



(43) International Publication Date
4 April 2013 (04.04.2013)

(10) International Publication Number
WO 2013/048856 A2

- (51) International Patent Classification: Not classified
- (21) International Application Number:
PCT/US2012/056255
- (22) International Filing Date:
20 September 2012 (20.09.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/248,232 29 September 2011 (29.09.2011) US
- (71) Applicant (for all designated States except US): **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95052 (US).
- (72) Inventors: **LAKSHMANAMURTHY, Sridhar**; 1614 Grosbeak Avenue, Sunnyvale, California 94087 (US). **ADLER, Robert P.**; 2668 Forbes Avenue, Santa Clara, California 95051 (US). **HUNSAKER, Mikal C.**; 1646 Hedland Place, El Dorado Hills, California 95762 (US). **KLINGLESMTIH, Michael T.**; 3311 SW 1st Avenue, Portland, Oregon 97239 (US). **FANNING, Blaise**; 989 Smith Way, Folsom, California 95630 (US). **TAMARI, Eran**; Apartment 12, 14 Rimalt Street, 52281 Ramat Gan, TA (IL). **MURRAY, Joseph**; 11645 N 50th Street, Scottsdale, Arizona 85254 (US). **VERMA, Rohit R.**; 47867 Wabana Street, Fremont, California 94539 (US).
- (74) Agent: **ROZMAN, Mark J.**; 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, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, 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, 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, ML, MR, NE, SN, TD, TG).

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:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: COMMON IDLE STATE, ACTIVE STATE AND CREDIT MANAGEMENT FOR AN INTERFACE

(57) Abstract: In one embodiment, the present invention includes method for entering a credit initialization state of an agent state machine of an agent coupled to a fabric to initialize credits in a transaction credit tracker of the fabric. This tracker tracks credits for transaction queues of a first channel of the agent for a given transaction type. The agent may then assert a credit initialization signal to cause credits to be stored in the transaction credit tracker corresponding to the number of the transaction queues of the first channel of the agent for the first transaction type. Other embodiments are described and claimed.



WO 2013/048856 A2

**COMMON IDLE STATE, ACTIVE STATE AND
CREDIT MANAGEMENT FOR AN INTERFACE**

Background

[0001] 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) (hereafter 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) (hereafter the PCI specification).

[0002] 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.

[0003] One feature of interconnects in a PCIe™ system is the concept of various protocol layers, including a transaction layer, a link layer, and a physical layer. Responsibilities with respect to various signaling, power management, and credit mechanisms can be segmented across these layers. As a result, any one of these

functions requires execution of state machines or other logic in each of these layers, increasing complexity.

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 state transition diagram of an agent state machine in accordance with an embodiment of the present invention.

[0009] FIG. 6 is a state transition diagram of a fabric state machine in accordance with an embodiment of the present invention.

[0010] FIG. 7 is a more detailed block diagram of a portion of a SoC in accordance with an embodiment of the present invention.

[0011] FIG. 8 is a block diagram of another SoC in accordance with an embodiment of the present invention.

Detailed Description

[0012] Embodiments provide a state machine, referred to herein as an idle state machine (ISM), that can be a common architected interface to support multiple operations including credit initialization, active mode data exchange, and power management. This single state machine, which can be implemented via an agent-based state machine and a corresponding fabric-based state machine, can provide for all interactions via an on-die interconnect between an agent and a fabric with regard to these operations. In this way, all of the above functions can be implemented in a single layer, since on-die communications between different

agents can avoid link integrity mechanisms commonly used for off-die communications (e.g., as implemented in a link layer).

[0013] In general, an idle state machine in accordance with an embodiment of the present invention provides a handshake between an agent and the fabric. Specific states and their transitions with regard to these two state machines are described in detail below. Fundamentally, the ISM states can be partitioned into three categories: states to support credit initialization; states to support the active mode of the interface; and idle state support. For each of these three major categories, the ISM provides an architected handshake between agent and fabric that supports the major activities that are associated with each category. The result is that the ISM provides a solution that neatly ties together the concepts of credit management, active mode data transmission, and idle state power and clock gating into a single state machine.

