



- (51) **International Patent Classification:**
G06F 13/40 (2006.01) *G06F 15/78* (2006.01)
G06F 13/42 (2006.01)
- (21) **International Application Number:**
PCT/US2016/030945
- (22) **International Filing Date:**
5 May 2016 (05.05.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/170,904 4 June 2015 (04.06.2015) US
14/865,005 25 September 2015 (25.09.2015) US
- (71) **Applicant:** INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) **Inventors:** KLINGLESMTIH, Michael T.; 3311 SW 1st Ave, Portland, Oregon 97239 (US). HUNSAKER, Mikal C.; 220 Imad Place, El Dorado Hills, California 95762 (US). KNOLLA, William; 1900 Prairie City Rd, M/S: FM6-107, Folsom, California 95630 (US). SINGH, Hartej; 1232 Meredith Way, Folsom, California 95630 (US).
- (74) **Agents:** ROZMAN, Mark J. et al.; Trop, Pruner & Hu, P.C., 1616 S. Voss Rd., Ste. 750, Houston, Texas 77057-2631 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) **Title:** HANDLING A PARTITION RESET IN A MULTI-ROOT SYSTEM




FIG. 12

(57) **Abstract:** In one embodiment, a method includes: receiving, via a sideband interface of a multi-root agent associated with a first root space and a second root space, a reset prepare signal to inform the multi-root agent that the first root space is to be reset; sending, via the sideband interface, an acknowledgement signal to acknowledge the reset prepare signal; receiving one or more transactions for the first root space from a fabric coupled to the multi-root agent; and terminating the one or more transactions responsive to the reset prepare signal, where the first root space is in a reset state when the one or more transactions are received. Other embodiments are described and claimed.

**Declarations under Rule 4.17:**

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))

HANDLING A PARTITION RESET IN A MULTI-ROOT SYSTEM

[0001] This application claims priority to United States Provisional Patent Application No. 62/170,904, filed on June 4, 2015 and United States Patent Application No. 14/865,005, filed on September 25, 2015, in the names of Michael T. Klinglesmith, Mikal C. Hunsaker, William Knolla, and Hartej Singh, entitled HANDLING A PARTITION RESET IN A MULTI-ROOT SYSTEM, the disclosures of which are hereby incorporated by reference.

Background

[0002] Mainstream processor chips, both in high performance and low power segments, are increasingly integrating additional functionality such as graphics, display engines, security engines, PCIe™ ports (i.e., ports in accordance with the Peripheral Component Interconnect Express (PCI Express™ (PCIe™)) Specification Base Specification version 2.0 (published 2007) (this and related specifications are referred to hereafter as the PCIe™ specification) and other PCIe™ based peripheral devices, while maintaining legacy support for devices compliant with a PCI specification such as the Peripheral Component Interconnect (PCI) Local Bus Specification, version 3.0 (published 2002) (this and related specifications are referred to hereafter as the PCI specification).

[0003] Such designs are highly segmented due to varying requirements from the server, desktop, mobile, embedded, ultra-mobile and mobile Internet device segments. Different markets seek to use single chip system-on-chip (SoC) solutions that combine at least some of processor cores, memory controllers, input/output controllers and other segment specific acceleration elements onto a single chip. However, designs that accumulate these features are slow to emerge due to the difficulty of integrating different intellectual property (IP) blocks on a single die. This is especially so, as IP blocks can have various requirements and design uniqueness, and can require many specialized wires, communication protocols and so forth to enable their incorporation into an SoC. As a result, each SoC or other advanced semiconductor device that is developed requires a great amount of design complexity and customization to incorporate different IP blocks into a single device.

This is so, as a given IP block typically needs to be re-designed to accommodate interface and signaling requirements of a given SoC.

Brief Description of the Drawings

[0004] FIG. 1 is a block diagram of a basic interconnect architecture in accordance with an embodiment of the present invention.

[0005] FIG. 2 is a block diagram of further details of an interconnect architecture in accordance with an embodiment of the present invention.

[0006] FIG. 3 is a high level block diagram of a SoC in accordance with an embodiment of the present invention.

[0007] FIG. 4 is a block diagram of a system in accordance with another embodiment of the present invention.

[0008] FIG. 5 is a block diagram of a sideband interconnection in accordance with an embodiment of the present invention.

[0009] FIG. 6 is a block diagram of a multi-root system in accordance with an embodiment.

[0010] FIG. 7 is a block diagram of a multi-root system in accordance with another embodiment.

[0011] FIG. 8 is a block diagram of a multi-root system in accordance with yet another embodiment.

[0012] FIG. 9 is a more detailed block diagram of a portion of a SoC in accordance with an embodiment of the present invention.

[0013] FIG. 10 is a block diagram of another SoC in accordance with an embodiment the present invention.

[0014] FIG. 11 is a flow diagram of a method in accordance with an embodiment of the present invention.

[0015] FIG. 12 is a flow diagram of a method in accordance with another embodiment of the present invention.

[0016] FIG. 13 is a block diagram of an example system with which embodiments may be used.

[0017] FIG. 14 is a block diagram of a representative computer system.

Detailed Description

[0018] Embodiments can be used in many different types of systems. As examples, implementations described herein may be used in connection with semiconductor devices such as processors or other semiconductor devices that can be fabricated on a single semiconductor die. In particular implementations, the device may be a system-on-chip (SoC) or other advanced processor or chipset that includes various homogeneous and/or heterogeneous processing agents, and additional components such as networking components, e.g., routers, controllers, bridge devices, devices, memories and so forth.

[0019] Some implementations may be used in a semiconductor device that is designed according to a given specification such as an integrated on-chip system fabric (IOSF) specification issued by a semiconductor manufacturer to provide a standardized on-die interconnect protocol for attaching intellectual property (IP) blocks within a chip, including a SoC. Such IP blocks can be of varying types, including general-purpose processors such as in-order or out-of-order cores, fixed function units, graphics processors, IO controllers, display controllers, media processors among many others. By standardizing an interconnect protocol, a framework is thus realized for a broad use of IP agents in different types of chips. Accordingly, not only can the semiconductor manufacturer efficiently design different types of chips across a wide variety of customer segments, it can also, via the specification, enable third parties to design logic such as IP agents to be incorporated in such chips. And furthermore, by providing multiple options for many facets of the interconnect protocol, reuse of designs is efficiently accommodated. Although embodiments are described herein in connection with this IOSF specification, understand the scope of the present invention is not limited in this regard and embodiments can be used in many different types of systems.

[0020] Referring now to FIG. 1, shown is a block diagram of a basic interconnect architecture in accordance with an embodiment of the present invention. As shown in FIG. 1, system 10 may be a portion of a system-on-chip or any other semiconductor device such as a highly integrated processor complex or an integrated IO hub, and includes a fabric 20 that acts as an interconnect between various components. In the implementation shown, these components include IP agents 30 and 40, which can be independent IP blocks to provide various functionality such as compute capabilities, graphics capabilities, media processing capabilities and so forth. These IP agents are thus IP blocks or logical devices having an interface that is compliant with the IOSF specification, in one embodiment. As further seen, fabric 20 also interfaces to a bridge 50. Although not shown for ease of illustration in the embodiment of FIG. 1, understand that bridge 50 may act as an interface to other system components, e.g., on the same chip or on one or more different chips.

