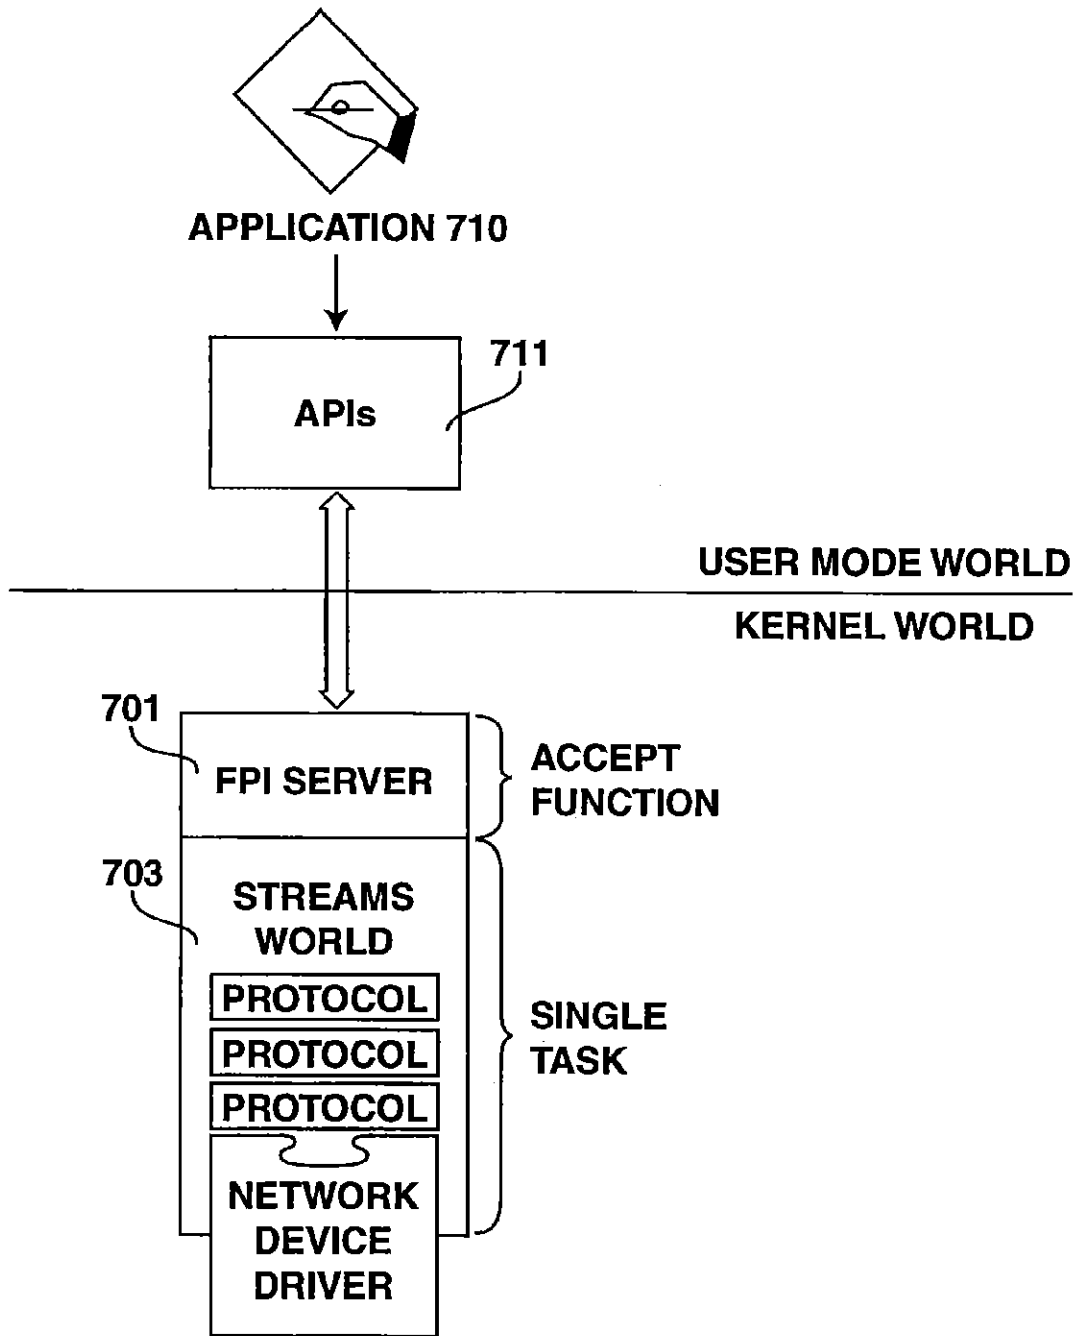


FIG. 7




FIG. 8

**METHOD AND APPARATUS FOR
HANDLING I/O REQUESTS UTILIZING
SEPARATE PROGRAMMING INTERFACES
TO ACCESS SEPARATE I/O SERVICES**

FIELD OF THE INVENTION

The invention relates to the field of computer systems; particularly, the present invention relates to handling service requests generated by application programs.

BACKGROUND OF THE INVENTION

Application programs running in computer systems often access system resources, such as input/output (I/O) devices. These system resources are often referred to as services. Certain sets of services (e.g., devices) have similar characteristics. For instance, all display devices or all ADB devices have similar interface requirements.

To gain access to I/O resources, applications generate service requests to which are sent through an application programming interface (API). The service requests are converted by the API to a common set of functions that are forwarded to the operating system to be serviced. The operating system then sees that service requests are responded to by the appropriate resources (e.g., device). For instance, the operating system may direct a request to a device driver.

One problem in the prior art is that service requests are not sent directly to the I/O device or resource. All service requests from all applications are typically sent through the same API. Because of this, all of the requests are converted into a common set of functions. These common set of functions do not have meaning for all the various types of I/O devices. For instance, a high level request to play a sound may be converted into a write function to a sound device. However, the write function is not the best method of communicating sound data to the sound device. Thus, another conversion of write data to a sound data format may be required. Also, some functions do not have a one-to-one correspondence with the function set of some I/O devices. Thus, it would be desirable to avoid this added complexity and to take advantage of the similar characteristics of classes of I/O devices when handling I/O requests, while providing services and an environment in which to run those services that is tuned to the specific device needs and requirements.

SUMMARY OF THE INVENTION

A method and apparatus for handling I/O requests is described. In the present invention, the I/O requests are handled by the computer system having a bus and a memory coupled to the bus that stores data and programming instructions. The programming instructions include application programs and an operating system. A processing unit is coupled to the bus and runs the operating system and application programs by executing programming instructions. Each application programs have multiple separate programming interfaces available to access multiple sets of I/O services provided through the operating system via service requests.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

FIG. 1 a block diagram of one embodiment in the computer system of the present invention.

FIG. 2 is an overview of the I/O architecture of the present invention.

FIG. 3 illustrates a flow diagram of I/O service request handling according to the teachings of the present invention.

FIG. 4 illustrates an overview of the I/O architecture of the present invention having selected families accessing other families.

FIG. 5 illustrates extended programming family interface of the present invention.

FIG. 6 illustrates plug-in modules of different families that share code and/or data.

FIG. 7 illustrates a single task activation model according to the teachings of the present invention.

FIG. 8 illustrates a task-per-plug-in model used as an activation model according to the teachings of the present invention.

**DETAILED DESCRIPTION OF THE PRESENT
INVENTION**

A method and apparatus handling service requests is described. In the following detailed description of the present invention numerous specific details are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may

comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose machines may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these machines will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

Overview of the Computer System of the Present Invention

Referring to FIG. 1, an overview of a computer system of the present invention is shown in block diagram form. The present invention may be implemented on a general purpose microcomputer, such as one of the members of the Apple family of personal computers, one of the members of the IBM personal computer family, or one of several other computer devices which are presently commercially available. Of course, the present invention may also be implemented on a multi-user system while encountering all of the costs, speed, and function advantages and disadvantages available with these machines.

As illustrated in FIG. 1, the computer system of the present invention generally comprises a local bus or other communication means **100** for communicating information, a processor **103** coupled with local bus **100** for processing information, a random access memory (RAM) or other dynamic storage device **104** (commonly referred to as a main memory) coupled with local bus **100** for storing information and instructions for processor **103**, and a read-only memory (ROM) or other non-volatile storage device **106** coupled with local bus **100** for storing non-volatile information and instructions for processor **103**.

The computer system of the present invention also includes an input/output (I/O) bus or other communication means **101** for communication information in the computer system. A data storage device **107**, such as a magnetic tape and disk drive, including its associated controller circuitry, is coupled to I/O bus **101** for storing information and instructions. A display device **121**, such as a cathode ray tube, liquid crystal display, etc., including its associated controller circuitry, is also coupled to I/O bus **101** for displaying information to the computer user, as well as a hard copy device **124**, such as a plotter or printer, including its associated controller circuitry for providing a visual representation of the computer images. Hard copy device **124** is coupled with processor **103**, main memory **104**, non-volatile memory **106** and mass storage device **107** through I/O bus **101** and bus translator/interface unit **140**. A modem **108** and an ethernet local area network **109** are also coupled to I/O bus **101**.

Bus interface unit **140** is coupled to local bus **100** and I/O bus **101**, and acts as a gateway between processor **103** and the I/O subsystem. Bus interface unit **140** may also provide translation between signals being sent from units on one of the buses to units on the other bus to allow local bus **100** and I/O bus **101** to co-operate as a single bus.

An I/O controller **130** is coupled to I/O bus **101** and controls access to certain I/O peripherals in the computer system. For instance, I/O controller **130** is coupled to controller device **127** that controls access to an alphanumeric input device **122** including alpha-numeric and other keys, etc., for communicating information and command

selections to processor **103**, and a cursor control **123**, such as a trackball, stylus, mouse, or trackpad, etc., for controlling cursor movement. The system also includes a sound chip **125** coupled to I/O controller **130** for providing audio recording and play back. Sound chip **125** may include a sound circuit and its driver which are used to generate various audio signals from the computer system. I/O controller **130** may also provide access to a floppy disk and driver **126**. The processor **103** controls I/O controller **130** with its peripherals by sending commands to I/O controller **130** via local bus **100**, interface unit **140** and I/O bus **101**.

Batteries or other power supply **152** may also be included to provide power necessary to run the various peripherals and integrated circuits in the computer system. Power supply **152** is typically a DC power source that provides a constant DC power to various units, particularly processor **103**. Various units such as processor **103**, display **121**, etc., also receive clocking signals to synchronize operations within the computer systems. These clocking signals may be provided by a global clock generator or multiple clock generators, each dedicated to a portion of the computer system. Such a clock generator is shown as clock generator **160**. In one embodiment, clock generator **160** comprise a phase-locked loop (PLL) that provides clocking signals to processor **103**.

I/O controller **140** includes control logic to coordinate the thermal management. Several additional devices are included within the computer system to operate with the control logic within I/O controller **140**. A timer **150**, a switch **153** and a decoder **154** are included to function in connection with the control logic. In one embodiment, decoder **154** is included within bus interface unit **140** and timer **150** is included in I/O controller **130**.

Switch **153** is a p-channel power MOSFET, which has its gate connected to the power signal **182**, its source to the power supply and its drain to processor's V_{DD} pin.

In one embodiment, processor **103** is a member of the PowerPC™ family of processors, such as those manufactured by Motorola Corporation of Schaumburg, Ill. The memory in the computer system is initialized to store the operating system as well as other programs, such as file directory routines and application programs, and data inputted from I/O controller **130**. In one embodiment, the operating system is stored in ROM **106**, while RAM **104** is utilized as the internal memory for the computer system for accessing data and application programs. Processor **103** accesses memory in the computer system via an address bus within bus **100**. Commands in connection with the operation of memory in the computer system are also sent from the processor to the memory using bus **100**. Bus **100** also includes a bi-directional data bus to communicate data in response to the commands provided by processor **103** under the control of the operating system running on it.

Of course, certain implementations and uses of the present invention may neither require nor include all of the above components. For example, in certain implementations a keyboard or cursor control device for inputting information to the system may not be required. In other implementations, it may not be required to provide a display device displaying information. Furthermore, the computer system may include additional processing units.

The operating system running on processor **103** takes care of basic tasks such as starting the system, handling interrupts, moving data to and from memory **104** and peripheral devices via input/output interface unit **140**, and managing the memory space in memory **104**. In order to take care of such operations, the operating system provides

multiple execution environments at different levels (e.g., task level, interrupt level, etc.). Tasks and execution environments are known in the art.

Overview of the Present Invention

In one embodiment, the computer system runs a kernel-based, preemptive, multitasking operation system in which applications and I/O services, such as drivers, operate in separate protection domains (e.g., the user and kernel domains, respectively). The user domain does not have direct access to data of the kernel domain, while the kernel domain can access data in the user domain.