[0014] 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.

[0015] 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.

[0016] 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.

[0017] 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.

[0018] 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.

[0019] 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.

[0020] 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.

[0021] 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.

[0022] 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.

[0023] 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.

[0024] 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.

[0025] 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.

[0026] 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.

[0027] 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.

[0028] 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.

[0029] In one embodiment, each agent may include an idle state machine on each primary and sideband message interface that it supports. For every ISM implemented by an agent, there is a corresponding ISM implemented in the fabric (or a router). The following Tables 1-4 set forth a description of the different states for the agent and fabric/router ISMs (Tables 1 and 3, respectively), and state transitions for the ISMs (Tables 2 and 4, respectively). These tables also refer to the state transition diagrams of FIGS. 5 and 6 for the agent ISM and fabric ISM, respectively. Note that the discussion is with reference to a single agent and fabric interconnection, but understand that multiple agents and fabrics can be present, with

multiple agents interfacing with a single fabric (each by its own agent ISM and corresponding fabric ISM).

Table 1

State	Description
CREDIT_REQ	Agent enters this state during run-time to initialize or re-initialize credits.
CREDIT_INIT	The agent initializes interface credits in this state. The fabric ISM can be in CREDIT_ACK state when an agent starts initializing.
CREDIT_DONE	The agent has completed initializing its interface credits.
IDLE	The agent is IDLE. An agent can locally gate its interface clock and deassert its interface clock request when it is in this state.
ACTIVE_REQ	The agent requests to use the interface or is acknowledging the request of the fabric to use the interface.
ACTIVE	The interface is active. In this state the agent can issue puts or credit updates. In this state, the agent can receive puts, grants, and credit updates.
IDLE_REQ	The agent no longer requires the use of the interface, has met the ISM IDLE request rules and is requesting to enter IDLE.

Table 2

Present State	Next State	Description
CREDIT_REQ	CREDIT_INIT	The fabric ISM is in the CREDIT_ACK state.
CREDIT_INIT	CREDIT_DONE	The agent has completed credit initialization on its target interface and the fabric ISM is in the CREDIT_INIT state. The agent ensures that the rules for interface clock gating have been met before entering the CREDIT_DONE state.
CREDIT_DONE	IDLE	The fabric ISM is in the IDLE state, indicating that

		it has completed request credit initialization.
IDLE	CREDIT_REQ	All target queues are empty, all credits have been returned and the agent seeks to re-initialize credits or the fabric ISM is in the CREDIT_REQ state. The agent ensures that the rules for interface clock gating have been met before entering the CREDIT_REQ state.
IDLE	ACTIVE_REQ	The agent seeks the use of the interface to communicate with the fabric OR the fabric ISM is in the ACTIVE_REQ state AND the agent ISM arc from IDLE to ACTIVE is not implemented. The agent ensures the clock is valid and that all interface clock gating rules are followed before exiting IDLE.
IDLE	ACTIVE	The fabric ISM is in the ACTIVE_REQ state.
ACTIVE_REQ	ACTIVE	The fabric ISM is in the ACTIVE_REQ state.
ACTIVE_REQ	CREDIT_REQ	The fabric ISM is in the CREDIT_REQ state.
ACTIVE	IDLE_REQ	The agent satisfies all conditions under which this transition may or must be taken.
IDLE_REQ	IDLE	The fabric ISM is in the IDLE state.
IDLE_REQ	ACTIVE	The fabric ISM is in the IDLE_NAK state.

Table 3

State	Description
CREDIT_REQ	The fabric enters this state to request to initialize or re-initialize credits.
CREDIT_INIT	During this state, the fabric initializes its

	interface credits.
CREDIT_ACK	Fabric acknowledges an agent's request for credit initialization.
IDLE	Fabric is allowed to locally gate its interface clock only in this state.
ACTIVE_REQ	The fabric requests use of the interface or is acknowledging the request of the agent to the interface. In this state the fabric can receive puts and credit updates.
ACTIVE	The interface is active. In this state the fabric can issue puts, grants, credit updates and can decrement target credits. In this state, the fabric can receive puts and credit updates.
IDLE_NAK	Fabric is rejecting an agent's request to enter IDLE. In this state the fabric can issue puts, grants, credit updates and can decrement target credits. In this state, the fabric can receive puts, and credit updates.