[0021] As will be described further below, each of the elements shown in FIG. 1, namely the fabric, the IP agents, and the bridge may include one or more interfaces to handle communication of various signals. These interfaces may be defined according to the IOSF specification, which defines signals for communication on these interfaces, protocols used for information exchange between agents, arbitration and flow control mechanisms used to initiate and manage information exchange, supported address decoding and translation capabilities, messaging for in-band or out-of-band communication, power management, test, validation and debug support.

[0022] The IOSF specification includes 3 independent interfaces that can be provided for each agent, namely a primary interface, a sideband message interface and a testability and debug interface (design for test (DFT), design for debug (DFD) interface). According to the IOSF specification, an agent may support any combination of these interfaces. Specifically, an agent can support 0-N primary interfaces, 0-N sideband message interfaces, and optional DFx interfaces. However, according to the specification, an agent must support at least one of these 3 interfaces.

[0023] Fabric 20 may be a hardware element that moves data between different agents. Note that the topology of fabric 20 will be product specific. As examples, a fabric can be implemented as a bus, a hierarchical bus, a cascaded hub or so forth. Referring now to FIG. 2, shown is a block diagram of further details of an interconnect architecture in accordance with an embodiment of the present invention. As shown in FIG. 2, the IOSF specification defines three distinct fabrics, namely a primary interface fabric 112, a DFX fabric 114, and a sideband fabric 116. Primary interface fabric 112 is used for all in-band communication between agents and memory, e.g., between a host processor such as a central processing unit (CPU) or other processor and an agent. Primary interface fabric 112 may further enable communication of peer transactions between agents and supported fabrics. All transaction types including memory, input output (IO), configuration, and in-band messaging can be delivered via primary interface fabric 112. Thus the primary interface fabric may act as a high performance interface for data transferred between peers and/or communications with upstream components.

[0024] In various implementations, primary interface fabric 112 implements a split transaction protocol to achieve maximum concurrency. That is, this protocol provides for a request phase, a grant phase, and a command and data phase. Primary interface fabric 112 supports three basic request types: posted, non-posted, and completions, in various embodiments. Generally, a posted transaction is a transaction which when sent by a source is considered complete by the source and the source does not receive a completion or other confirmation message regarding the transaction. One such example of a posted transaction may be a write transaction. In contrast, a non-posted transaction is not considered completed by the source until a return message is received, namely a completion. One example of a non-posted transaction is a read transaction in which the source agent requests a read of data. Accordingly, the completion message provides the requested data.

[0025] In addition, primary interface fabric 112 supports the concept of distinct channels to provide a mechanism for independent data flows throughout the system. As will be described further, primary interface fabric 112 may itself include a master interface that initiates transactions and a target interface that receives transactions.

The primary master interface can further be sub-divided into a request interface, a command interface, and a data interface. The request interface can be used to provide control for movement of a transaction's command and data. In various embodiments, primary interface fabric 112 may support PCI ordering rules and enumeration.

[0026] In turn, sideband interface fabric 116 may be a standard mechanism for communicating all out-of-band information. In this way, special-purpose wires designed for a given implementation can be avoided, enhancing the ability of IP reuse across a wide variety of chips. Thus in contrast to an IP block that uses dedicated wires to handle out-of-band communications such as status, interrupt, power management, fuse distribution, configuration shadowing, test modes and so forth, a sideband interface fabric 116 according to the IOSF specification standardizes all out-of-band communication, promoting modularity and reducing validation requirements for IP reuse across different designs. In general, sideband interface fabric 116 may be used to communicate non-performance critical information, rather than for performance critical data transfers, which typically may be communicated via primary interface fabric 112.

[0027] As further illustrated in FIG. 2, IP agents 130, 140, and 150 may each include a corresponding primary interface, a sideband interface and a DFX interface. However, as discussed above, each agent need not include every one of these interfaces, and a given IP agent may include only a single interface, in some embodiments.

[0028] Using an IOSF specification, various types of chips can be designed having a wide variety of different functionality. Referring now to FIG. 3, shown is a high level block diagram of a SoC in accordance with an embodiment of the present invention. As shown in FIG. 3, SoC 200 may include various components, all of which can be integrated on a single semiconductor die to provide for various processing capabilities at high speeds and low power, consuming a comparatively small amount of real estate. As seen in FIG. 3, SoC 200 includes a plurality of cores 205₀ – 205_n. In various embodiments, cores 205 can be relatively simple in-order cores or more complex out-of-order cores. Or a combination of in-order and out-of-

order cores can be present in a single SoC. As seen, cores 205 can be interconnected via a coherent interconnect 215, which further couples to a cache memory 210, e.g., a shared last level cache (LLC). Although the scope of the present invention is not limited in this regard, in one embodiment coherent interconnect 215 may be in accordance with the Quick Path Interconnect (QPI)[™] specification available from Intel Corporation, Santa Clara, California.

[0029] As further seen in FIG. 3, coherent interconnect 215 may communicate via a bridge 220 to a fabric 250, which may be an IOSF fabric. Coherent interconnect 215 may further communicate via an integrated memory controller 215 to an off-chip memory (not shown for ease of illustration the embodiment of FIG. 3), and further through bridge 230 to fabric 250.

[0030] As further seen in FIG. 3, various components can couple to fabric 250 including a content processing module (CPM) 240 which can be used for performing various operations such as security processing, cryptographic functions and so forth. In addition, a display processor 245 can be part of a media processing pipeline that renders video for an associated display.

[0031] As further seen, fabric 250 may further couple to an IP agent 255. Although only a single agent is shown for ease of illustration in the FIG. 3 embodiment, understand that multiple such agents are possible in different embodiments. In addition, to enable communication with other on-chip devices, fabric 250 may further communicate with a PCIe[™] controller 260 and a universal serial bus (USB) controller 265, both of which can communicate with various devices according to these protocols. Finally, shown in the embodiment of FIG. 3 is a bridge 270, which can be used to communicate with additional components of other protocols, such as an open core protocol (OCP) or an ARM advanced microcontroller bus architecture (AMBA) protocol. Although shown with these particular components in the embodiment of FIG. 3, understand that the scope of the present invention is not limited in this way and in different embodiments additional or different components may be present.

[0032] Furthermore, understand that while shown as a single die SoC implementation in FIG. 3, embodiments can further be implemented in a system in which multiple chips communicate with each other via a non-IOSF interface. Referring now to FIG. 4, shown is a block diagram of a system in accordance with another embodiment of the present invention. As shown in FIG. 4, the system may include a SoC 200', which may include many components similar to those discussed above with regard to FIG. 3, and an additional off-die interface 275. Accordingly, SoC 200' can communicate with another chip 280 which may include various functionality to enable communication between these two chips, as well as to various off-chip devices such as different peripherals according to one or more different specifications. Specifically, a second chip 280 is shown to include an off-die interface 282 to enable communication with SoC 200', and which in turn communicates with a fabric 290, which may be an IOSF fabric according to an embodiment of the present invention. As seen, fabric 290 may further be coupled to various controllers in communication with off-chip devices, including a PCIe™ controller 292, a USB controller 294, and a bridge 296.