The computer system of the present invention uses one or more separate families to provide I/O services to the system. Each I/O family provides a set of I/O services to the system. For instance, a SCSI family and its SCSI interface modules (SIMs) provide SCSI based services, while a file systems family and its installable file systems provide file management services. In one embodiment, an I/O family is implemented by multiple modules and software routines.

Each family defines a family programming interface (FPI) designed to meet the particular needs of that family. An FPI provides access to a given family's plug-ins, which are dynamically loaded pieces of software that each provide an instance of the service provided by a family. For example, within the file systems family (File Manager), a plug-in implements file-system-specific services. In one embodiment, plug-ins are a superset of device drivers, such that all drivers are plug-ins, but not all plug-ins are drivers.

Access to services is available only through an I/O family's programming interface. In one embodiment, hardware is not directly accessible to application software, nor is it vulnerable to application error. Applications have access to hardware services only through an I/O family's programming interface. Also, the context within which an I/O service runs and the method by which it interacts with the system is defined by the I/O family to which it belongs.

FIG. 2 illustrates the relationship between an application, several I/O families, and their plug-ins. Referring to FIG. 2, an application 201 requests services through one or more family FPIs, shown in FIG. 2 as File Manager API 202, Block Storage API 203, and SCSI Manager API 204. The File Manager API 202, Block Storage API 203, and SCSI Manager API 204 are available to one or more applications in the user domain.

In one embodiment, the service requests from application 201 (and other applications) are sent through File Manager API 202, Block Storage API 203, and/or SCSI Manager API 204, etc., and flow as messages to family FPI servers 205-207, which reside in the kernel domain. In one embodiment, the messages are delivered using a kernel-supplied messaging service.

Any communication method may be used to communicate service requests to I/O families. In one embodiment, kernel messaging is used between the FPI libraries and the FPI server for a given family, between different families, and between plug-ins of one family and another family. The communication method used should be completely opaque to a client requesting a family service.

Each of the FPI servers 205-207 permit access to a distinct set of services. For example, File Manager FPI server 205 handles service for the file manager family of services. Similarly, the Block Storage FPI server 206 handles service requests for the block storage family of services.

Note that FIG. 2 shows three families linked by kernel messages. Messages flow from application level through a family to another family, and so on. For instance, a service

request may be communicated from application level to the file system family, resulting in one or more requests to the block storage family, and finally one or more to the SCSI family to complete a service request. Note that in one embodiment, there is no hierarchical relationship among families; all families are peers of each other.

Families in the Present Invention

A family provides a distinct set of services to the system. For example, one family may provide network services, while another provides access to a variety of block storage mediums. A family is associated with a set of devices that have similar characteristics, such as all display devices or all ADB devices.

In one embodiment, each family is implemented in software that runs in the computer system with applications. A family comprises software that includes a family programming interface and its associated FPI library or libraries for its clients, an FPI server, an activation model, a family expert, a plug-in programming interface for its plug-ins, and a family services library for its plug-ins.

FIG. 3 illustrates the interaction between these components. Referring to FIG. 3, a family programming interface (FPI) 301 provides access to the family's services to one or more applications, such as application 302. The FPI 301 also provides access to plug-ins from other families and to system software. That is, an FPI is designed to provide callers with services appropriate to a particular family, whether those calls originate from in the user domain or the operating system domain.

For example, when an application generates data for a video device, a display FPI tailored to the needs of video devices is used to gain access to display services. Likewise, when an application desires to input or output sound data, the application gains access to a sound family of services through an FPI. Therefore, the present invention provides family programming interfaces tailored to the needs of specific device families.

Service requests from application 302 (or other applications) are made through an FPI library 303. In one embodiment, the FPI library 303 contains code that passes requests for service to the family FPI server 304. In one embodiment, the FPI library 303 maps FPI function calls into messages (e.g., kernel messages) and sends them to the FPI server 304 of the family for servicing. In one embodiment, a family 305 may provide two versions of its FPI library 303, one that runs in the user domain and one that runs in the operating system kernel domain.

In one embodiment, FPI server 304 runs in the kernel domain and responds to service requests from family clients (e.g., applications, other families, etc.). FPI server 304 responds to a request according to the activation model (not shown) of the family 305. In one embodiment, the activation model comprises code that provides the runtime environment of the family and its plug-ins. For instance, FPI server 304 may put a request in a queue or may call a plug-in directly to service the request. As shown, the FPI server 304 forwards a request to the family 305 using a procedure call. Note that if FPI library 303 and the FPI server 304 use kernel messaging to communicate, the FPI server 304 provides a message port.

Each family 305 includes an expert (not shown) to maintain knowledge of the set of family devices. In one embodiment, the expert comprises code within a family 305 that maintains knowledge of the set of family plug-ins within the system. At system startup and each time a change occurs, the expert is notified.

In one embodiment, the expert may maintain the set of family services using a central device registry in the system.

The expert scans the device registry for plug-ins that belong to its family. For example, a display family expert looks for display device entries. When a family expert finds an entry for a family plug-in, it instantiates the plug-in, making it available to clients of the family. In one embodiment, the system notifies the family expert on an ongoing basis about new and deleted plug-ins in the device registry. As a result, the set of plug-ins known to and available through the family remains current with changes in system configuration.

Note that family experts do not add or alter information in the device registry nor do they scan hardware. In one embodiment, the present invention includes another level of families (i.e., low-level families) whose responsibility is to discover devices by scanning hardware and installing and removing information for the device registry. These low-level families are the same as the families previously discussed above (i.e., high level family) in other ways, i.e. they have experts, services, an FPI, a library, an activation model and plug-ins. The low-level families' clients are usually other families rather than applications. In one embodiment, families are insulated from knowledge of physical connectivity. Experts and the device registry are discussed in more detail below.

A plug-in programming interface (PPI) 306 provides a family-to-plug-in interface that defines the entry points a plug-in supports so that it can be called and a plug-in-to-family interface that defines the routines plug-ins call when certain events, such as an I/O completion, occur. In addition, PPI 306 defines the path through which the family and its plug-in exchange data.

A family services library (not shown) is a collection of routines that provide services to the plug-ins of a family. The services are specific to a given family and they may be layered on top of services provided by the kernel. Within a family, the methods by which data is communicated, memory is allocated, interrupts are registered and timing services are provided may be implemented in the family services library. Family services libraries may also maintain state information needed by a family to dispatch and manage requests.

For example, a display family services library provides routines that deal with vertical blanking (which is a concern of display devices). Likewise, SCSI device drivers manipulate command blocks, so the SCSI family services library contains routines that allow block manipulation. A family services library that provides commonly needed routines simplifies the development of that family's plug-ins.

Through the PPI 306, a call is made to a plug-in 307. In one embodiment, a plug-in, such as plug-in 307, comprises dynamically loaded code that runs in the kernel's address space to provide an instance of the service provided by a family. For example, within the file systems family, a plug-in implements file-system-specific services. The plug-ins understand how data is formatted in a particular file system such as HFS or DOS-FAT. On the other hand, it is not the responsibility of file systems family plug-ins to obtain data from a physical device. In order to obtain data from a physical device, a file system family plug-in communicates to, for instance, a block storage family. In one embodiment, block storage plug-ins provide both media-specific drivers, such as a tape driver, a CD-ROM driver, or hard disk driver, and volume plug-ins that represent partitions on a given physical disk. Block storage plug-ins in turn may make SCSI family API calls to access data across the SCSI bus on a physical disk. Note that in the present invention, plug-ins are a superset of device drivers. For instance, plug-ins may include code that does not use hardware. For instance, file

system and block storage plug-ins are not drivers (in that drivers back hardware).

Applications, plug-ins from other I/O families, and other system software can request the services provided by a family's plug-ins through the family's FPI. Note also that plug-ins are designed to operate in the environment set forth by their family activation model.

In one embodiment, a plug-in may comprises two code sections, a main code section that runs in a task in the kernel domain and an interrupt level code section that services hardware interrupts if the plug-in is, for instance, a device driver. In one embodiment, only work that cannot be done at task level in the main code section should be done at interrupt level. In one embodiment, all plug-ins have a main code section, but not all have interrupt level code sections.

The main code section executes and responds to client service requests made through the FPI. For example, sound family plug-ins respond to sound family specific requests such as sound playback mode setting (stereo, mono, sample size and rate), sound play requests, sound play cancellation, etc. The interrupt level code section executes and responds to interrupts from a physical device. In one embodiment, the interrupt level code section performs only essential functions, deferring all other work to a higher execution levels.

Also because all of the services associated with a particular family are tuned to the same needs and requirements, the drivers or plug-ins for a given family may be as simple as possible.

30 Family Programming Interfaces

In the present invention, a family provides either a user-mode or a kernel-mode FPI library, or both, to support the family's FPI. FIG. 4 illustrates one embodiment of the I/O architecture of the present invention. Referring to FIG. 4, three instances of families 401-403 are shown operating in the kernel environment. Although three families are shown, the present invention may have any number of families.

In the user mode, two user-mode FPI libraries, Xlib_u 404 and Zlib_u 405, are shown that support the FPIs for families X and Z, respectively. In the kernel environment, two kernel-mode FPI libraries, Ylib_k 406 and Zlib_k 407, for families Y and Z, respectively, are shown.

Both the user-mode and the kernel-mode FPI libraries present the same FPI to clients. In other words, a single FPI is the only way family services can be accessed. In one embodiment, the user-mode and kernel mode libraries are not the same. This may occur when certain operations have meaning in one mode and not the other. For example, operations that are implemented in the user-mode library, such as copying data across address-space boundaries, may be unnecessary in the kernel library.

In response to service requests, FPI libraries 404 and 405 map FPI functions into messages for communication from the user mode to the kernel mode. In one embodiment, the messages are kernel messages.

The service requests from other families are generated by plug-ins that make calls on libraries, such as FPI libraries 406 and 407. In one embodiment, FPI libraries 406 and 407 map FPI functions into kernel messages and communicate those messages to FPI servers such as Y FPI server 409 and Z FPI server 410 respectively. Other embodiments may use mechanisms other than kernel messaging to communicate information.

In the example, the Z family 403 has both a user-mode library 405 and a kernel-mode library 407. Therefore, the services of the Z family may be accessed from both the user mode and the kernel mode.

In response to service request messages, X FPI server 408, Y FPI server 409 and Z FPI server 410 dispatch requests for services to their families. In one embodiment, each of FPI servers 408-410 receives a kernel message, maps the message into a FPI function called by the client, and then calls the function in the family implementation (414-416).

In one embodiment, there is a one-to-one correspondence between the FPI functions called by clients and the function called by FPI servers 408-410 as a result. The calls from FPI servers 408-410 are transferred via interfaces 411-413. For instance, X interface 411 represents the interface presented to the FPI server 408 by the X family 414. It is exactly the same as the FPI available to applications or other system software. The same is true of Y interface 412 and Z interface 413.

The X family implementation 414 represents the family activation model that defines how requests communicated from server 408 are serviced by the family and plug-in(s). In one embodiment, X family implementation 414 comprises family code interfacing to plug-in code that completes the service requests from application 400 via server 408. Similarly, the Y family implementation 415 and Z family implementation 416 define their family's plug-in activation models.

X plug-in 417, Y plug-in 418 and Z plug-in 419 operate within the activation model mandated by the family and provide code and data exports. The required code and data exports and the activation model for each family of drivers is family specific and different.

Extending Family Programming Interfaces

A plug-in may provide a plug-in-specific interface that extends its functionality beyond that provided by its family. This is useful in a number of situations. For example, a block storage plug-in for a CD-ROM device may provide a block storage plug-in interface required of the CD-ROM device as well as an interface that allows knowledgeable application software to control audio volume and to play, pause, stop, and so forth. Such added capabilities require a plug-in-specific API.

If a device wishes to export extended functionality outside the family framework, a separate message port is provided by the device and an interface library for that portion of the device driver. FIG. 5 illustrates the extension of a family programming interface.

Referring to FIG. 5, a plug-in module, Z plug-in 501, extends beyond the Z family boundary to interface to family implementation D 502 as well. A plug-in that has an extended API offers features in addition to those available to clients through its family's FPI. In order to provide extra services, the plug-in provides additional software shown in FIG. 5 as an interface library Dlib_z 503, the message port code D FPI server 504, and the code that implements the extra features D 505.

Sharing Code and Data Between Plug-ins

In one embodiment, two or more plug-ins can share data or code or both, regardless of whether the plug-ins belong to the same family or to different families. Sharing code or data is desirable when a single device is controlled by two or more families. Such a device needs a plug-in for each family. These plug-ins can share libraries that contain information about the device state and common code. FIG. 6 illustrates two plug-ins that belong to separate families and that share code and data.