Table 4

Present State	Next State	Description
CREDIT_REQ	CREDIT_ACK	The agent ISM is in the CREDIT_REQ state. Following interface reset de-assertion, the fabric ISM may choose to add additional delay before acknowledging the agent ISM CREDIT_REQ state to ensure that reset has fully propagated to all destinations.
CREDIT_ACK	CREDIT_INIT	The agent ISM is in the CREDIT_INIT state.
CREDIT_INIT	IDLE	The agent ISM is in the CREDIT_DONE state and the fabric has completed credit initialization for that agent on that interface.
IDLE	ACTIVE_REQ	The agent ISM is in the ACTIVE_REQ state OR the fabric requires the interface to communicate

		with the agent and the agent is in the IDLE state. The fabric ensures that the interface clock is valid and will not be gated after exiting the IDLE state
IDLE	CREDIT_REQ	The agent ISM is in the CREDIT_REQ state. The fabric ensures that the interface clock is valid and will not be gated after exiting the IDLE state.
ACTIVE_REQ	ACTIVE	The agent is in the ACTIVE state
ACTIVE_REQ	CREDIT_REQ	The agent ISM is in the CREDIT_REQ state. Note: as the system is quiesced prior to credit initialization, this arc should never need to be taken, but is included here for robustness.
ACTIVE	IDLE	The agent ISM is in the IDLE_REQ state, the fabric does not require the interface to communicate with the agent, and the fabric has not put a new transaction to the agent at the same time the agent indicated it was in the IDLE_REQ state.
ACTIVE	IDLE_NAK	The agent ISM is in the IDLE_REQ state, and the fabric requires the use of the IOSF interface to communicate with the agent, or the fabric put a transaction to the agent at the same time that the agent indicated it was in the IDLE_REQ state.
IDLE_NAK	ACTIVE	The agent ISM is in the ACTIVE state.

[0030] In the embodiment of the ISMs shown in FIGS. 5 and 6, the interface credit initialization phase can be used to initialize various interface credits including request

and transaction credits within a system. This phase may occur shortly after an interface signal reset has de-asserted and involves the agent ISM states of CREDIT_REQ, CREDIT_INIT, and CREDIT_DONE and the fabric ISM states of CREDIT_REQ, CREDIT_ACK, and CREDIT_INIT. In some embodiments, the ISM state machine seamlessly supports two different reset states, IDLE and CREDIT_REQ. The choice of which state is the reset state can be implementation specific, with the IDLE state being the preferred choice due to favorable reset and power isolation properties. If the reset state is chosen to be the IDLE state, then an internal CREDITS_INITIALIZED flag is included within the agent to indicate that a transition from IDLE to CREDIT_REQ is to occur after reset in order to initialize credits.

[0031] Once both the agent and fabric are in the CREDIT_REQ state, the fabric ISM transitions to the CREDIT_ACK state when it is ready for credit initialization to begin. The agent ISM will then transition to the CREDIT_INIT state and proceed with transaction credit initialization. In one embodiment, the transaction credit initialization state may cause the agent to send information to the fabric regarding its available transaction queues so that the fabric can initialize its credit transaction counters with this information. In one embodiment, a credit initialization grant can be issued from a request put arbiter of the agent. The number of credits (e.g., corresponding to the number of transaction queues it has) may be indicated by maintaining this credit initialization grant signal active for a number of clock cycles corresponding to the number of transaction queues. In other implementations, the number of credits can be indicated by interface-specific signaling.

[0032] The fabric ISM next transitions to the CREDIT_INIT state where it will proceed with request credit initialization. In one embodiment, the request credit initialization state may cause the fabric to send information to the agent regarding its available request queues so that the agent can initialize its credit request counters with this information. In one embodiment, a credit initialization grant can be issued from a fabric arbiter of the fabric, and the number of credits (e.g., corresponding to the number of request queues it has) may be indicated by maintaining this credit

initialization grant signal active for a number of clock cycles corresponding to the number of request queues.