[0033] As discussed above, in various embodiments all out-of-band communications may be via a sideband message interface. Referring now to FIG. 5, shown is a block diagram of a sideband interconnection in accordance with an embodiment of the present invention. As shown in FIG. 5, sideband interface system 175 includes multiple routers 180 and 190, which are shown in the embodiment of FIG. 5 as being coupled via a point-to-point (PTP) interconnect 185. In turn, each router can be coupled to various endpoints, which can be, for example, IP agents or other components of a given system. Specifically, router 180 couples to a plurality of endpoints 186a-186e and router 190 couples to a plurality of endpoints 196x-196z.

[0034] In various embodiments, multiple independent domains or roots may be provided in a single semiconductor device such as an SoC or other processor. More specifically, embodiments enable boot time assignment of agents of the device to one of multiple root spaces. In an embodiment, these root spaces may include a secure root space and a user root space. Hardware and software within a root

space can be considered an independent computing system. The physical hardware that implements the multiple root spaces has shared components (including in an embodiment a shared fabric, system agent, memory and other components). Every command with a primary interconnect may include a root space attribute, which may be a given field of a command.

[0035] As used herein, a root is a collection of resources including one or more cores, a host bridge, system memory, zero or more root complex integrated endpoints, zero or more root complex event collectors, and zero or more root ports. Within a root there may be an entire computer system. Software executing within that root need not be aware that there is more than one root. In a multi-root system, software may run without knowledge of the other roots, and without the need for any changes in the existing software model including microcode, basic input output system (BIOS), virtual machine monitors (VMMs), operating systems (OSs), device drivers, and applications.

[0036] A shared fabric implements a complete set of address decode logic for each root space that is independent of the other root spaces. The address decoder used for a transaction is selected based on a root space (RS) indicator (which may be implemented as a plurality of bits to identify a given one of multiple root spaces) in a command (CMD). The system agent may implement some limited address decode functionality for the alternate root spaces and a full decode for the host root space, in an embodiment. Various queue structures (e.g., channels) may be shared between root spaces or may be independent for each root space based on SoC configuration choice. Embodiments thus provide a scalable, architectural solution that can be implemented in a standard configurable fabric.

[0037] A multi-root system includes two or more roots that share one or more physical resources. One root may be designated as a privileged root. This privileged root is the first to boot and is responsible for configuring the system. The privileged root completes configuration of the other roots before they are released from reset. Any resource that is not statically assigned to a root (e.g., strap or fuse) is configured by the privileged root. The privileged root is also responsible for coordinating the re-assignment of any resources that can be dynamically re-

assigned to different roots. In an embodiment, system configuration is accomplished by programming registers to assign resources to the different roots. In some cases, resource assignment can be static for some resources and dynamic for others. Embodiments may isolate roots by associating transactions with a given root space field (RS field), to enable comparison of the RS value with the RS field.

[0038] In order to support the existing software model a root contains certain hardware resources that are owned and controlled by the root. During initialization of the hardware before software is given control of the root, the resources of any root may be accessed or configured by the privileged root. For example, a CPU may include a set of configuration registers (CRs). CRs are owned by the hardware within the root and can only be accessed from microcode executing on the cores of the root. Similarly, each CPU has a set of machine specific registers (MSRs). MSRs are accessible from software executing on the CPU cores of the root.

[0039] A single root agent is assigned to and used by a single root at any given time. This assignment can be either static or dynamic. Dynamic assignment of a single root agent may be coordinated with the software executing on the roots involved. Assignment of a single root agent to a root may be implemented by setting the RS value for the agent. The RS field in every primary command is then used in conjunction with address decode mechanisms to ensure the agent is only accessed for transactions associated with the assigned root. In some cases, a single root agent may be multi-root aware. Such agent may include a mechanism such as a register or strap for setting the RS value of the agent, which may be programmed by a privileged root. An address decode logic of the agent may then enforce root-based security policies.

[0040] If a single root agent is not multi-root aware, it is the fabric's responsibility to implement hardware logic, e.g., an address decode logic, to enable the agent to be assigned a RS value and to enforce given security policies.

[0041] In an embodiment, a multi-root agent can be exposed to two or more roots. This may be accomplished with a dedicated primary interface per root, or with a shared primary interface with dedicated channels. In the case of dedicated

interfaces, each interface is assigned an RS value. For dedicated channels, each channel may be associated with an RS value. In another case, a shared interface may be provided with a shared channel. On the shared channel, the agent associates each transaction with an RS value.

[0042] When a multi-root agent is used by more than one root concurrently, the agent may be configured with appropriate hardware (e.g., straps or registers) to expose dedicated resources to each root. In an embodiment, software running in a root and interacting with the hardware may be unaware that it is interacting with a multi-root agent.

[0043] A host bridge provides the connection between the CPU(s), system memory, and the I/O subsystem, where the I/O subsystem is south of the host bridge, and the CPU(s) are north of the host bridge. A single root host bridge contains resources for one root only. It is statically assigned to a root space. All transactions originating from the north of a single root host bridge are implicitly associated with the root space of the host bridge. Either the host bridge or a fabric below the host bridge is responsible for tagging transactions with the correct RS field, and implementing the multi-root security policies.

[0044] If the host bridge includes a bridge to a sideband interconnect, that bridge includes the correct RS value in a security attribute indicator (SAI) field. The system memory behind a single root host bridge can only be accessed with the associated root space in an embodiment. Other roots may use a piece of stolen system memory only through a multi-root non-transparent bridge. The bridge may map the transaction from the originating root space into the destination root space.

[0045] A multi-root host bridge receives and processes transactions on behalf of several roots. The division of host bridge resources between root spaces is either statically assigned in hardware or configured at initialization time. The host bridge enforces security policies to ensure the isolation of resources between different roots. A multi-root host bridge that supports assignment of various CPUs on the north to different root spaces tracks the origin of incoming transactions from the north and tags them with the correct RS field. Any cache coherent flows through the

host bridge are isolated to the originating root space. System memory of a multi-root host bridge may be divided between multiple root spaces. The host bridge is responsible for maintaining the mapping between the host physical addresses associated with different root spaces and the memory physical address used as the physical address for system memory.

[0046] In an embodiment, a multi-root non-transparent bridge is a multi-root agent connected to multiple roots. Such bridge provides one or more communication mechanisms for software in one root to interact with software in another root (e.g., mailbox registers, doorbells, memory translation windows and so forth).

[0047] In an embodiment, address decode is performed in the fabric on behalf of an agent that positively decodes against the entire RS field. If the fabric is connected to a single root agent that does not support the RS field, the fabric is responsible for implementing multi-root security policies on behalf of the agent. In an embodiment, such policies includes full decode and comparison of the RS field in the command to the RS value for the agent.

[0048] In an embodiment, a sideband fabric uses a single set of destination IDs that are independent of a root in which an endpoint resides. Any transaction that bridges from a primary interconnect to a sideband interconnect may include a RS field from the primary transaction into a given field of the sideband message, such as a SAI field to enable this information to be passed along to an appropriate destination.

[0049] In various embodiments, a shared (multi-root) fabric and one or more connected IP logic blocks may be used to support independent reset domains (e.g., resetting a host root while keeping another root powered up).