Plug-ins can share code and data through shared libraries. Using shared libraries for plug-ins that share code or data allows the plug-ins to be instantiated independently without

encountering problems related to simultaneous instantiation. Referring to FIG. 6, the first plug-in 601 to be opened and initialized obtains access to the shared libraries. At this point, the first plug-in 601 does not share access. When the second plug-in 602 is opened and initialized, a new connection to the shared libraries is created. From that point, the two plug-ins contend with each other for access to the shared libraries.

Sharing code or data may also be desirable in certain special cases. For instance, two or more separate device drivers may share data as a way to arbitrate access to a shared device. An example of this is a single device that provides network capabilities and real time clock. Each of these functions belong to a distinct family but may originate in a single physical device.

Activation Models in the Present Invention

An activation model defines how the family is implemented and the environment within which plug-ins of the family execute. In one embodiment, the activation model of the family defines the tasking model a family uses, the opportunities the family plug-ins have to execute and the context of those opportunities (for instance, are the plug-ins called at task time, during privileged mode interrupt handling, and so forth), the knowledge about states and processes that a family and its plug-ins are expected to have, and the portion of the service requested by the client that is performed by the family and the portion that is performed by the plug-ins.

Each model provides a distinctly different environment for the plug-ins to the family, and different implementation options for the family software. Examples of activation models include the single-task model, the task-per-plug-in model, and the task-per-request model. Each is described in further detail below. Note that although three activation models are discussed, the choice of activation model is a design choice and different models may be used based on the needs and requirements of the family.

In one embodiment, the activation model uses kernel messaging as the interface between the FPI libraries that family clients link to and the FPI servers in order to provide the asynchronous or synchronous behavior desired by the family client. Within the activation model, asynchronous I/O requests are provided with a task context. In all cases, the implementation of the FPI server depends on the family activation model.

The choice of activation model limits the plug-in implementation choices. For example, the activation model defines the interaction between a driver's hardware interrupt level and the family environment in which the main driver runs. Therefore, plug-ins conform to the activation model employed by its family.

Single-Task Model

One of the activation models that may be employed by a family is referred to herein as the single-task activation model. In the single-task activation model, the family runs as a single monolithic task which is fed from a request queue and from interrupts delivered by plug-ins. Requests are delivered from the FPI library to an accept function that enqueues the request for processing by the family's processing task and wakes the task if it is sleeping. Queuing, synchronization, and communication mechanism within the family follow a set of rules specified by the family.

The interface between the FPI Server and a family implementation using the single-task model is asynchronous. Regardless of whether the family client called a function synchronously or asynchronously, the FPI server calls the family code asynchronously. The FPI server maintains a set

of kernel message IDs that correspond to messages to which the FPI server has not yet replied. The concept of maintaining kernel message IDs corresponding to pending I/O server request messages is well-known in the art;

Consider as an example family **700**, which uses the single-task activation model, shown in FIG. 7. Referring to FIG. 7, an application **710** is shown generating a service request to the family's APIs **711**. APIs **711** contain at least one library in which service requests are mapped to FPI functions. The FPI functions are forwarded to the family's FPI server **701**. FPI server **701** dispatches the FPI function to family implementation **703**, which includes various protocols and a network device driver that operate as a single task. Each protocol layer provides a different level of service.

The FPI server **701** is an accept function that executes in response to the calling client via the FPI library (not shown). An accept function, unlike a message-receive-based kernel task, is able to access data within the user and kernel bands directly. The accept function messaging model requires that FPI server **701** be re-entrant because the calling client task may be preempted by another client task making service requests.

When an I/O request completes within the family's environment, a completion notification is sent back to the FPI server **701**, which converts the completion notification into the appropriate kernel message ID reply. The kernel message ID reply is then forwarded to the application that generated the service request.

With a single-task model, the family implementation is insulated from the kernel in that the implementation does it not have kernel structures, IDs, or tasking knowledge. On the other hand, the relationship between FPI server **701** and family code **702** is asynchronous, and has internal knowledge of data structures and communication mechanisms of the family.

The single-task model may be advantageously employed for families of devices that have one of several characteristics: (1) each I/O request requires little effort of the processing unit. This applies not only to keyboard or mouse devices but also to DMA devices to the extent that the processing unit need only set up the transfer, (2) no more than one I/O request is handled at once, such that, for instance, the family does not allow interleaving of I/O requests. This might apply to sound, for example, or to any device for which exclusive reservation is required (i.e., where only one client can use a device at a time). The opposite of a shared resource. Little effort for the processor exists where the processor initiates an I/O request and then is not involved until the request completes, or (3) the family to be implemented provides its own scheduling mechanisms independent of the underlying kernel scheduling. This applies to the Unix™ stream programming model.

Task-Per-Plug-In Model

For each plug-in instantiated by the family, the family creates a task that provides the context within which the plug-in operates.

FIG. 8 illustrates the task-per-plug-in model. Referring to FIG. 8, an application **801** generates service requests for the family, which are sent to FPI **802**. Using an FPI library, the FPI **802** generates a kernel message according to the family activation model **804** and a driver, such as plug-in driver **805**.

In one embodiment, the FPI server **803** is a simple task-based message-receive loop or an accept function. FPI server **803** receives requests from calling clients and passes those requests to the family code **804**. The FPI server **803** is

responsible for making the data associated with a request available to the family, which in turn makes it available to the plug-in that services the request. In some instances, this responsibility includes copying or mapping buffers associated with the original request message to move the data from user address space to the kernel level area.

The family code **804** consists in part of one or more tasks, one for each family plug-in. The tasks act as a wrapper for the family plug-ins such that all tasking knowledge is located in the family code. A wrapper is a piece of code that insulates called code from the original calling code. The wrapper provides services to the called code that the called code is not aware of.

When a plug-in's task receives a service request (by whatever mechanisms the family implementation uses), the task calls its plug-in's entry points, waits for the plug-in's response, and then responds to the service request.

The plug-in performs the work to actually service the request. Each plug-in does not need to know about the tasking model used by the family or how to respond to event queues and other family mechanisms; it only needs to know how to perform its particular function.

For concurrent drivers, all queuing and state information describing an I/O request is contained within the plug-in code and data and within any queued requests. The FPI library forwards all requests regardless of the status of outstanding I/O requests to the plug-in. When the client makes a synchronous service request, the FPI library sends a synchronous kernel message. This blocks the requesting client, but the plug-in's task continues to run within its own task context. This permits clients to make requests of this plug-in even while another client's synchronous request is being processed.

In some cases of a family, a driver (e.g., **805**) can be either concurrent or nonconcurrent. Nevertheless, clients of the family may make synchronous and asynchronous requests, even though the nonconcurrent drivers can handle only one request at a time. The device manager FPI server **803** knows that concurrent drivers cannot handle multiple requests concurrently. Therefore, FPI server **803** provides a mechanism to queue client requests and makes no subsequent requests to a task until the task signals completion of an earlier I/O request.

When a client calls a family function asynchronously, the FPI library sends an asynchronous kernel message to the FPI server and returns to the caller. When a client calls a family function synchronously, the FPI library sends a synchronous kernel message to the FPI server and does not return to the caller until the FPI server replies to the message, thus blocking the caller's execution until the I/O request is complete.

In either case, the behaviors of the device manager FPI server **803** is exactly the same: for all incoming requests, it either queues the request or passes it to the family task, depending on whether the target plug-in is busy. When the plug-in signals that the I/O operation is complete, the FPI server **803** replies to the kernel message. When the FPI library receives the reply, it either returns to the synchronous client, unblocking its execution or it notifies the asynchronous client about the I/O completion.

The task-per-plug-in model is intermediate between the single-task and task-per-request models in terms of the number of tasks it typically uses. The task-per-plug-in model is advantageously used where the processing of I/O requests varies widely among the plug-ins.

Task-Per-Request Model

The task-per-request model shares the following characteristics with the two activation models already discussed:

(1) the FPI library to FPI server communication provides the synchronous or asynchronous calling behavior requested by family clients, and (2) the FPI library and FPI server use kernel messages to communicate I/O requests between themselves. However, in the task-per-request model, the FPI server's interface to the family implementation is completely synchronous.

In one embodiment, one or more internal family request server tasks, and, optionally, an accept function, wait for messages on the family message port. An arriving message containing information describing an I/O request awakens one of the request server tasks, which calls a family function to service the request. All state information necessary to handle the request is maintained in local variables. The request server task is blocked until the I/O request completes, at which time it replies to the kernel message from the FPI library to indicate the result of the operation. After replying, the request server task waits for more messages from the FPI library.

As a consequence of the synchronous nature of the interface between the FPI server and the family implementation, code calling through this interface remains running as a blockable task. This calling code is either the request server task provided by the family to service the I/O (for asynchronous I/O requests) or the task of the requester of the I/O (for certain optimized synchronous requests).

The task-per-request model is advantageously employed for a family where an I/O request can require continuous attention from the processor and multiple I/O requests can be in progress simultaneously. A family that supports dumb, high bandwidth devices is a good candidate for this model. In one embodiment, the file manager family uses the task-per-request model. This programming model requires the family plug-in code to have tasking knowledge and to use kernel facilities to synchronize multiple threads of execution contending for family and system resources.

Unless there are multiple task switches within a family, the tasking overhead is identical within all of the activation models. The shortest task path from application to I/O is completely synchronous because all code runs on the caller's task thread.

Providing at least one level of asynchronous call between an application and an I/O request results in better latency results from the user perspective. Within the file system, a task switch at a file manager API level allows a user-visible application, such as the Finder™, to continue. The file manager creates an I/O tasks to handle the I/O request, and that task is used via synchronous calls by the block storage and SCSI families to complete their part in I/O transaction processing.

The Device Registry of the Present Invention

The device registry of the present invention comprises an operating system naming service that stores system information. In one embodiment, the device registry is responsible for driver replacement and overloading capability so that drivers may be updated, as well as for supporting dynamic driver loading and unloading.

In one embodiment, the device registry of the present invention is a tree-structured collection of entries, each of which can contain an arbitrary number of name-value pairs called properties. Family experts examine the device registry to locate devices or plug-ins available to the family. Low-level experts, discussed below, describe platform hardware by populating the device registry with device nodes for insertion of devices that will be available for use by applications.

In one embodiment, the device registry contains a device subtree pertinent to the I/O architecture of the present invention. The device tree describes the configuration and connectivity of the hardware in the system. Each entry in the device tree has properties that describe the hardware repre-

sented by the entry and that contain a reference to the driver in control of the device.

Multiple low-level experts are used, where each such expert is aware of the connection scheme of physical devices to the system and installs and removes that information in the device tree portion of the device registry. For example a low-level expert, referred to herein as a bus expert or a motherboard expert, has specific knowledge of a piece of hardware such as a bus or a motherboard. Also, a SCSI bus expert scans a SCSI bus for devices, and installs an entry into the device tree for each device that it finds. The SCSI bus expert knows nothing about a particular device for which it installs an entry. As part of the installation, a driver gets associated with the entry by the SCSI bus expert. The driver knows the capabilities of the device and specifies that the device belongs to a given family. This information is provided as part of the driver or plug-in descriptive structure required of all plug-ins as part of their PPI implementation.

Low-level experts and family experts use a device registry notification mechanism to recognize changes in the system configuration and to take family-specific action in response to those changes.

An example of how family experts, low-level experts, and the device registry service operate together to stay aware of dynamic changes in system configuration follows: Suppose a motherboard expert notices that a new bus, a new network interface and new video device have appeared within the system. The motherboard expert adds a bus node, a network node, and a video node to the device tree portion of the device registry. The device registry service notifies all software that registered to receive notifications of these events.

Once notified that changes have occurred in the device registry, the networking and video family experts scan the device registry and notice the new entry belonging to their family type. Each of the experts adds an entry in the family subtree portion of the device registry.

The SCSI bus expert notices an additional bus, and probes for SCSI devices. It adds a node to the device registry for each SCSI device that it finds. New SCSI devices in the device registry result in perusal of the device registry by the block storage family expert. The block storage expert notices the new SCSI devices and loads the appropriate drivers, and creates the appropriate device registry entries, to make these volumes available to the file manager. The file manager receives notification of changes to the block storage family portion of the device registry, and notifies the Finder™ that volumes are available. These volumes then appear on the user's desktop.

Whereas, many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that the particular embodiment shown and described by way of illustration are in no way to be considered limiting. Therefore, reference to the details of the various embodiments are not intended to limit the scope of the claims which themselves recite only those features regarded as essential to the invention.

Thus, a method and apparatus for handling I/O requests in a computer system has been described.

We claim:

1. A computer system comprising:
a bus;

at least one memory coupled to the bus for storing data and programming instructions that include applications and an operating system; and

a processing unit coupled to the bus and running the operating system and applications by executing programming instructions, wherein an application has a first plurality of tailored distinct programming inter-

15

faces available to access a plurality of separate sets of computer system services provided through the operating system of the computer system via service requests.

2. The computer system defined in claim 3 wherein each of the first plurality of tailored distinct programming interfaces are tailored to a type of I/O service provided by each set of I/O services.