[0033] When the agent has completed request credit initialization, it transitions the agent ISM to the CREDIT_DONE state. When the fabric arbiter has completed request credit initialization (and the agent ISM is in the CREDIT_DONE state), it will return to the IDLE state, at which time the agent ISM will return to the IDLE state. The fabric can reset its transaction credit counters on all arcs into the CREDIT_ACK state. Similarly, the agent can reset its request credit counters on the arc from CREDIT_REQ to CREDIT_INIT.

[0034] As further seen in FIGS. 5 and 6, both the agent and fabric ISMs include an architected ACTIVE state in which agents and fabrics may initiate transaction puts, which are requests to initiate transactions to each other. Once credit initialization has been completed and both the agent and fabric ISMs are in the IDLE state, it is possible for either side to initiate the sequence of state transitions that will result in both ISMs being in the ACTIVE state that allows them to send transactions.

[0035] If the agent is the side that first desires to send a transaction, it will first transition from the IDLE to ACTIVE_REQ state. When the fabric ISM determines that the agent ISM is in the ACTIVE_REQ state, it will also transition from the IDLE to ACTIVE_REQ state. The sequence completes with the agent ISM transitioning to the ACTIVE state when it determines that the fabric ISM is in the ACTIVE_REQ state, and the fabric ISM transitions to the ACTIVE state when it see the agent in the ACTIVE state.

[0036] If the fabric is the side that first desires to send a transaction, it will first transition from the IDLE to ACTIVE_REQ state. Once the agent ISM sees the fabric ISM in the ACTIVE_REQ state, it will also transition from the IDLE to ACTIVE_REQ, or potentially, directly to ACTIVE state, e.g., based on an agent implementation. If the agent transitions first to ACTIVE_REQ, the sequence completes with the agent ISM transitioning to the ACTIVE state. Once the fabric determines that the agent is in the ACTIVE state, it transitions to the ACTIVE state in order to complete the handshake.

[0037] In various embodiments, to expand the opportunities for idle state power management, the ISM protocol defines two specific conditions when an agent and its ISM must initiate entry into the IDLE state from the ACTIVE state. As a first set of preconditions to this required IDLE state entry by the agent ISM, the following conditions first are met: the interface has been inactive for a minimum of N clocks after the fabric ISM has entered into the ACTIVE state; the interface is not in the middle of a transaction; and the agent has advertised a minimum number of credits. As a second set of preconditions to this required IDLE state entry, the following conditions first are met: the agent ISM must transition from ACTIVE to IDLE_REQ state between clocks N+1 and M if it has returned all credits to the fabric, has no outstanding inbound or outbound non-posted requests, and the conditions for transitioning to the IDLE_REQ state have been satisfied for the preceding M-N clocks. That is, an agent that has not returned all credits or has outstanding non-posted requests may (but is not required to) transition to the IDLE_REQ state.

[0038] Note that the parameter values of M and N can be protocol specific. In one embodiment, N can be chosen as 16 and M as 32. When the agent has transitioned to its IDLE_REQ state, the fabric may either acknowledge this request to go to the IDLE state by transitioning to its IDLE state or it may not acknowledge (NAK) the request by transitioning to its IDLE_NAK state. If the fabric acknowledges the agent's request to go to the IDLE state, then the agent ISM transitions from the IDLE_REQ to IDLE state. If the fabric NAKs the agent's IDLE request, then the agent transitions from the IDLE_REQ state back to the ACTIVE state, and the wait for the M and N clocks can begin again.

[0039] The fabric not acknowledges (NAK) the agent's request if, as seen on the interface, the fabric puts a transaction in the same cycle or any cycle after the agent's ISM has transitioned to the IDLE_REQ state. In this way, a race condition can be prevented that might lead to performance degradation if the agent is allowed to go to the IDLE state due to a previous idle request that came at the same time that the fabric gave the agent more work to do.