[0050] When resetting a root, shared components (e.g., an IOSF fabric) remain out of reset. In addition, the fabric drains/discards transactions associated with the domain to be reset. Agents coupled to the fabric assigned entirely to one domain use a wire signal, namely a power okay (POK) signal, to the fabric to inform the fabric that they are to be power gated and to enter into reset. It becomes the fabric's responsibility to terminate transactions targeting that agent. In contrast, an agent coupled to the fabric that supports multiple domains receives a sideband reset

prepare message, which may be received from a power control agent. Responsive to this message, the agent becomes responsible to terminate transactions for the domain in reset, while continuing to accept transactions for the domain not in reset. Thus rather than having a dedicated fabric for each domain, a shared fabric across multiple independent reset domains may be used to independently enable reset of the different roots.

[0051] At the beginning of a reset flow (reset entry), a power management agent (e.g., a power management controller (PMC)) is responsible for sending a reset prepare message on a sideband channel to every agent that supports multiple reset domains. The encoding in the message informs the agent which partition is being reset. At the end of the reset flow (reset exit), the PMC is responsible for sending a boot prepare message on the sideband channel to every agent that supports multiple reset domains (and responsive to which the agent may send a boot prepare acknowledge signal or message). At the beginning of the reset flow (reset entry), the PMC is responsible for sending a ForcePwrGatePOK message on the sideband channel to every agent that does not support multiple reset domains.

[0052] A single root agent is an agent that is only aware of a single root space. Such agent may be configured to use the POK signal to disconnect from the fabric prior to a partition reset. When the agent receives the ForcePwrGatePOK message from the PMC, it advances its primary interface ISM to the idle state and then asserts the POK signal to the fabric.

[0053] A multi-root agent is an agent that supports multiple reset domains. Such agent may be configured to receive a reset prepare message before one of the roots is reset. After receiving the reset prepare message, the agent responds with a reset prepare acknowledgement message on the sideband channel. After the reset prepare acknowledgement message, the agent is responsible for terminating any transaction from the domain that is being reset. In an embodiment, posted and completion transactions may be dropped to be terminated. In an embodiment an unsupported request (UR) completion message may be returned for a non-posted transaction to be terminated. The agent may continue to terminate transactions for

the root space of the reset domain until it receives a boot prepare message from the PMC, in an embodiment.

[0054] In an embodiment, a fabric is not directly informed that a partition reset is occurring. The fabric's role is to respond correctly to the POK signal of all agents. When the fabric detects an agent's POK signal asserted, the fabric may be configured to terminate any transaction targeting the agent. In an embodiment, there may be specific rules for the transition of POK with respect to agent and fabric ISMs to ensure that the fabric does not received a POK transition in the middle of a transaction.

[0055] Embodiments thus enable a SoC implementing multiple independent roots or domains (e.g., user and secure) to be independently reset within a system.

[0056] Referring now to FIG. 6, shown is a block diagram of a multi-root system in accordance with an embodiment. In the embodiment shown in FIG. 6, multi-root system 300 may be implemented as a SoC or other multicore processor, e.g., implemented in a single processor package. In other cases, system 300 may be implemented as a set of discrete components such as adapted to a motherboard or other circuit board of a given computing system, which can range from a small portable device such as smartphone, tablet computer, or so forth, to a server computer. As illustrated, system 300 includes a first root 310 and a second root 360. In an embodiment, these different roots can be of different privilege levels such that one of the roots (e.g., first root 310) is configured as a privileged root while second root 360 is configured as an unprivileged root. Note the terms "root" and "root space" are used interchangeably herein, and thus roots 310 and 360 also are referred to herein as root spaces.

[0057] In the privileged root, various operations to configure a system and to execute privileged code, e.g., of an operating system, hypervisor or so forth may be implemented, while in turn the unprivileged root may execute a variety of unprivileged user applications. Note that the blocks shown in clear form are of first root 310 while the blocks shown in black form are associated with second root 360. In turn, shared resources are illustrated in FIG. 6 as shaded blocks.

[0058] As illustrated, first root 310 includes a plurality of CPUs 315a-315b. CPUs 315 couple to a host bridge 320 that in turn is coupled to a system memory 325, which in an embodiment may be a DRAM. As seen host bridge 320 may include a converged security manageability engine (CSME). Host bridge 320 further couples to a fabric 350, which in an embodiment may be an IOSF fabric that is a shared resource between the multiple roots. As such, shared fabric 350 provides interconnection to agents of both root spaces. Still with reference to first root 310, fabric 350 couples to a root port 330 and a plurality of root complex integrated endpoints 340a and 340b.

[0059] Referring now to second root 360, a single CPU 365 is present. Of course in other embodiments, multiple CPUs may be associated with this root also. CPU 365 couples to a host bridge 370 that in turn couples to a system memory 375. In turn, host bridge 370 further couples to shared fabric 350, and with regard to components of second root 360, fabric 350 couples to root complex integrated endpoints 380a and 380b. Understand while shown with this particular implementation in the embodiment of FIG. 6, many variations and alternatives are possible.

[0060] Referring now to FIG. 7, shown is a block diagram of a multi-root system in accordance with another embodiment. In the embodiment shown in FIG. 7, multi-root system 400 may be implemented as a SoC or other multicore processor, e.g., implemented in a single processor package, or as a set of discrete components such as adapted to a motherboard or other circuit board of a given computing system. As illustrated, system 400 includes a first root 410 and a second root 460. In an embodiment, these different roots can be of different privilege levels such that one of the roots (e.g., first root 410) is configured as a privileged root while second root 460 is configured as an unprivileged root. As above, the blocks shown in clear form are of first root 410 while the blocks shown in black form are associated with second root 460. In turn, shared resources are illustrated in FIG. 7 as shaded blocks.

[0061] As illustrated, first root 410 includes a plurality of CPUs 415a-415b. CPUs 415 couple to a host bridge 452 that in turn is coupled to a system memory 454. Host bridge 452 further couples to a fabric 450, which in an embodiment may be an

IOSF fabric. As seen in FIG. 7, host bridge 452, system memory 454 and fabric 450 are shared resources between the multiple roots. As such, shared fabric 450 provides interconnection to agents of both root spaces. Still with reference to first root 410, fabric 450 couples to a root port 430 and a plurality of root complex integrated endpoints 440a and 440b.

[0062] Referring now to second root 460, a single CPU 465 is present. Of course in other embodiments multiple CPUs may be associated with this root also. CPU 465 also couples to host bridge 452. With regard to components of second root 460, fabric 450 couples to root complex integrated endpoints 480a and 480b. Understand while shown with this particular implementation in the embodiment of FIG. 7, many variations and alternatives are possible.

[0063] Referring now to FIG. 8, shown is a block diagram of a multi-root system in accordance with yet another embodiment. In the embodiment shown in FIG. 8, multi-root system 500 may be implemented as a SoC or other multicore processor or as a set of discrete components. As illustrated, system 500 includes a first root 510 and a second root 560. In an embodiment, these different roots can be of different privilege levels such that one of the roots (e.g., first root 510) is configured as a privileged root while second root 560 is configured as an unprivileged root. As above, the blocks shown in clear form are of first root 510 while the blocks shown in black form are associated with second root 560. In turn, shared resources are illustrated in FIG. 8 as shaded blocks.