3. A computer system comprising:

a bus;

at least one memory coupled to the bus for storing data and programming instructions that include applications and an operating system, wherein the operating system comprises a plurality of servers, and each of the first plurality of programming interfaces transfer service requests to one of the plurality of servers, wherein each of the plurality of servers responds to service requests from clients of the separate sets of I/O services; and

a processing unit coupled to the bus and running the operating system and applications by executing programming instructions, wherein an application has a first plurality of tailored distinct programming interfaces available to access a plurality of separate sets of I/O services provided through the operating system via service requests.

4. The computer system defined in claim 3 wherein service requests are transferred as messages in a messaging system.

5. The computer system defined in claim 4 wherein each of the plurality of servers supports a message port.

6. The computer system defined in claim 3 wherein at least one of the plurality of servers is responsive to service requests from applications and from at least one other set of I/O services.

7. The computer system defined in claim 3 wherein the operating system further comprises a plurality of activation models, wherein each of the plurality of activation models is associated with one of the plurality of servers to provide a runtime environment for the set of I/O services to which access is provided by said one of the plurality of servers.

8. The computer system defined in claim 7 wherein at least one instance of a service is called by one of the plurality of servers for execution in an environment set forth by one of the plurality of activation models.

9. A computer system comprising:

a bus;

at least one memory coupled to the bus for storing data and programming instructions that comprise applications and an operating system;

a processing unit coupled to the bus and running the operating system and applications by executing programming instructions, wherein the operating system provides computer system services through a tailored distinct one of a plurality of program structures, each tailored distinct program structure comprising:

a first programming interface for receiving service requests for a set of computer system I/O services of a first type,

a first server coupled to receive service requests and to dispatch service requests to the computer system I/O services,

an activation model to define an operating environment in which a service request is to be serviced by the set of computer system I/O services, and

at least one specific instance of the set of computer system I/O services that operate within the activation model.

16

10. The computer system defined in claim 9 wherein the first programming interface is responsive to request from applications and from other program structures.

11. The computer system defined in claim 9 wherein the first programming interface comprises at least one library for converting functions into messages.

12. The computer system defined in claim 9 wherein the first server receives a message corresponding a service request from the first programming interface, maps the message into a function called by the client, and then calls the function.

13. The computer system defined in claim 9 wherein the message comprises a kernel message.

14. A computer system comprising:

a bus;

at least one memory coupled to the bus for storing data and programming instructions that comprise applications and an operating system;

a processing unit coupled to the bus and running the operating system and applications by executing programming instructions, wherein the operation system provides input/output (I/O) services through a tailored distinct one of plurality of program structures, each tailored distinct program structure comprising:

a first programming interface for receiving service requests for a set of I/O services of a first type,

a first server coupled to receive service requests and to dispatch service requests to the I/O services,

an activation model to define operating environment in which a service request is to be serviced by the set of I/O services, and

at least one specific instance of the set of I/O services that operate within the activation model, wherein one of the said at least one specific instances comprises a service that accesses another program structure, and further wherein said one of said at least one specific instances communicates to said another program structure of a second type using a message created using a library sent to the server of said another program structure.

15. The computer system defined in claim 9 wherein two or more I/O services share code or data.

16. The computer system defined in claim 15 wherein said two or more I/O services are different types.

17. The computer system defined in claim 9 wherein the program structure further comprises a storage mechanism to maintain identification of available services to which access is provided via the first server.

18. A computer implemented method of accessing I/O services of a first type, said computer implemented method comprising the steps of:

generating a service request for a first type of I/O services;

a tailored distinct family server, operating in an operating system environment and dedicated to providing access to service requests for the first type of I/O service, receiving and responding to the service request based on an activation model specific to the first type of I/O services; and

a processor running an instance of the first type of I/O services that is interfaces to the file server to satisfy the service request.

19. The method defined in claim 18 wherein the service request is generated by an application.

20. The method defined in claim 18 wherein the service request is generated by an instance of an I/O service running in the operating system environment.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 5,915,131
DATED : June 22, 1999
INVENTOR(S) : Knight, et al.

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

In column 15 at line 14 delete "the" and insert -- a --

In column 15 at line 54 delete ":" and insert -- ; --

In column 16 at line 20 delete "operation" and
insert -- operating --

In column 16 at line 58 delete "interfaces" and
insert -- interfaced --

Signed and Sealed this

Eighteenth Day of January, 2000

Attest:



Q. TODD DICKINSON

Attesting Officer

Commissioner of Patents and Trademarks

EXHIBIT E



US005920726A

United States Patent [19]

Anderson

[11] Patent Number: 5,920,726

[45] Date of Patent: Jul. 6, 1999

[54] SYSTEM AND METHOD FOR MANAGING POWER CONDITIONS WITHIN A DIGITAL CAMERA DEVICE

5,477,264	12/1995	Sarbadhikari et al.	348/231
5,493,335	2/1996	Parulski et al.	348/233
5,560,022	9/1996	Dunstan et al.	395/750.01
5,634,000	5/1997	Wicht	395/182.08

[75] Inventor: Eric C. Anderson, San Jose, Calif

OTHER PUBLICATIONS

[73] Assignee: Apple Computer, Inc., Cupertino, Calif.

Martyn Williams, Review-NEC PC-DC401 Digital Still Camera, AppleLink Newbytes, Mar. 15, 1996, pp. 1-3.

[21] Appl. No.: 08/873,412

Primary Examiner—Ayaz R. Sheikh

Assistant Examiner—Xuan M. Thai

[22] Filed: Jun. 12, 1997

Attorney, Agent, or Firm—Carr & Ferrell LLP; Gregory J. Koerner

[51] Int. Cl.⁶ G06F 1/30

[52] U.S. Cl. 395/750.01

[58] Field of Search 395/750.01, 750.02, 395/750.03, 750.04, 750.05, 750.06, 182.22, 182.2, 182.12, 737, 575

[57] ABSTRACT


A system and method for recovering from a power failure in a digital camera comprises a power manager for detecting and handling power failures, an interrupt handler for responsively incrementing a counter device, service routines which register to receive notification of the power failure, and a processor for evaluating the counter and providing notification of the power failure to the service routines which may then assist the digital camera to recover from the power failure.

[56] References Cited

U.S. PATENT DOCUMENTS

5,283,792	2/1994	Davies, Jr. et al.	371/66
5,359,728	10/1994	Rusmack et al.	395/575
5,386,552	1/1995	Garney	395/575
5,475,428	12/1995	Hintz et al.	348/263
5,475,441	12/1995	Parulski et al.	348/552