[0040] Once the agent's ISM has reached the IDLE state, it may internally gate its own clock. In order to support trunk level clock gating of all devices that share an

interface clock with the agent, the agent may also follow an additional clock request protocol to signal to the rest of the fabric that it no longer needs its clock. The clock request protocol is controlled by two signals, CLKREQ and CLKACK. Upon entering the IDLE state, the agent determines whether both CLKREQ and CLKACK are active (e.g., logic high). Once the agent determines that it no longer needs its interface clock, it may indicate this to the fabric by deasserting the CLKREQ signal. In response to the CLKREQ signal being deasserted, the fabric deasserts the CLKACK signal as an acknowledgement of the agent's request.

[0041] Although the SoCs of FIGS. 3 and 4 are at a high level, understand that additional functionality may be present. Referring now to FIG. 7, 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. 7, 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.

[0042] 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. 7, understand that various engines such as a manageability engine and a virtualization engine can also be directly coupled to hub 715.

[0043] 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. 7, 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.

[0044] As further seen in FIG. 7, 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. 7, 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. 7, understand the scope of the present invention is not limited in this regard.

[0045] Still other implementations are possible. Referring now to FIG. 8, shown is a block diagram of another SoC in accordance with an embodiment the present invention. As shown in FIG. 8, SoC 800 may be configured for use, e.g., in server systems. As seen in FIG. 8, SoC may include a platform controller hub (PCH) 840, which may generally include components such as seen in the embodiment of FIG. 7. 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. 8).

[0046] As further seen in FIG. 8, 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.

[0047] 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.

[0048] 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. 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.

[0049] 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 fabric coupled to each of a plurality of agents via an interface having a fabric
3 arbiter and a plurality of queues each associated with a corresponding queue of the
4 agent, the fabric further including a fabric idle state machine (ISM) to perform credit
5 management, active mode management, and idle state management; and
6 the plurality of agents coupled to the fabric, each of the plurality of agents
7 having an interface including a request put arbiter, a plurality of queues, and a
8 plurality of credit trackers each to track credits for one of the queues of the
9 corresponding fabric interface, and further including an agent ISM to perform credit
10 management, active mode management, and idle state management, via a
11 handshake protocol with the fabric ISM.
- 1 2. The apparatus of claim 1, wherein each of the interfaces of the fabric further
2 includes a plurality of credit trackers each to track credits for one of the queues of
3 the corresponding agent interface.
- 1 3. The apparatus of claim 2, wherein the agent ISM is to enter a credit
2 initialization state to initialize credits in the credit trackers of the fabric interface.
- 1 4. The apparatus of claim 2, wherein the credit initialization is to store a number
2 of credits in a first one of the credit trackers corresponding to the number of queues
3 of a first channel of the agent interface for a first transaction type.
- 1 5. The apparatus of claim 4, wherein the agent ISM is to assert a credit
2 initialization signal for a number of clock cycles corresponding to the number of
3 queues of the first channel of the agent interface for the first transaction type.
- 1 6. The apparatus of claim 1, wherein the fabric ISM is to enter a credit
2 initialization state to initialize credits in the credit trackers of a first agent.

1 7. The apparatus of claim 6, wherein the credit initialization is to store a number
2 of credits in a first one of the credit trackers corresponding to the number of queues
3 of a first channel of the fabric interface for a first transaction type, the fabric ISM to
4 assert a credit initialization signal for a number of clock cycles corresponding to the
5 number of queues of the first channel of the fabric interface for the first transaction
6 type.

1 8. The apparatus of claim 1, wherein the agent ISM is to transition from an idle
2 state to an active state to enable a first request put corresponding to a first
3 transaction to be sent to the fabric, wherein the transition is to cause the fabric ISM
4 to transition from an idle state to an active state.

1 8. The apparatus of claim 1, further comprising an on-die interconnect to couple
2 the fabric and a first agent, wherein the fabric ISM and the agent ISM operate at a
3 physical layer without a link layer.

1 10. The apparatus of claim 1, wherein the agent ISM is to transition from an active
2 state to an idle request state if the interface of the agent has been inactive for a
3 predetermined number of clock cycles, the interface is not communicating a
4 transaction, and the agent has advertised minimum credits.