[0064] As illustrated, first root 510 includes a plurality of CPUs 515a-515b. CPUs 515 couple to a host bridge 520 that in turn is coupled to a system memory 525. Host bridge 520 further couples to a fabric 550, which in an embodiment may be an IOSF fabric. As seen in FIG. 8, host bridge 520 and system memory 525 are resources of first root 510. Shared fabric 550 is a shared resource which provides interconnection to agents of both root spaces through a shared bridge 555. Still with reference to first root 510, fabric 550 couples to a root port 530 and a root complex integrated endpoint 540a.

[0065] Referring now to second root 560, a single CPU 565 is present. Of course in other embodiments multiple CPUs may be associated with this root also. CPU 565 couples to a host bridge 570 and a system memory 575, which are resources of second root 560. With regard to components of second root 560, fabric 590 couples to root complex integrated endpoints 580b and 580c, and as shown via bridge 555 and fabric 550, also couples to root complex integrated endpoint 580a. Understand while shown with this particular implementation in the embodiment of FIG. 8, many variations and alternatives are possible.

[0066] Although the SoCs above are at a high level, understand that additional functionality may be present. Referring now to FIG. 9, shown is a more detailed block diagram of a portion of a SoC in accordance with an embodiment of the present invention. As shown in FIG. 9, the portion of SoC 700 shown may correspond to non-core portions coupled below a memory controller hub or other interface logic that can in turn interface to multiple processor cores, as well as to system memory.

[0067] Thus as seen, an off-die interface 710 (which in one embodiment can be a direct media interface (DMI)) may couple to a hub 715, e.g., an input/output hub that in turn provides communication between various peripheral devices. Although not shown for ease of illustration in FIG. 9, understand that various engines such as a manageability engine and a virtualization engine can also be directly coupled to hub 715.

[0068] To provide connection to multiple buses, which may be multi-point or shared buses in accordance with the IOSF specification, an IOSF controller 720 may couple between hub 715 and bus 730, which may be an IOSF bus that thus incorporates elements of the fabric as well as routers. In the embodiment shown in FIG. 9, first IOSF bus 730 may have coupled to it various controllers to provide for control of off-chip devices. Specifically, seen is a PCI controller 722, a SATA controller 724, and a USB controller 726. In turn, a second IOSF bus 750 may couple to a system management bus 752 and to a real time clock 754.

[0069] As further seen in FIG. 9, first IOSF bus 730 may couple to an IOSF bridge 735 for both primary and sideband information that in turn provides interconnection to a third bus 740, e.g., of a different protocol, to which various controllers and components may be attached. In the embodiment shown in FIG. 9, such components include a flash controller 741 to provide an interface to a non-volatile memory, a legacy device 742, which may implement various legacy functions, e.g., of a PCI specification and further may include an interrupt controller and timer. In addition, interfaces for audio 743, USB 744, gigabyte Ethernet (GbE) 745, serial peripheral interface (SPI) 746 and PCI 747 may all be provided. Although shown with this particular implementation in the embodiment of FIG. 9, understand the scope of the present invention is not limited in this regard.

[0070] Still other implementations are possible. Referring now to FIG. 10, shown is a block diagram of another SoC in accordance with an embodiment the present invention. As shown in FIG. 10, SoC 800 may be configured for use, e.g., in server systems. As seen in FIG. 10, SoC may include a platform controller hub (PCH) 840, which may generally include components such as seen in the embodiment of FIG. 9. Namely, multiple IOSF buses 730 and 740 may be present, along with a bridge 735 to couple the buses. Bus 730 may include various agents coupled to it, including a PCIe controller 722, SATA controller 724, and a USB controller 726. In turn, via an IOSF controller 720, communication may occur via an additional bus 718, which may communicate with upstream devices, such as cores or other processing units (not shown for ease of illustration in the embodiment of FIG. 10).

[0071] As further seen in FIG. 10, for providing communications with other server-based components, an additional IOSF bus 820 may be provided, which in turn can communicate with an IOSF controller 822 and an upstream switch port 824 (e.g., an X16 port) that may be coupled to an upstream bus 825. Also coupled to bus 820 may be multiple downstream switch ports 826 and 828.

[0072] Furthermore, to enable communications, e.g., with storage units of a server-based system, a switch port 830 may couple between bus 820 and another IOSF bus 850, which in turn may be coupled to a storage controller unit (SCU) 855, which may be a multi-function device for coupling with various storage devices.

[0073] To effect partition resets in a multi-root system, first the system is configured for multi-root operation. To this end, multiple root spaces may be defined. Such definition may be implemented on configuration of an SoC or other processor, or the definition can be performed dynamically. For example, one or more agents can be dynamically associated with different root spaces at different times.

[0074] Referring now to FIG. 11, shown is a flow diagram of a method in accordance with an embodiment of the present invention. As shown in FIG. 11, method 1000 may be performed within an SoC to enable multi-root space configuration and operation. As seen, method 1000 begins by associating at least a first processor and at least a first agent (which may be an endpoint or other agent of an SoC) with a first root space value (block 1010). To this end, each such component can be associated with this first root space value by storage of the given value in a storage included in or associated with the component, such as a configuration storage. By this association, a first root space is defined to include these components. Similar operations may occur at block 1020 to associate at least a second processor and at least a second agent with a second root space value, to define a second root space.

[0075] Still referring to FIG. 11, these components all may be coupled together, e.g., via a shared fabric, which may be implemented in an embodiment as an IOSF fabric. At this point, the SoC is configured for normal operation.

[0076] Thus as shown in FIG. 11, during operation a transaction may be received in this shared fabric (block 1040). Based at least in part on a root space value of a root space field of the transaction, the shared fabric may route the transaction to a given one of the multiple root spaces (and more specifically to a particular destination agent within the root space) (block 1050). Understand while shown at this high level in the embodiment of FIG. 11, many variations and alternatives are possible.

[0077] In such a multi-root space system, partition resets can be performed, as described herein. Referring now to FIG. 12, shown is a flow diagram of a method in accordance with another embodiment of the present invention. As shown in FIG. 12, method 1100 may be implemented within an SoC that provides for multi-root space

operation. As seen, method 1100 begins by receiving a reset prepare signal for a first root space (block 1110). In an embodiment, this reset prepare signal may be received in a multi-root agent via a sideband interface, and which may be sent from a router in turn coupled to a power controller of the SoC. Responsive to receipt of this reset prepare signal, the agent may send an acknowledgement signal (block 1120), via the sideband interface to the router (that in turn provides the acknowledgement to the power controller). At this point, the root space associated with the reset prepare signal, namely the first root space, may thus enter a reset state. As such, any multi-root space devices may be controlled to properly operate in this reset state.

[0078] Still with reference to FIG. 12, next it is determined whether a boot prepare signal is received for the first root space (diamond 1130). This signal may be received in the agent via a sideband interface. If such signal is not received, control next passes to diamond 1150 where it is determined whether a transaction is received in the agent for this first root space. If so, control passes to block 1160 where the transaction may be terminated, as this first root space is in the reset state.