18 Claims, 11 Drawing Sheets






FIG. 1




FIG. 2


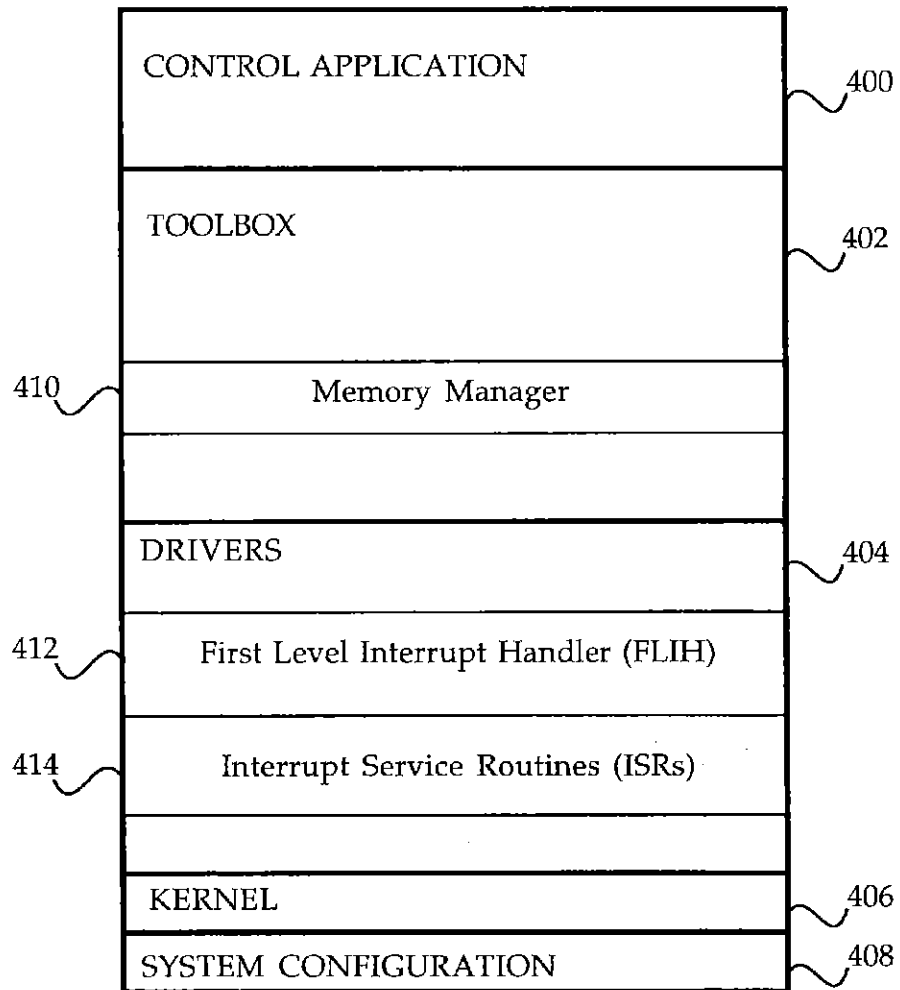


FIG. 3



350 ↗

FIG. 4

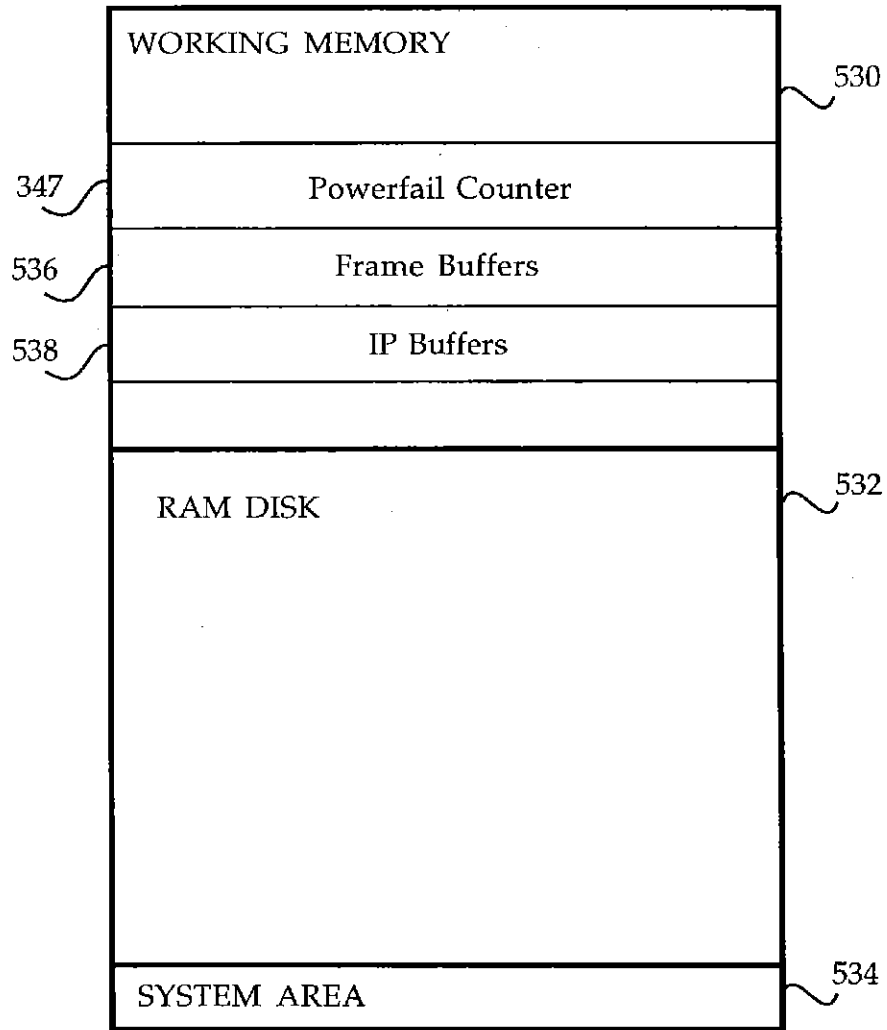


FIG. 5


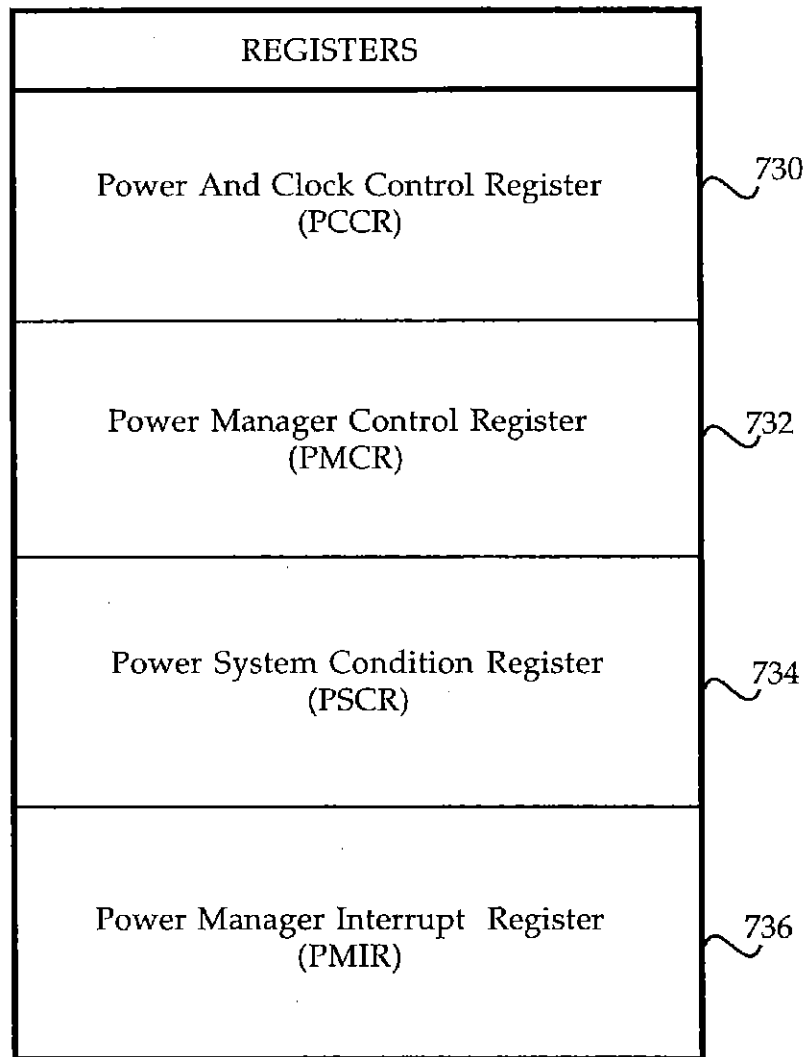


FIG. 6



646 ↗

FIG. 7

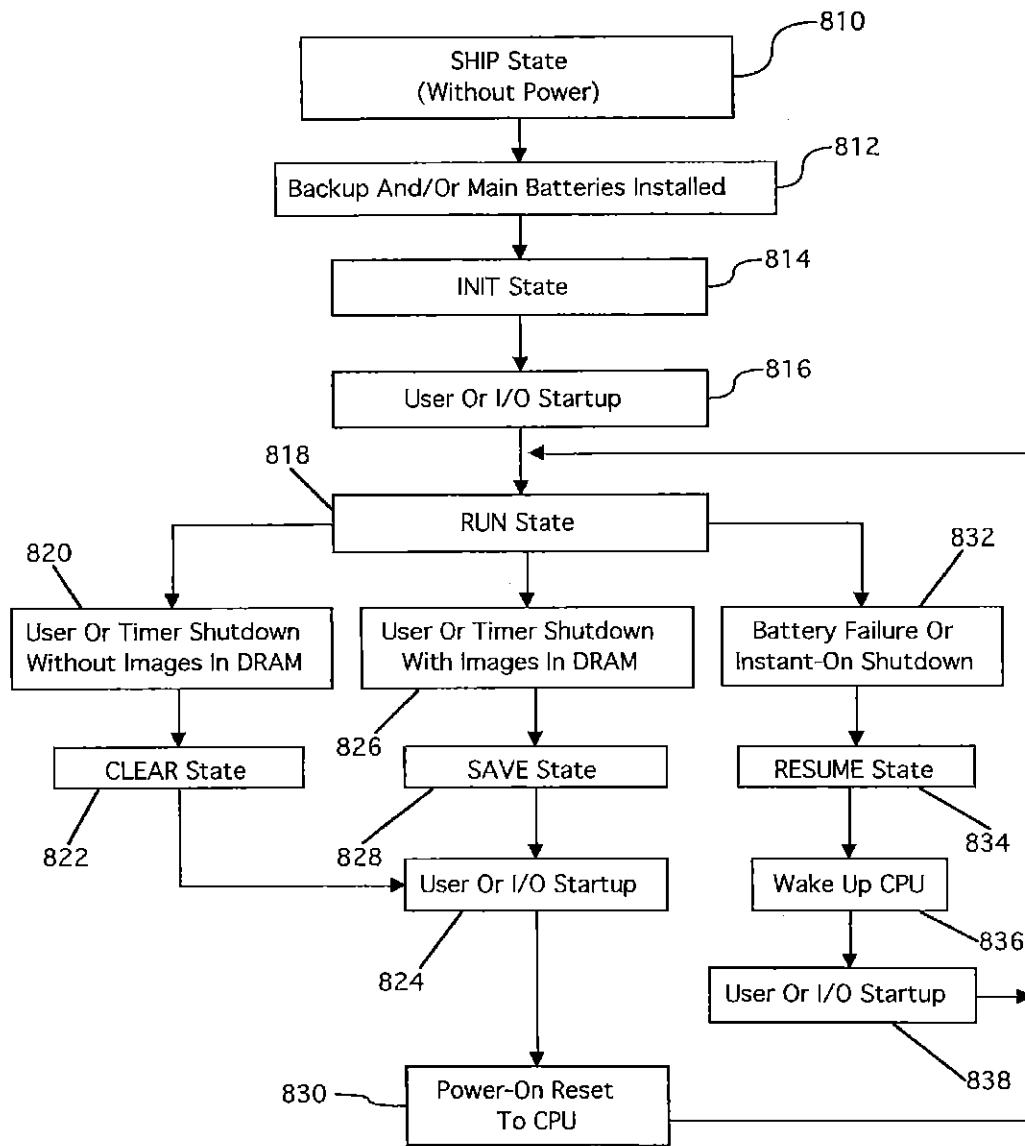


FIG. 8




FIG. 9




FIG. 10

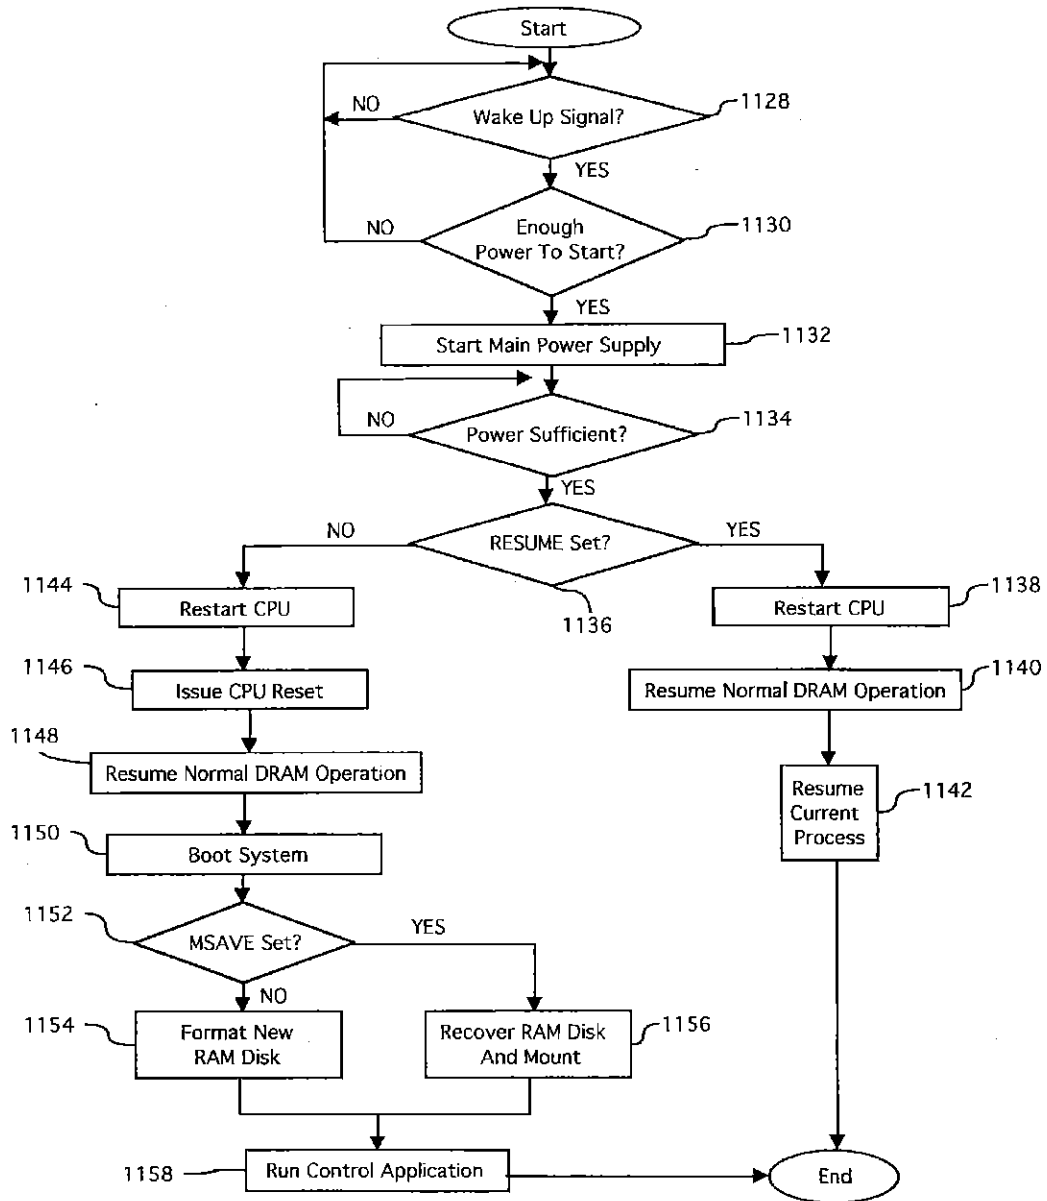


FIG. 11

SYSTEM AND METHOD FOR MANAGING POWER CONDITIONS WITHIN A DIGITAL CAMERA DEVICE

CROSS-REFERENCE TO RELATED APPLICATIONS

This application relates to co-pending U.S. patent application Ser. No. 08/702,246, entitled "System And Method For Recovering From A Power Failure Within A Digital Camera Device," filed on Aug. 23, 1996, and to co-pending U.S. patent application Ser. No. 08/719,264, entitled "System And Method For Conserving Power Within A Backup Battery Device," filed on Sep. 24, 1996, and to copending U.S. patent application Ser. No. 08/628,549, entitled "System And Method Using An LCD Indicator To Provide Status Of A Digital Camera Storage Device," filed on Apr. 10, 1996, which are hereby incorporated by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates generally to digital cameras and more particularly to a system and method for managing power conditions within a digital camera device.

2. Description of the Background Art

Current photographic technologies include various digital camera devices which capture image data by electronically scanning selected target objects. Digital camera devices typically process and compress the captured image data before storing the processed image data into internal or external memory devices. Furthermore, these digital camera devices may utilize multiple software routines running within a multi-threading environment to perform the various steps of capturing, processing, compressing and storing the image data.

Protecting the captured image data during the processing and compression stages (prior to final storage in non-volatile memory) is an important consideration of both camera manufacturers and camera users. Camera designers must therefore anticipate the occurrence of any events which might endanger the integrity of the captured image data.

A power failure in a digital camera device is one example of an event which might seriously jeopardize unprotected image data within the digital camera. For example, the digital camera may be performing a critical process at the time a power failure occurs. The intervening power failure may destroy the effect of the critical process and thus damage the image data or cause the digital camera to malfunction.

Furthermore, a power failure may interrupt various camera functions which typically resume their respective tasks whenever power is restored to the digital camera. The interrupted functions, however, would be unaware that a power failure had intervened. The interrupted functions would thus be unaware of the hardware reset which results from reapplying power after the power failure. This confusion between the system software and hardware would potentially endanger camera operations.

For the foregoing reasons, and because of other serious consequences of power failures in digital cameras, an improved system and method is needed for managing power conditions within a digital camera device, according to the present invention.

SUMMARY OF THE INVENTION

In accordance with the present invention, a system and method are disclosed for recovering from a power failure

within a digital camera device. The preferred embodiment of the present invention includes central processing unit, a powerfail counter, a first-level interrupt handler, various interrupt service routines, a power manager and a voltage sensor

In the preferred embodiment, the first-level interrupt handler initially sets the powerfail counter to a value of zero. Various interrupt service routines (each corresponding to a specific camera function or operation) may register themselves with the first-level interrupt handler (which coordinates all interrupts within the digital camera) to receive notification of an intervening power failure.