1 11. The apparatus of claim 10, wherein the agent ISM is thereafter to enter an idle
2 state, and to clock gate a clock of the agent when in the idle state.

1 12. The apparatus of claim 11, wherein the agent is to deassert a clock request
2 signal to the fabric to indicate the clock gate.

1 13. A method comprising:
2 entering a credit initialization state of an agent state machine of an agent
3 coupled to a fabric to initialize credits in a credit tracker of the fabric that tracks
4 credits for queues of a first channel of the agent for a first transaction type; and

5 asserting a credit initialization signal to cause credits to be stored in the credit
6 tracker corresponding to the number of the queues of the first channel of the agent
7 for the first transaction type.

1 14. The method of claim 13, further comprising entering the credit initialization
2 state of the agent state machine responsive to a fabric state machine of the fabric
3 being in a credit acknowledge state.

1 15. The method of claim 13, entering a credit initialization state of the fabric state
2 machine to initialize credits in a credit tracker of the agent that tracks credits for
3 queues of a first channel of the fabric for the first transaction type.

1 16. The method of claim 15, further comprising storing a number of credits in a
2 first one of the credit trackers corresponding to the number of the queues of the
3 fabric for the first channel for the first transaction type, the fabric state machine to
4 assert a credit initialization signal for a number of clock cycles corresponding to the
5 number of the queues.

1 17. The method of claim 13, further comprising entering an idle state of the agent
2 state machine after performing the credit initialization.

1 18. The method of claim 17, further comprising thereafter entering an active
2 request state of the agent state machine to cause the fabric state machine to enter
3 into an active request state, and thereafter entering an active state of the agent state
4 machine to cause the fabric state machine to enter into an active state, and
5 thereafter sending a request for a transaction from an interface of the agent to an
6 interface of the fabric if a credit tracker of the agent associated with a type of the
7 transaction indicates an available queue in the fabric.

1 19. The method of claim 18, further comprising transitioning from an active state
2 to an idle request state of the agent state machine if the interface of the agent has

3 been inactive for a predetermined number of clock cycles, the interface of the agent
4 is not communicating a transaction, and the agent has advertised minimum credits.

1 20. The method of claim 13, further comprising asserting the credit initialization
2 signal for a number of clock cycles corresponding to the number of the queues of the
3 agent of the first channel for the transaction type.

1 21. A system-on-chip (SoC) comprising:
2 a fabric having a fabric state machine;
3 a plurality of agents coupled to the fabric, each of the plurality of agents
4 including an agent state machine, wherein the fabric state machine and a first agent
5 state machine of a first agent are to enter into a plurality of handshake credit states
6 to initialize credits in a credit tracker of the fabric that tracks credits for queues of a
7 first channel of the first agent for a first transaction type and to initialize credits in a
8 credit tracker of the first agent that tracks credits for queues of a first channel of the
9 fabric for the first transaction type, enter into an idle state after the plurality of
10 handshake credit states, and responsive to an indication from the first agent state
11 machine, enter into a plurality of handshake active states to enable communication
12 of a transaction between the first agent and the fabric.

1 22. The SoC of claim 21, wherein the first agent state machine is to transition
2 from the idle state to an active state to enable a request put corresponding to a first
3 transaction to be sent to the fabric.

1 23. The SoC of claim 22, wherein the transition is to cause the fabric state
2 machine to transition from the idle state to the active state.