[0079] Continuing with FIG. 12 at some point the agent will receive (via the sideband interface) a boot prepare signal for the first root space (as determined at diamond 1130). Responsive to receipt of this signal, control passes to block 1140 where a boot prepare acknowledgement signal is sent via the sideband interface, to enable the acknowledgement to be sent from a router to the power controller. Understand that responsive to this boot prepare acknowledgement signal, the first root space may now return to an active state. Thus when incoming transactions for the first root space are received in the agent, they can be handled, e.g., as indicated in the transaction. Understand while shown at this high level in the embodiment of FIG. 12, many variations and alternatives are possible.

[0080] Referring now to FIG. 13, shown is a block diagram of example system with which embodiments may be used. In the illustration of FIG. 13, system 1300 may be mobile low-power system such as a tablet computer, 2:1 tablet, phablet or other convertible or standalone tablet system. As illustrated, a SoC 1310 is present and may be configured to operate as an application processor for the device. SoC 1310

may be configured to include multiple root spaces as described herein, and in which one or more agents may be dynamically associated with one or more of the root spaces. Such multi-root-aware device in turn may be configured to implement partial reset operations as described herein, in cooperation with a power controller and/or fabric.

[0081] A variety of devices may couple to SoC 1310. In the illustration shown, a memory subsystem includes a flash memory 1340 and a DRAM 1345 coupled to SoC 1310. In addition, a touch panel 1320 is coupled to the SoC 1310 to provide display capability and user input via touch, including provision of a virtual keyboard on a display of touch panel 1320. To provide wired network connectivity, SoC 1310 couples to an Ethernet interface 1330. A peripheral hub 1325 is coupled to SoC 1310 to enable interfacing with various peripheral devices, such as may be coupled to system 1300 by any of various ports or other connectors.

[0082] In addition to internal power management circuitry and functionality within SoC 1310, a power management integrated circuit (PMIC) 1380 is coupled to SoC 1310 to provide platform-based power management, e.g., based on whether the system is powered by a battery 1390 or AC power via an AC adapter 1395. In addition to this power source-based power management, PMIC 1380 may further perform platform power management activities based on environmental and usage conditions. Still further, PMIC 1380 may communicate control and status information to SoC 1310 to cause various power management actions within SoC 1310.

[0083] Still referring to FIG. 13, to provide for wireless capabilities, a WLAN unit 1350 is coupled to SoC 1310 and in turn to an antenna 1355. In various implementations, WLAN unit 1350 may provide for communication according to one or more wireless protocols.

[0084] As further illustrated, a plurality of sensors 1360 may couple to SoC 1310. These sensors may include various accelerometer, environmental and other sensors, including user gesture sensors. Finally, an audio codec 1365 is coupled to SoC 1310 to provide an interface to an audio output device 1370. Of course

understand that while shown with this particular implementation in FIG. 13, many variations and alternatives are possible.

[0085] Referring now to FIG. 14, shown is a block diagram of a representative computer system such as notebook, Ultrabook™ or other small form factor system. A processor 1410, in one embodiment, includes a microprocessor, multi-core processor, multithreaded processor, an ultra low voltage processor, an embedded processor, or other known processing element. In the illustrated implementation, processor 1410 acts as a main processing unit and central hub for communication with many of the various components of the system 1400. As one example, processor 1410 is implemented as a SoC that provides for multiple root spaces and their configuration and control, as described herein.

[0086] Processor 1410, in one embodiment, communicates with a system memory 1415. As an illustrative example, the system memory 1415 is implemented via multiple memory devices or modules to provide for a given amount of system memory.

[0087] To provide for persistent storage of information such as data, applications, one or more operating systems and so forth, a mass storage 1420 may also couple to processor 1410. In various embodiments, to enable a thinner and lighter system design as well as to improve system responsiveness, this mass storage may be implemented via a SSD or the mass storage may primarily be implemented using a hard disk drive (HDD) with a smaller amount of SSD storage to act as a SSD cache to enable non-volatile storage of context state and other such information during power down events so that a fast power up can occur on re-initiation of system activities. Also shown in FIG. 14, a flash device 1422 may be coupled to processor 1410, e.g., via a serial peripheral interface (SPI). This flash device may provide for non-volatile storage of system software, including a basic input/output software (BIOS) as well as other firmware of the system.

[0088] Various input/output (I/O) devices may be present within system 1400. Specifically shown in the embodiment of FIG. 14 is a display 1424 which may be a high definition LCD or LED panel that further provides for a touch screen 1425. In

one embodiment, display 1424 may be coupled to processor 1410 via a display interconnect that can be implemented as a high performance graphics interconnect. Touch screen 1425 may be coupled to processor 1410 via another interconnect, which in an embodiment can be an I²C interconnect. As further shown in FIG. 14, in addition to touch screen 1425, user input by way of touch can also occur via a touch pad 1430 which may be configured within the chassis and may also be coupled to the same I²C interconnect as touch screen 1425.

[0089] For perceptual computing and other purposes, various sensors may be present within the system and may be coupled to processor 1410 in different manners. Certain inertial and environmental sensors may couple to processor 1410 through a sensor hub 1440, e.g., via an I²C interconnect. In the embodiment shown in FIG. 14, these sensors may include an accelerometer 1441, an ambient light sensor (ALS) 1442, a compass 1443 and a gyroscope 1444. Other environmental sensors may include one or more thermal sensors 1446 which in some embodiments couple to processor 1410 via a system management bus (SMBus) bus.

[0090] Also seen in FIG. 14, various peripheral devices may couple to processor 1410 via a low pin count (LPC) interconnect. In the embodiment shown, various components can be coupled through an embedded controller 1435. Such components can include a keyboard 1436 (e.g., coupled via a PS2 interface), a fan 1437, and a thermal sensor 1439. In some embodiments, touch pad 1430 may also couple to EC 1435 via a PS2 interface. In addition, a security processor such as a trusted platform module (TPM) 1438 may also couple to processor 1410 via this LPC interconnect.

[0091] System 1400 can communicate with external devices in a variety of manners, including wirelessly. In the embodiment shown in FIG. 14, various wireless modules, each of which can correspond to a radio configured for a particular wireless communication protocol, are present. One manner for wireless communication in a short range such as a near field may be via a NFC unit 1445 which may communicate, in one embodiment with processor 1410 via an SMBus. Note that via this NFC unit 1445, devices in close proximity to each other can communicate.

[0092] As further seen in FIG. 14, additional wireless units can include other short range wireless engines including a WLAN unit 1450 and a Bluetooth™ unit 1452. Using WLAN unit 1450, Wi-Fi™ communications can be realized, while via Bluetooth™ unit 1452, short range Bluetooth™ communications can occur. These units may communicate with processor 1410 via a given link.

[0093] In addition, wireless wide area communications, e.g., according to a cellular or other wireless wide area protocol, can occur via a WWAN unit 1456 which in turn may couple to a subscriber identity module (SIM) 1457. In addition, to enable receipt and use of location information, a GPS module 1455 may also be present. Note that in the embodiment shown in FIG. 14, WWAN unit 1456 and an integrated capture device such as a camera module 1454 may communicate via a given link.