The power manager monitors the voltage sensor to detect a power failure within the digital camera. After detecting a power failure in which the camera operating power is less than a specified threshold value, the power manager generates a powerfail interrupt. The central processing unit responsively performs a powerfail powerdown sequence to preserve image data contained within the digital camera at the time of the intervening power failure. The power manager removes operating power from all non-critical subsystems and switches the critical subsystems to a backup power supply. The central processing unit and the camera's volatile memory are thus maintained in a static low-power mode, with all states preserved intact.

In the preferred embodiment, the first-level interrupt handler increments the powerfail counter to record the intervening power failure. The first level interrupt handler then notifies the registered interrupt service routines about the power failure restart and corresponding hardware reset.

After the power failure is remedied, the central processing unit performs a restart sequence to preserve any stored image data and to return the digital camera to a normal operational mode. The camera powerup sequence is performed in response to the contents of registers in the power manager. The power manager registers advantageously contain corresponding bits which indicate the conditions present in the camera at the time that the camera shutdown occurred. The present invention thus preserves the integrity of captured image data and effectively assists the digital camera to recover from an intervening power failure.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a digital camera according to the present invention;

FIG. 2 is a block diagram of the preferred embodiment for the imaging device of FIG. 1;

FIG. 3 is a block diagram of the preferred embodiment for the computer of FIG. 1;

FIG. 4 is a memory map showing the preferred embodiment of the Read-Only Memory of FIG. 3;

FIG. 5 is a memory map showing the preferred embodiment of the Dynamic Random-Access Memory of FIG. 3;

FIG. 6 is a block diagram of the preferred embodiment of the power manager of FIG. 3;

FIG. 7 is a block diagram of the preferred embodiment of the registers of FIG. 6;

FIG. 8 is a flowchart of the preferred power states for the power manager of FIG. 3;

FIG. 9 is a flowchart of preferred general method steps for recovering from a power failure according to the present invention;

FIG. 10 is a flowchart of preferred method steps for performing a powerfail powerdown sequence according to the present invention; and

FIG. 11 is a flowchart of preferred method steps for performing a powerup sequence according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

The present invention discloses a system and method for managing power conditions within a digital camera device and comprises a power manager for detecting and handling power failures, an interrupt handler for providing notice of power failures, service routines which register to receive notification of the power failure, and a processor for responsively controlling the digital camera during recovery from the power failure.

Referring now to FIG. 1, a block diagram of a camera 110 is shown according to the present invention. Camera 110 preferably comprises an imaging device 114, a system bus 116 and a computer 118. Imaging device 114 is optically coupled to an object 112 and electrically coupled via system bus 116 to computer 118. Once a photographer has focused imaging device 114 on object 112 and, using a capture button or some other means, instructed camera 110 to capture an image of object 112, computer 118 commands imaging device 114 via system bus 116 to capture raw image data representing object 112. The captured raw image data is transferred over system bus 116 to computer 118 which performs various image processing functions on the image data before storing it in its internal memory. System bus 116 also passes various status and control signals between imaging device 114 and computer 118.

Referring now to FIG. 2, a block diagram of the preferred embodiment of imaging device 114 is shown. Imaging device 114 preferably comprises a lens 220 having an iris, a filter 222, an image sensor 224, a timing generator 226, an analog signal processor (ASP) 228, an analog-to-digital (A/D) converter 230, an interface 232, and one or more motors 234.

U.S. patent application Ser. No. 08/355,031, entitled "A System and Method For Generating a Contrast Overlay as a Focus Assist for an Imaging Device," filed on Dec. 13, 1994 is incorporated herein by reference and provides a detailed discussion of the preferred elements of imaging device 114. Briefly, imaging device 114 captures an image of object 112 via reflected light impacting image sensor 224 along optical path 236. Image sensor 224 responsively generates a set of raw image data representing the captured image 112. The raw image data is then routed through ASP 228, A/D converter 230 and interface 232. Interface 232 has outputs for controlling ASP 228, motors 234 and timing generator 226. From interface 232, the raw image data passes over system bus 116 to computer 118.

Referring now to FIG. 3, a block diagram of the preferred embodiment for computer 118 is shown. System bus 116 provides connection paths between imaging device 114, power manager 342, central processing unit (CPU) 344, dynamic random-access memory (DRAM) 346, input/output interface (I/O) 348, read-only memory (ROM) 350, and connector 352. In the preferred embodiment, removable memory 354 may also connect to system bus 116 via connector 352.

Power manager 342 communicates via line 366 with power supply 356 and coordinates power management operations for camera 110 as discussed below in conjunction with FIGS. 6-11. CPU 344 typically includes a conventional processor device for controlling the operation of camera 110. In the preferred embodiment, CPU 344 is capable of con-

currently running multiple software routines to control the various processes of camera 110 within a multi-threading environment. DRAM 346 is a contiguous block of dynamic memory which may be selectively allocated to various storage functions by computer 118. DRAM 346 includes a powerfail counter 347 which is incremented each time a power failure occurs in power supply 356. DRAM 346 and powerfail counter 347 are further discussed below in conjunction with FIGS. 9-11.

I/O 348 is an interface device allowing communications to and from computer 118. For example, I/O 348 permits an external host computer (not shown) to connect to and communicate with computer 118. I/O 348 also permits a camera 110 user to communicate with camera 110 via a set of externally-mounted user controls and via an external LCD display panel. ROM 350 typically comprises a conventional nonvolatile read-only memory which stores a set of computer-readable program instructions to control the operation of camera 110. ROM 350 is further discussed below in conjunction with FIG. 4. Removable memory 354 serves as an additional image data storage area and is preferably a non-volatile device, readily removable and replaceable by a camera 110 user via connector 352. Thus, a user who possesses several removable memories 354 may replace a full removable memory 354 with an empty removable memory 354 to effectively expand the picture-taking capacity of camera 110. In the preferred embodiment of the present invention, removable memory 354 is typically implemented using a flash disk.

Power supply 356 supplies operating power to the various components of camera 110. In the preferred embodiment, power supply 356 provides operating power to a main power bus 362 and also to a secondary power bus 364. The main power bus 362 provides power to imaging device 114, I/O 348, ROM 350 and removable memory 354. The secondary power bus 364 provides power to power manager 342, CPU 344 and DRAM 346.

Power supply 356 is connected to main batteries 358 and also to backup batteries 360. In the preferred embodiment, a camera 110 user may also connect power supply 356 to an external power source. During normal operation of power supply 356, the main batteries 358 provide operating power to power supply 356 which then provides the operating power to camera 110 via both main power bus 362 and secondary power bus 364.

During a power failure mode in which the main batteries 358 have failed (when their output voltage has fallen below a minimum operational voltage level) the backup batteries 360 provide operating power to power supply 356 which then provides the operating power only to the secondary power bus 364 of camera 110. Selected components of camera 110 (including DRAM 346) are thus protected against a power failure in the main batteries 358.

Power supply 356 preferably also includes a flywheel capacitor connected to the power line coming from the main batteries 358. If the main batteries 358 suddenly fail, the flywheel capacitor temporarily maintains the voltage from the main batteries 358 at a sufficient level, so that computer 118 can protect any image data currently being processed by camera 110 before shutdown occurs.

Voltage sensor 359 detects the voltage supplied by main batteries 358 and responsively provides the detected voltage reading to power manager 342. The operation of power manager 342, power supply 356 and voltage sensor 359 are further discussed below in conjunction with FIGS. 6-11.

Referring now to FIG. 4, a memory map showing the preferred embodiment of read-only memory (ROM) 350 is

shown. In the preferred embodiment, ROM 350 includes control application 400, toolbox 402, drivers 404, kernel 406 and system configuration 408. Control application 400 comprises program instructions for controlling and coordinating the various functions of camera 110. Toolbox 402 contains selected function modules including memory manager 410 which is controlled by control application 400 and responsively allocates DRAM 346 storage locations depending upon the needs of computer 118 and the sets of received image data.

Drivers 404 control various components of camera 110 and include a first level interrupt handler (FLIH) 412 and various interrupt service routines (ISRs) 414. In the preferred embodiment, FLIH 412 is a software routine which coordinates all interrupts within camera (110. FLIH 412 typically handles ordinary non-critical interrupts and also handles non-maskable critical interrupts such as a power failure in main batteries 358. FLIH 412 preferably communicates with the various ISRs 414 which are each designed to handle a specific corresponding interrupt within camera 110. FLIH 412 notifies the appropriate ISRs 414 via a "signal" when the interrupts occur. A signal is a mechanism used by multi-tasking operating systems for interprocess communications and synchronization.

For example, a camera 110 user may request zoom motor 234 to perform a zoom operation using lens 220. When the requested zoom process is complete, an interrupt is generated to indicate that zoom motor 234 and lens 220 have reached their destination positions. The particular ISR 414 which corresponds to the foregoing zoom process then responsively handles the generated interrupt and provides a status update to higher-level routines, if necessary. In preferred embodiment, kernel 406 provides a range of basic underlying services for the camera 110 operating system. System configuration 408 performs initial start-up routines for camera 110, including the boot routine and initial system diagnostics.

Referring now to FIG. 5, a memory map showing the preferred embodiment of dynamic random-access memory (DRAM) 346 is shown. In the preferred embodiment, DRAM 346 includes working memory 530, RAM disk 532 and system area 534. Working memory 530 includes a powerfail counter 347, frame buffers 536 (for initially storing sets of raw image data received from imaging device 114) and image processing (IP) buffers 538 (for temporarily storing image data during the image processing and compression 420 process). In the preferred embodiment, power fail counter 347 stores a value which first-level interrupt handler 412 preferably increments each time voltage sensor 359 detects a power failure in main batteries 358. Powerfail counter 347 is further discussed below in conjunction with FIGS. 9-11. Working memory 530 may also contain various stacks, data structures and variables used by CPU 344 while executing the software routines used within computer 118.

RAM disk 532 is a memory area used for storing raw and compressed image data and typically is organized in a "sectored" format similar to that of conventional hard disk drives. In the preferred embodiment, RAM disk 532 uses a well-known and standardized file system to permit external host computer systems, via I/O 348, to readily recognize and access the data stored on RAM disk 532. System area 534 typically stores data regarding system errors (for example, why a system shutdown occurred) for use by CPU 344 upon a restart of computer 118.

Referring now to FIG. 6, a block diagram of the preferred embodiment of power manager 342 (FIG. 3) is shown.

Power manager 342 includes bus interface 640, registers 646, state machine 650 and LCD generator 654. Bus interface 640 is connected, via line 660, to system bus 116 and may thus handle slave access of registers 646 within power manager 342 (typically by CPU 344).

Registers 646 include a PCCR register, a PMCR register, a PSCR register and a PMIR register each described below in conjunction with FIG. 7. Registers 646 also include external input pins (not shown) that asynchronously affect selected bit transitions within registers 646. These input pins include PMRST_, MBFAIL_, MBALERT_, PWRSW_, IOSYS_, and USRRST_. The effect of these pin transitions is also dependent on the current state of power manager 342. The following discussion describes the above-referenced register 646 input pins and the results of asserting each of these input pins via line 662.

A logic level of value "0" on the PMRST_ pin causes all bits in the PCCR, PMIR and PMCR registers to be cleared to their inactive states. This action occurs regardless of the current state of power manager 342. In response, power manager 342 will then transition to the IDLE_OFF state (described below). A "1 to 0" transition on the MBFAIL pin causes all bits in the PCCR register to be immediately deasserted if power manager 342 is in the IDLE_ON state (described below). This transition will also set the PwrFail bit to 1 in the PMCR register, the MBFail bit to 1 in the PMIR register, and assert the IRQ0_ pin. If power manager 342 is in any other state than IDLE_ON, any transition on this pin will be ignored. A "1 to 0" transition on the MBALERT pin will cause the Strobe power and control bits in the PCCR to be immediately deasserted if power manager 342 is in the IDLE_ON state. This transition will also set the MBAlert bit in the PMIR register to 1 and activate the IRQ1_ power alert interrupt. If power manager 342 is in any other state than IDLE_ON, any transition on this pin will be ignored.

In the preferred embodiment, a momentary switch may be used to cause a "1 to 0" transition on the PWRSW_ pin. The momentary "1 to 0" transition on the PWRSW_ pin will immediately set the Uspwr bit to 1 in the PMCR register if power manager 342 is in the IDLE_OFF state. It will also initiate the PowerUp state sequencing. Furthermore, the momentary "1 to 0" transition on the PWRSW_ pin will immediately set the PDRReq bit in the PMIR register if power manager 342 is in the IDLE_ON state. It will also cause the assertion of the IRQ2_ pin. If power manager 342 is in any other state than IDLE_ON or IDLE_OFF, the momentary transition on this pin will be ignored.