[0094] To provide for audio inputs and outputs, an audio processor can be implemented via a digital signal processor (DSP) 1460, which may couple to processor 1410 via a high definition audio (HDA) link. Similarly, DSP 1460 may communicate with an integrated coder/decoder (CODEC) and amplifier 1462 that in turn may couple to output speakers 1463 which may be implemented within the chassis. Similarly, amplifier and CODEC 1462 can be coupled to receive audio inputs from a microphone 1465 which in an embodiment can be implemented via dual array microphones (such as a digital microphone array) to provide for high quality audio inputs to enable voice-activated control of various operations within the system. Note also that audio outputs can be provided from amplifier/CODEC 1462 to a headphone jack 1464. Although shown with these particular components in the embodiment of FIG. 14, understand the scope of the present invention is not limited in this regard.

[0095] The following examples pertain to further embodiments.

[0096] In one example, an apparatus comprises: a semiconductor die that includes but is not limited to: a plurality of agents, at least a first subset of the plurality of agents associated with a first root space value to define a first root space and at least a second subset of the plurality of agents associated with a second root space value to define a second root space, where a first agent of the plurality of agents is

associated with the first root space and the second root space; and a fabric coupled to the plurality of agents via a fabric primary interface, the fabric primary interface including at least one master interface to initiate transactions and at least one target interface to receive transactions.

[0097] In an example, the first agent is to receive a reset prepare signal to inform the first agent that the first root space is to be reset, and responsive thereto, send an acknowledgement signal to acknowledge the reset prepare signal.

[0098] In an example, the first agent is to receive one or more transactions for the first root space and terminate the one or more transactions responsive to the reset prepare signal, where the first root space is in a reset state when the one or more transactions are received.

[0099] In an example, the first agent is to receive a boot prepare signal to inform the first agent that the reset of the first root space has completed, and responsive thereto, send a boot prepare acknowledge signal.

[0100] In an example, the first agent is to receive and handle one or more additional transactions for first root space received in the first agent after receipt of the boot prepare signal.

[0101] In an example, the first agent, while the first root space is in the reset state, is to drop a posted transaction and a completion transaction, and send an unsupported request completion for a non-posted transaction.

[0102] In an example, the first agent is to receive and handle at least one transaction for the second root space received in the first agent while the first root space is in the reset state.

[0103] In an example, the apparatus further comprises a router having a sideband message interface to couple to the first agent via a sideband interconnect, and where the first agent is to receive the reset prepare signal via the sideband interconnect from a power controller of a SoC comprising the semiconductor die.

[0104] In an example, a second agent of the second subset of the plurality of agents, responsive to a power gate signal from a power controller to indicate that the

second root space is to be reset, is to send a signal to the fabric to inform the fabric that the second agent is to be power gated, where the fabric is to terminate one or more transactions directed to the second agent, responsive to receipt of the signal from the second agent.

[0105] In an example, the fabric is to enable independent reset of the plurality of root spaces.

[0106] In another example, a method comprises: receiving, via a sideband interface of a multi-root agent associated with a first root space and a second root space, a reset prepare signal to inform the multi-root agent that the first root space is to be reset; sending, via the sideband interface, an acknowledgement signal to acknowledge the reset prepare signal; receiving one or more transactions for the first root space from a fabric coupled to the multi-root agent; and terminating the one or more transactions responsive to the reset prepare signal, where the first root space is in a reset state when the one or more transactions are received.

[0107] In an example, the method further includes: receiving, via the sideband interface, a boot prepare signal, the boot prepare signal to inform the multi-root agent that the reset has completed; sending a boot prepare acknowledgement signal; and handling one or more additional transactions via first root space received in the multi-root agent after receipt of the boot prepare signal.

[0108] In an example, the method further comprises dropping a posted transaction and a completion transaction to terminate the posted transaction and the completion transaction.

[0109] In an example, the method further comprises sending an unsupported request completion for a non-posted transaction to terminate the non-posted transaction.

[0110] In an example, the method further comprises receiving and handling at least one transaction for the second root space while the first root space is in the reset state.

[0111] In an example, the method further comprises receiving the reset prepare signal from a power controller of the SoC.

[0112] In another example, a computer readable medium including instructions is to perform the method of any of the above examples.

[0113] In another example, a computer readable medium including data is to be used by at least one machine to fabricate at least one integrated circuit to perform the method of any one of the above examples.

[0114] In another example, an apparatus comprises means for performing the method of any one of the above examples.

[0115] In another example, a SoC comprises: at least one core to execute instructions; a coherent interconnect coupled to the at least one core; a plurality of agents, including a first agent associated with a first root space value, a second agent associated with a second root space value, and a third agent associated with the first root space value and the second root space value, the first root space value to define a first root space and the second root space value to define a second root space; a fabric coupled to the coherent interconnect and to at least some of the plurality of agents; and a router coupled to at least some of the plurality of agents via a sideband interconnect.

[0116] In an example, the first agent is to receive one or more transactions for the first root space from the fabric and terminate the one or more transactions when the first root space is in a reset state, and after receipt of a boot prepare signal to inform the first agent that a reset of the first root space has completed, to receive and handle one or more additional transactions for first root space.

[0117] In an example, at least one component is coupled to the fabric via a bridge, where the at least one component is of an open core protocol (OCP) or an ARM advanced microcontroller bus architecture (AMBA) protocol.

[0118] In an example, the at least one core comprises a first in-order core and a first out-of-order core.

[0119] Understand that various combinations of the above examples are possible.

[0120] Embodiments may be used in many different types of systems. For example, in one embodiment a communication device can be arranged to perform the various methods and techniques described herein. Of course, the scope of the present invention is not limited to a communication device, and instead other embodiments can be directed to other types of apparatus for processing instructions, or one or more machine readable media including instructions that in response to being executed on a computing device, cause the device to carry out one or more of the methods and techniques described herein.

[0121] Embodiments may be implemented in code and may be stored on a non-transitory storage medium having stored thereon instructions which can be used to program a system to perform the instructions. Embodiments also may be implemented in data and may be stored on a non-transitory storage medium, which if used by at least one machine, causes the at least one machine to fabricate at least one integrated circuit to perform one or more operations. Still further embodiments may be implemented in a computer readable storage medium including information that, when manufactured into a SoC or other processor, is to configure the SoC or other processor to perform one or more operations. The storage medium may include, but is not limited to, any type of disk including floppy disks, optical disks, solid state drives (SSDs), compact disk read-only memories (CD-ROMs), compact disk rewritables (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[0122] While the present invention has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations therefrom. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this present invention.

What is claimed is:

- 1 1. An apparatus comprising:
2 a semiconductor die including but not limited to:
3 a plurality of agents, at least a first subset of the plurality of agents
4 associated with a first root space value to define a first root space and at least a
5 second subset of the plurality of agents associated with a second root space value to
6 define a second root space, wherein a first agent of the plurality of agents is
7 associated with the first root space and the second root space; and
8 a fabric coupled to the plurality of agents via a fabric primary interface,
9 the fabric primary interface including at least one master interface to initiate
10 transactions and at least one target interface to receive transactions.
- 1 2. The apparatus of claim 1, wherein the first agent is to receive a reset prepare
2 signal to inform the first agent that the first root space is to be reset, and responsive
3 thereto, send an acknowledgement signal to acknowledge the reset prepare signal.
- 1 3. The apparatus of claim 2, wherein the first agent is to receive one or more
2 transactions for the first root space and terminate the one or more transactions
3 responsive to the reset prepare signal, wherein the first root space is in a reset state
4 when the one or more transactions are received.
- 1 4. The apparatus of claim 3, wherein the first agent is to receive a boot prepare
2 signal to inform the first agent that the reset of the first root space has completed,
3 and responsive thereto, send a boot prepare acknowledge signal.
- 1 5. The apparatus of claim 4, wherein the first agent is to receive and handle one
2 or more additional transactions for first root space received in the first agent after
3 receipt of the boot prepare signal.