In an alternate embodiment without the above-mentioned momentary switch, a "1 to 0" transition on the PWRSW_ pin will immediately set the Uspwr bit to 1 in the PMCR register if power manager 342 is in the IDLE_OFF state. It will also initiate the PowerUp state sequencing. If power manager 342 is in any other state than IDLE_OFF, the "1 to 0" transition on this pin will be ignored.

A "0 to 1" transition on the PWRSW_ pin will immediately set the PDRReq bit in the PMIR register if power manager 342 is in the IDLE_ON state. It will also cause the assertion of the IRQ2_ pin. If power manager 342 is in any other state than IDLE_ON, the "0 to 1" transition on this pin will be ignored.

A "1 to 0" transition on the IOSYS_ pin will immediately set the IOPwr bit to 1 in the PMCR register if power manager 342 is in the IDLE_OFF state. It will also initiate the PowerUp state sequencing. A "1 to 0" transition on the IOSYS_ pin will immediately set the IOReq bit in the PMIR

register if power manager 342 is in the IDLE_ON state. This transition will also cause the assertion of the IRQ3 pin. A "1 to 0" transition on the USRRST pin will immediately set the USRRST bit to 1 in the PMCR register if power manager 342 is in the IDLE_ON state. This transition will also cause power manager 342 to issue a CPU 344 soft reset (SRST) for 30 us (one clock pulse of the 32 KHz real-time clock), starting at the next rise of PMCLK. CPU 344 will then hold SRST low for 16 ms. The PCCR and PMIR registers are also cleared. Additionally, all bits in the PMCR are cleared except for the USRRST bit. If power manager 342 is in any other state than IDLE_ON, any transition on this pin will be ignored. Power manager 342 will remain in the IDLE_ON state.

Registers 646 generate a series of interrupts onto line 664 in response to various conditions and states in camera 110. Registers 646 and the generated interrupts are further discussed below in conjunction with FIGS. 7 and 8. Registers 646 also provide signals to LCD generator 654 via line 670. LCD generator 654 responsively generates and provides LCD outputs to an LCD display unit (not shown) via line 672. The LCD display unit is preferably mounted on the exterior surface of camera 110 and forms part of a user interface for camera 110.

In the preferred embodiment, power manager 342 has three main transition events called PowerUp, PowerDown and PowerDown w/Save. The PowerUp event occurs when power manager 342 is in the IDLE_OFF state and one of three external events occur. These external events occur when the user turns on the camera (PWRSW_low), when a timer wakeup (TEXP_high) is signaled, and when the host attempts to connect to the camera (IOSYS_low). The PowerDown event occurs when power manager 342 is in the IDLE_ON state and the software writes a 1 to the PwrDwn bit in the PMCR and the MemSave bit in the PMCR is set to 0. The PowerDown w/Save occurs when power manager 342 is in the IDLE_ON state and the software writes a 1 to the PwrDwn bit in the PMCR and the MemSave bit in the PMCR is set to 1.

Power manager 342 includes state machine 650 which preferably has eight main synchronous states. The following is a description of the main states and the events that trigger a transition into the particular state. The IDLE_OFF state is the initial state after PMRST is released. For the PowerUp sequence, power manager 342 starts from IDLE_OFF, moves through two more states and ends with the IDLE_ON state.

The IDLE_OFF_SAVE state is entered from the IDLE_ON state when the MemSave option is specified. For the PowerUp sequence, power manager 342 starts from IDLE_OFF_SAVE, moves through two more states and ends with the IDLE_ON state. The IDLE_ON is the end state for a PowerUp sequence. In the PowerDown sequence, power manager 342 starts from IDLE_ON and ends in IDLE_OFF or IDLE_OFF_SAVE.

In the preferred embodiment, there are two types of PowerDown sequences called PWRDWN_NORM and PWRDWN_SAVE. The PWRDWN_NORM sequence executes various steps to PowerDown camera 110 before entering an idle state. This sequence is executed if the software writes a 1 to the PwrDwn bit in the PMCR when power manager 342 is in the IDLE_ON state and the MemSave bit in the PMCR is programmed to 0. The PWRDWN_SAVE sequence executes various steps to PowerDown camera 110 before entering an idle state. This sequence is entered if the software writes a 1 to the PwrDwn

bit in the PMCR when power manager 342 is in the IDLE_ON state and the MemSave bit in the PMCR is programmed to 1.

The CNT_PRE_WAIT state is used to wait for voltage sensor 359 to start operation. When the desired time has passed, power manager 342 will continue the PowerUp sequence and move to the CNT_WAIT state which is used to wait for the test of main batteries 358 to complete. When the desired time has passed, power manager 342 will continue the PowerUp sequence and move to the POR_IRQ_WAIT state.

The CNT_PRE_WAIT_SAVE is also used to wait for voltage sensor 359 to start operation. When the specified time has passed, power manager 342 continues the (PowerUp sequence and also switches to the CNT_WAIT_SAVE state. The CNT_WAIT_SAVE state is used to wait for testing of main batteries 358 to complete while the MemSave is specified. When the desired time has passed, power manager 342 will continue the PowerUp sequence and move to the POR_IRQ_WAIT state. The POR_IRQ_WAIT state is used to wait one PMCLK cycle for power manager 342 to drive the POR signal or to deassert the IRQ0 interrupt before transitioning to the IDLE_ON state.

State machine 650 receives, via line 668, a real-time clock which preferably operates at a 32 KHz rate. In the preferred embodiment, state machine 650 also generates a power control signal which is supplied to power supply 356 via line 366 to advantageously control the operation of power supply 356 according to the present invention.

Referring now to FIG. 7, a block diagram of the preferred embodiment of registers 646 is shown. In the preferred embodiment, registers 646 include power and clock control register (PCCR) 730, power manager control register (PMCR) 732, power system condition register (PSCR) 734 and power manager interrupt register (PMIR) 736.

PCCR 730 is an 8-bit read/write register which includes the following binary fields. Bits 6-7 of PCCR 730 correspond to a field named I/O. The I/O field is an I/O Enable and these bits control which I/O system is engaged and in full power mode. If the I/O field is 00, then no I/O is enabled and if the I/O field is 01, then the serial port is enabled. If the I/O field is 10, then USB is enabled and if the I/O field is 11, then IRDA is enabled.

Bit 5 of PCCR 730 corresponds to a field named LCDCtrl. The LCDCtrl field is a LCD and backlight power control. Bit 4 of PCCR 730 corresponds to a field named AudioEn. The AudioEn field is an audio system clock and power enable. Bit 3 of PCCR 730 corresponds to a field named ICHClk. The ICHClk field is an image capture head clock enable. Bit 2 of PCCR 730 corresponds to a field named ICHPwr. The ICHPwr field is an image capture head power control (and also controls LED power). Bits 0-1 of PCCR 730 correspond to a field named ICHStrb. The ICHStrb field is an image capture head strobe power control and level. Both bits are cleared by the Power Alert condition.

All bits of PCCR 730 are placed in their deasserted state by the Power Fail condition. The ICHstrb bits are cleared by the Power Alert condition. All bits correspond to pins on the power manager 342 ASIC. Software controls the setting or clearing of the bits to power down subsystems. All bits are placed in their deasserted state during PMRST.

PMCR 732 is an 8-bit read/write register which is the main control center of power manager 342. All bits are cleared to zero during PMRST. PMCR 732 preferably includes the following binary fields. Bit 7 of PMCR 732

corresponds to a field named RESUME. The RESUME field is set and cleared by software and is set for power failure shutdown or "instant-on shutdown". RESUME prevents reset of CPU 344 on start-up. Bit 6 of PMCR 732 corresponds to a field named MSAVE. The MSAVE field is set and cleared by software and is set to maintain power to DRAM 346 and CPU 344 in Shut Down mode. MSAVE is also used to generate an LCD status indication during main power off/on.

Bit 5 of PMCR 732 corresponds to a field named DiskIU. The DiskIU field is set and cleared by software and is set when the DiskInUse indicator on the status LCD is required. The status LCD maintains this display with main power off/on. Bit 4 of PMCR 732 corresponds to a field named PwrDwn. The PwrDwn field is set by software to initiate power down from power manager 342. CPU 344 enters Power Down Mode or Deep Sleep Mode 1 PMCLK cycle following the next rising edge of PMCLK after setting this bit. This bit is cleared on Power Up sequence by power manager 342. While power manager 342 is clearing this bit during the PowerOn sequence, writes to this bit will be disabled.

Bit 3 of PMCR 732 corresponds to a field named UserPwr. The UserPwr field is set during the normal power up sequence if initiated by a user pressing the main power switch (PWRSW input). UserPwr is cleared by software writing a 1 to this bit. Bit 2 of PMCR 732 corresponds to a field named IOPwr. The IOPwr field is set during the normal power up sequence if initiated by the external communications port 348 (IOSYS_ input). If neither UserPwr, IOPwr, or UserRst is set, software will assume a timer initiated restart has occurred. Bit 1 of PMCR 732 corresponds to a field named UserRst. The UserRst field is set by the user pressing the RESET button (USERRST_ input) and is cleared by software by writing a 1 to this bit during boot. If neither PwrFail or UserRst is set, software assumes a normal power up. Bit 0 of PMCR 732 corresponds to a field named PwrFail. The PwrFail field is set at the "0 to 1" transition of the MBFail bit in PSCR 734. The PwrFail field is cleared by software writing a 1 to this bit after restart.

PSCR 734 is an 8-bit read-only register with the following fields. Bit 7 of PSCR 734 corresponds to a field named BBOK. The BBOK field indicates the backup batteries 360 contain sufficient operating power (above 4.5 volts). Bit 6 of PSCR 734 corresponds to a field named MBStart. The MBStart field corresponds to the main batteries 358 START level (above 6.0 volts). MBStart indicates that there is enough energy to do a camera 110 startup. Bit 5 of PSCR 734 corresponds to a field named MBLow. The MBLow field indicates that main batteries 358 are above 6.4 volts. Bit 4 of PSCR 734 corresponds to a field named MBGood. The MBGood field indicates that main batteries 358 are above 6.9 volts. Bit 3 of PSCR 734 corresponds to a field named MBFull. The MBFull field indicates that main batteries 358 are above 7.4 volts. Bit 2 of PSCR 734 corresponds to a field named ACPwr. The ACPwr field indicates that external AC power is being supplied to the camera via an external power adapter. Bit 1 of PSCR 734 corresponds to a field named MBAlert. MBAlert indicates an ALERT condition for main batteries 358 (below 5.4 volts). This bit reflects the current level of the MBALERT_ input. A "1 to 0" transition on the MBALERT_ pin generates an alert interrupt (IRQ1_ output), sets the Pfail bit in PMIR 736 to 1, and clears two bits in PCCR 730. Bit 0 of PSCR 734 corresponds to a field named MBFail. The MBFail field indicates a PwrFail in main batteries 358 (below 5.2 volts) and reflects the current level of the MBFAIL_ input. A "1 to 0" transition on the

MBFAIL_ pin generates a power failure interrupt (IRQ0_ output), sets the MBALERT bit in PMIR 736, deasserts all bits in PCCR 730, and sets the PwrFail bit in PMCR 732. PSCR 734 thus indicates the state of eight power system conditions.

PMIR 736 is an 8-bit read/write register with the following fields. Bits 4-7 of PMIR 736 correspond to a field named Reserved. The Reserved field always reads 0 and is reserved for potential future functionality in camera 110. Bit 3 of PMIR 736 corresponds to a field named IOReq. The IOReq field is a request from the external communications port 348 of camera 110. A "1 to 0" transition on the IOSYS_ pin while power manager 342 is in the IDLE_ON state generates an IOReq interrupt (IRQ3_ output). A write of 1 to this bit will immediately clear the interrupt.

Bit 2 of PMIR 736 corresponds to a field named PwrDwn. The PwrDwn field is a power down request generated by a camera 110 user (PWRSW pin). A "0 to 1" transition on the PWRSW_ pin generates a PwrDwn interrupt (IRQ2_ output). A write of 1 to this bit will immediately clear the interrupt. Bit 1 of PMIR 736 corresponds to a field named MBAlert. The MBAlert field indicates an ALERT condition (below 5.4 volts) for main batteries 358. A "1 to 0" transition on the MBALERT_ pin immediately generates an ALERT interrupt (IRQ1_ output) and clears two bits in the PCCR 730. A write of 1 to this bit will immediately clear the interrupt. Bit 0 of PMIR 736 corresponds to a field named MBFail. The MBFail field indicates a power failure condition (below 5.2 volts) in main batteries 358. A "1 to 0" transition on the MBFail_ pin generates a PwrFail interrupt (IRQ0_ output), deasserts all bits in the PCCR 730, and sets the PwrFail bit in the PMCR 732. A write of 1 to this bit will immediately clear the interrupt. PMIR 736 thus indicates the interrupt status of camera 110 and is also used to clear the interrupts. All bits are cleared during PMRST_.

Referring now to FIG. 8, a flowchart of the preferred power states for power manager 342 is shown. Initially, in step 810, camera 110 is in the SHIP state with neither main batteries 358 nor backup batteries 360 installed. In steps 812 and 814, the SHIP to INIT transition occurs when power is first applied to the camera 110, either from backup batteries 360, main batteries 358, or via external power.

In steps 816 and 818, the INIT to RUN transition occurs when the user turns on camera 110 (PWRSW low transition) or when the host computer attempts to connect to camera 110 (IOSYS low transition). CPU 344 is in POWER DOWN MODE prior to this transition. Following step 818, the FIG. 8 sequence proceeds either to step 820, 826 or 832, depending on the type of shutdown occurring in camera 110.

In steps 820 and 822, a CLEAR state may be entered if a shutdown occurs without images in DRAM 346. The RUN to CLEAR transition occurs when a user turns off camera 110 (PWRSW high transition), a host or script command to shut down camera 110 occurs, camera 110 times out after no activity, or camera 110 enters a timed shutdown. In the CLEAR state, the MSAVE bit in PMCR 732 is set to zero and the RESUME bit in PMCR 732 is also set to zero.

In steps 826 and 828, a SAVE state may be entered if a shutdown occurs with images in DRAM 346. The RUN to SAVE transition occurs when a user turns off camera 110 (PWRSW high transition, followed by software response to interrupt), a host or script command to shut down camera 110 occurs, camera 110 times out after no activity, or camera 110 enters a timed shutdown. In the SAVE state, the MSAVE bit in PMCR 732 is set to one and the RESUME bit in PMCR 732 is set to zero. In steps 832 and 834, a RESUME

state may be entered if a shutdown occurs due to a power failure in main batteries 358 or instant-on shutdown. In the RESUME state, the MSAVE bit in PMCR 732 is set to one and the RESUME bit in PMCR 732 is also set to one.

In steps 824, 830 and 818, a CLEAR to RUN transition may occur if a user turns on camera 110 (PWRSW low transition), a timer wakeup (TEXP) is signaled or if the host computer attempts to connect to camera 110 (IOSYS low transition). CPU 344 is in POWER DOWN MODE prior to this transition. In steps 824, 830 and 818, a SAVE to RUN transition may occur if a user turns on camera 110 (PWRSW low transition), a timer wakeup (TEXP) is signaled or when the host computer attempts to connect to camera 110 (IOSYS low transition). CPU 344 is in DEEP SLEEP MODE prior to this transition. In steps 834, 836 and 838, a RESUME to RUN transition may occur if a wake up of CPU 344 is signaled and a user turns on camera 110 (PWRSW low transition), a timer wakeup (TEXP) is signaled or the host computer attempts to connect to camera 110 (IOSYS low transition).

In FIG. 8, power manager 342 thus sequentially enters the SHIP state, the INIT state and the RUN state. Power manager 342 then may enter either the CLEAR state, the SAVE state or the RESUME state, depending on the particular shutdown conditions. Finally, power manager 342 typically transitions back to the RUN state to recommence normal operation, in accordance with the present invention.

Referring now to FIG. 9, a flowchart of preferred general method steps for recovering from a power failure is shown. Initially, a user applies 940 power to camera 110 by installing main batteries 358 and backup batteries 360, and then activating an externally-mounted power on-off switch. First-level interrupt handler (FLIH) 412 then sets 942 powerfail counter 347 to a value of zero. Next, various interrupt service routines 414 register 944 with the first level interrupt handler 412 to request notification in the event of a power failure in main batteries 358.

CPU 344 then runs 946 control application 400 to operate camera 110 in normal operation mode which captures, processes, compresses and stores sequential sets of image data. In normal operation mode, CPU 344 periodically requests the execution of various critical processes. In the preferred embodiment, CPU 344 repeatedly checks powerfail counter 347 to determine whether a critical process has been interrupted by an intervening power failure.

Next, voltage sensor 359 senses 948 the voltage level of main batteries 358 and provides power manager 342 with the sensed voltage level. Power manager 342 then determines 950 whether the voltage level of main batteries 358 is greater than a predetermined threshold voltage level. The threshold value is typically selected to be incrementally higher than the minimum operating voltage (to permit orderly shutdown of the camera 110 processes). If the voltage of main batteries 358 is greater than the selected threshold value, then the FIG. 9 process repeats the steps 946, 948 and 950.

However, if the voltage of main batteries 358 is not greater 950 than the predetermined threshold value, then power manager 342 generates 952 a powerfail interrupt. In the preferred embodiment, the powerfail interrupt may be disabled in rare cases in which a sequence of CPU 344 instructions must never be interrupted by a power failure. Any disabling of the powerfail interrupt, however, is restricted to a very short period of time. Next, CPU 344 receives the generated powerfail interrupt and responsively performs 954 a powerfail powerdown sequence to protect

the image data currently within camera 110. The powerfail powerdown sequence is further discussed below in conjunction with FIG. 10.

The camera 110 user may then replace 655 the main batteries 358 and activate the camera 110 power on/off switch. CPU 344 then performs 956 a powerup sequence to bring camera 110 back to normal operating mode while also preserving any existing image data. FLIH 412 then increments 958 powerfail counter 347 to indicate the occurrence of a power failure in main batteries 358. Alternately, powerfail counter 347 may be a hardware register which is incremented in power manager 342.

The first level interrupt handler 412 then notifies 960 any registered interrupt service routines 414 about the power failure restart so that the interrupt service routines 414 are aware that their corresponding hardware components have been reset by the power failure and the subsequent camera 110 powerup. The power failure notification allows the registered interrupt service routines 414 to run depending upon their relative task priority. Typically, this notification is accomplished through the use of a signal or semaphore which wakes up the interrupt service routine.

In alternate embodiments, the interrupt service routines may operate in cooperation with various other system routines. These cooperating routines thus may form various hierarchical networks which operate in synchronous or asynchronous modes. For example, a particular interrupt service routine may function in response to a device driver. The device driver, in turn, may function in response to an application program. In such cases, the interrupt service routines typically propagate their received power failure notification to any related routines in the network which require notification of the power failure restart. Finally, the FIG. 9 process then returns to step 946 and CPU 344 runs control application 400 to operate camera 110 in normal operation mode, as discussed above.

Referring now to FIG. 10, a flowchart of preferred method steps for performing a powerfail powerdown sequence according to the present invention is shown. Initially, power manager 342 sets 1010 a PFAIL bit which records the occurrence of a power failure so that computer 118 software routines may subsequently access this information when needed. Next, power manager 342 turns off 1012 all non-critical subsystems. Power manager 342 then signals 1014 CPU 344 with an interrupt and CPU 344 responsively stops 1016 the current process.

Next, CPU 344 sets 1018 the RESUME bit in power manager 342 to indicate that CPU 344 should not be reset in a subsequent powerup of camera 110. CPU 344 then forces 1020 a full refresh of DRAM 346 and then forces 1022 DRAM 346 into a self-refresh mode. Next, CPU 344 signals 1024 power manager 342 to shut down and then CPU 344 halts 1026 operation. After halting, CPU 344 still receives operating power from backup batteries 360 and is essentially stopped "in place." In this static mode, system, bus 116 is in a tri-state condition and the CPU 344 clock is stopped. All CPU 344 states, however, are still intact (for example, the registers, program counter, cache and stack are preserved intact) and image data in DRAM 346 is also preserved intact. Next, power manager 342 removes 1028 operating power from main power bus 362. The FIG. 10 powerfail powerdown sequence is then complete.

Referring now to FIG. 11, a flowchart of preferred method steps for performing a powerup sequence according to the present invention is shown. Initially, CPU 344 waits 1128 for a "wake up" signal which is typically generated in

response to the activation of a camera 110 power on-off switch. After the "wake up" signal is generated, power manager 342 determines 1130 whether power supply 356 is generating enough operating power to start camera 110. If sufficient operating power is present, power manager 342 starts 1132 power supply 356 in normal mode with the main batteries 358 providing operating power to power supply 356 which then responsively provides the operating power to main power bus 362 and also to secondary power bus 364. Next, power manager 342 determines 1134 whether the generated operating power is maintaining a sufficient voltage level.

If operating power is sufficient in camera 110, power manager 342 then determines 1136 whether a RESUME bit has been set. In the preferred embodiment, CPU 344 sets the RESUME bit in response to a power failure in order to indicate that CPU 344 should not be reset in a subsequent powerup of camera 110. If the RESUME bit has been set, power manager 342 restarts 838 the CPU 344. The CPU 344 then resumes 1140 normal operation of DRAM 346 and also resumes 1142 the camera 110 process which was interrupted by the intervening power failure.

If the RESUME bit has not been set, then power manager 342 restarts 1144 the CPU 344 and issues 1146 a CPU 344 reset. CPU 344 then resumes 1148 normal operation of DRAM 346 and boots 1150 the computer 110 system using the system configuration 408 routine. Next, CPU 344 determines 1152 whether a MSAVE bit has been set in power manager 342. In the preferred embodiment, CPU 344 sets the MSAVE bit to specify that RAM disk 532 contains image data that should be saved upon restart of computer 118. If the MSAVE bit has not been set, computer 118 formats 1154 a new RAM disk 532. CPU 344 then runs 1158 control application 400 for normal operation of camera 110. In step 1152, if the MSAVE bit has been set, then CPU 344 recovers and mounts 1156 RAM disk 532. CPU 344 then runs 1158 control application 400 for normal operation of camera 110. The powerup process of FIG. 11 then ends.

The invention has been explained above with reference to a preferred embodiment. Other embodiments will be apparent to those skilled in the art in light of this disclosure. For example, power manager 342 may contain various registers 646 other than those discussed above in the preferred embodiment. Therefore, these and other variations upon the preferred embodiment are intended to be covered by the present invention, which is limited only by the appended claims.

What is claimed is:

1. A system for managing power conditions in a digital camera device, comprising:
 - a processor coupled to said digital camera device for controlling said digital camera device; and
 - a power manager coupled to said processor, said power manager including registers for containing status information, interrupt information, and control information;
 - said power manager providing said status information, said interrupt information, and said control information to said processor for controlling said digital camera device.
2. The system of claim 1 wherein said conditions include a low power level condition to which said processor responsively performs a powerdown sequence and a powerup sequence to protect data including captured image data in said digital camera device.

3. The system of claim 1 wherein said power manager uses a sensor device to obtain said power state information from a power source coupled to said digital camera device.

4. The system of claim 1 wherein said power manager further comprises a control register for storing said control information, an interrupt register for storing said interrupt information, and a condition register for storing said status information.

5. The system of claim 1 wherein said control information indicates the contents of memory devices coupled to said digital camera device.

6. The system of claim 5 wherein said control information includes a RESUME bit and a MSAVE bit which said power manager uses to indicate shutdown states for said digital camera device.

7. The system of claim 1 wherein said power manager further comprises a state machine, a LCD generator and a bus interface.

8. The system of claim 1 wherein said processor generates selected interrupts to control said digital camera device in response to said interrupt information from said power manager.

9. A method for managing power conditions in a digital camera device, comprising the steps of:

- controlling said digital camera device with a processor coupled to said digital camera device;
- storing status information, interrupt information, and control information in a power manager coupled to said processor; and
- providing said status information, said interrupt information, and said control information to said processor for controlling said digital camera device.

10. The method of claim 9 wherein said conditions include a low power level condition to which said processor responsively performs a powerdown sequence and a powerup sequence to protect data including captured image data in said digital camera device.

11. The method of claim 9 wherein said power manager uses a sensor device to obtain said power state information from a power source coupled to said digital camera device.

12. The method of claim 9 wherein said power manager further comprises a control register for storing said control information, an interrupt register for storing said interrupt information, and a condition register for storing said status information.

13. The method of claim 9 wherein said control information indicates the contents of memory devices coupled to said digital camera device.

14. The method of claim 13 wherein said control information includes a RESUME bit and a MSAVE bit which said power manager uses to indicate shutdown states for said digital camera device.

15. The method of claim 9 wherein said power manager further comprises a state machine, a LCD generator and a bus interface.

16. The method of claim 9 wherein said processor generates selected interrupts to control said digital camera device in response to said interrupt information from said power manager.

17. A computer-readable medium comprising program instructions for managing power conditions in a digital camera device by performing the steps of:

- controlling said digital camera device with a processor coupled to said digital camera device;
- storing status information, interrupt information, and control information in a power manager coupled to said processor; and

15

providing said status information, said interrupt information, and said control information to said processor for controlling said digital camera device.

18. A system for managing power conditions in a digital camera device, comprising:

means for controlling said digital camera device with a processor coupled to said digital camera device,

16

means for storing status information, interrupt information, and control information in a power manager coupled to said processor; and

means for providing said status information, said interrupt information, and said control information to said processor for controlling said digital camera device.

* * * * *