1 6. The apparatus of claim 3, wherein the first agent, while the first root space is
2 in the reset state, is to drop a posted transaction and a completion transaction, and
3 send an unsupported request completion for a non-posted transaction.

1 7. The apparatus of claim 3, wherein the first agent is to receive and handle at
2 least one transaction for the second root space received in the first agent while the
3 first root space is in the reset state.

1 8. The apparatus of claim 2, further comprising a router having a sideband
2 message interface to couple to the first agent via a sideband interconnect, and
3 wherein the first agent is to receive the reset prepare signal via the sideband
4 interconnect from a power controller of a system on a chip (SoC) comprising the
5 semiconductor die.

1 9. The apparatus of claim 1, wherein a second agent of the second subset of the
2 plurality of agents, responsive to a power gate signal from a power controller to
3 indicate that the second root space is to be reset, is to send a signal to the fabric to
4 inform the fabric that the second agent is to be power gated, wherein the fabric is to
5 terminate one or more transactions directed to the second agent, responsive to
6 receipt of the signal from the second agent.

1 10. The apparatus of claim 1, wherein the fabric is to enable independent reset of
2 the plurality of root spaces.

1 11. A method comprising:
2 receiving, via a sideband interface of a multi-root agent associated with a first
3 root space and a second root space, a reset prepare signal to inform the multi-root
4 agent that the first root space is to be reset;
5 sending, via the sideband interface, an acknowledgement signal to
6 acknowledge the reset prepare signal;
7 receiving one or more transactions for the first root space from a fabric
8 coupled to the multi-root agent; and

9 terminating the one or more transactions responsive to the reset prepare
10 signal, wherein the first root space is in a reset state when the one or more
11 transactions are received.

1 12. The method of claim 11, further comprising:
2 receiving, via the sideband interface, a boot prepare signal, the boot prepare
3 signal to inform the multi-root agent that the reset has completed;
4 sending a boot prepare acknowledgement signal; and
5 handling one or more additional transactions via first root space received in
6 the multi-root agent after receipt of the boot prepare signal.

1 13. The method of claim 11, further comprising dropping a posted transaction and
2 a completion transaction to terminate the posted transaction and the completion
3 transaction.

1 14. The method of claim 11, further comprising sending an unsupported request
2 completion for a non-posted transaction to terminate the non-posted transaction.

1 15. The method of claim 11, further comprising receiving and handling at least
2 one transaction for the second root space while the first root space is in the reset
3 state.

1 16. The method of claim 11, further comprising receiving the reset prepare signal
2 from a power controller of the SoC.

1 17. A machine-readable storage medium including machine-readable instructions,
2 when executed, to implement a method as claimed in any one of claims 11 to 16.

1 18. An apparatus comprising means to perform a method as claimed in any one
2 of claims 11 to 16.

1 19. A system-on-chip (SoC) comprising:

2 at least one core to execute instructions;
3 a coherent interconnect coupled to the at least one core;
4 a plurality of agents, including a first agent associated with a first root space
5 value, a second agent associated with a second root space value, and a third agent
6 associated with the first root space value and the second root space value, the first
7 root space value to define a first root space and the second root space value to
8 define a second root space;
9 a fabric coupled to the coherent interconnect and to at least some of the
10 plurality of agents; and
11 a router coupled to at least some of the plurality of agents via a sideband
12 interconnect.

1 20. The SoC of claim 19, wherein the first agent is to receive one or more
2 transactions for the first root space from the fabric and terminate the one or more
3 transactions when the first root space is in a reset state, and after receipt of a boot
4 prepare signal to inform the first agent that a reset of the first root space has
5 completed, to receive and handle one or more additional transactions for first root
6 space.

1 21. The SoC of claim 19, further comprising at least one component coupled to
2 the fabric via a bridge, wherein the at least one component is of an open core
3 protocol (OCP) or an ARM advanced microcontroller bus architecture (AMBA)
4 protocol.

1 22. The SoC of claim 19, wherein the at least one core comprises a first in-order
2 core and a first out-of-order core.

1/14


10

FIG. 1




FIG. 2




FIG. 3

4/14


200`

FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8

700




FIG. 9

800




FIG. 10

11/14


1000

FIG. 11




FIG. 12

1300




FIG. 13

14/14




FIG. 14

A. CLASSIFICATION OF SUBJECT MATTER**G06F 13/40(2006.01)i, G06F 13/42(2006.01)i, G06F 15/78(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 13/40; G06F 15/173; G06F 13/368; G06F 13/36; G06F 13/42; G06F 13/00; G06F 15/78

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: multi-root, second root space, reset prepare signal, acknowledgement, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2013-0086296 A1 (SRIDHAR LAKSHMANAMURTHY et al.) 04 April 2013 See paragraphs [0016]-[0035], [0058], and [0067]; claims 1, 4, 12, and 19; and figure 3.	1-22
Y	US 2004-0098475 A1 (CARL ZEITLER et al.) 20 May 2004 See paragraphs [0008]-[0009], [0042] and [0061]-[0062]; claim 1; and figure 8.	1-22
Y	EP 1195689 A2 (BROADCOM CORPORATION) 10 April 2002 See paragraphs [0008] and [0019]; claims 1 and 11; and figure 1.	4-5, 9-10, 12, 19-22
A	US 2015-0019788 A1 (INTEL CORPORATION) 15 January 2015 See paragraphs [0047]-[0050]; and figures 3 and 7.	1-22
A	US 2012-0166690 A1 (JACK REGULA) 28 June 2012 See paragraphs [0097]-[0100] and figures 8-9.	1-22



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

17 August 2016 (17.08.2016)

Date of mailing of the international search report

17 August 2016 (17.08.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/030945

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0086296 A1	04/04/2013	US 2014-0258583 A1 US 8713240 B2	11/09/2014 29/04/2014
US 2004-0098475 A1	20/05/2004	AU 2003-286813 A1 CA 2506522 A1 EP 1563398 A2 JP 2006-506736 A US 7577755 B2 WO 2004-046946 A2 WO 2004-046946 A3	15/06/2004 03/06/2004 17/08/2005 23/02/2006 18/08/2009 03/06/2004 23/09/2004
EP 1195689 A2	10/04/2002	DE 60132961 D1 DE 60132961 T2 EP 1195689 A3 EP 1195689 B1 US 2004-0049620 A1 US 6633938 B1 US 6865633 B2	10/04/2008 26/02/2009 19/11/2003 27/02/2008 11/03/2004 14/10/2003 08/03/2005
US 2015-0019788 A1	15/01/2015	US 2013-0138858 A1 US 9053251 B2 US 9213666 B2	30/05/2013 09/06/2015 15/12/2015
US 2012-0166690 A1	28/06/2012	US 8521941 B2	27/08/2013