10




FIG. 1




FIG. 2

200

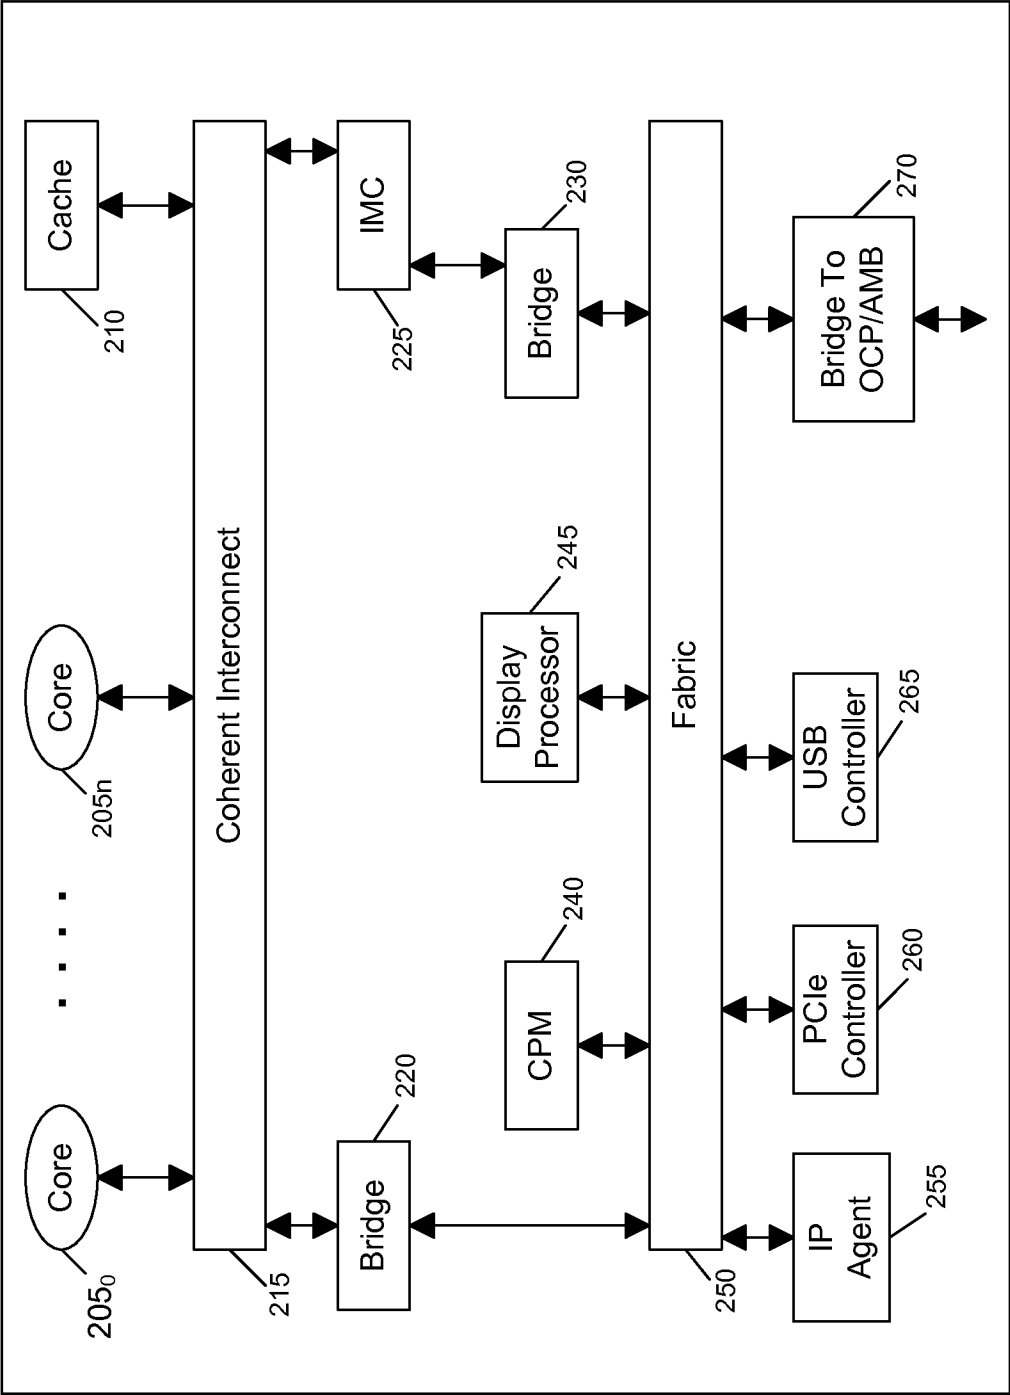


FIG. 3




FIG. 4

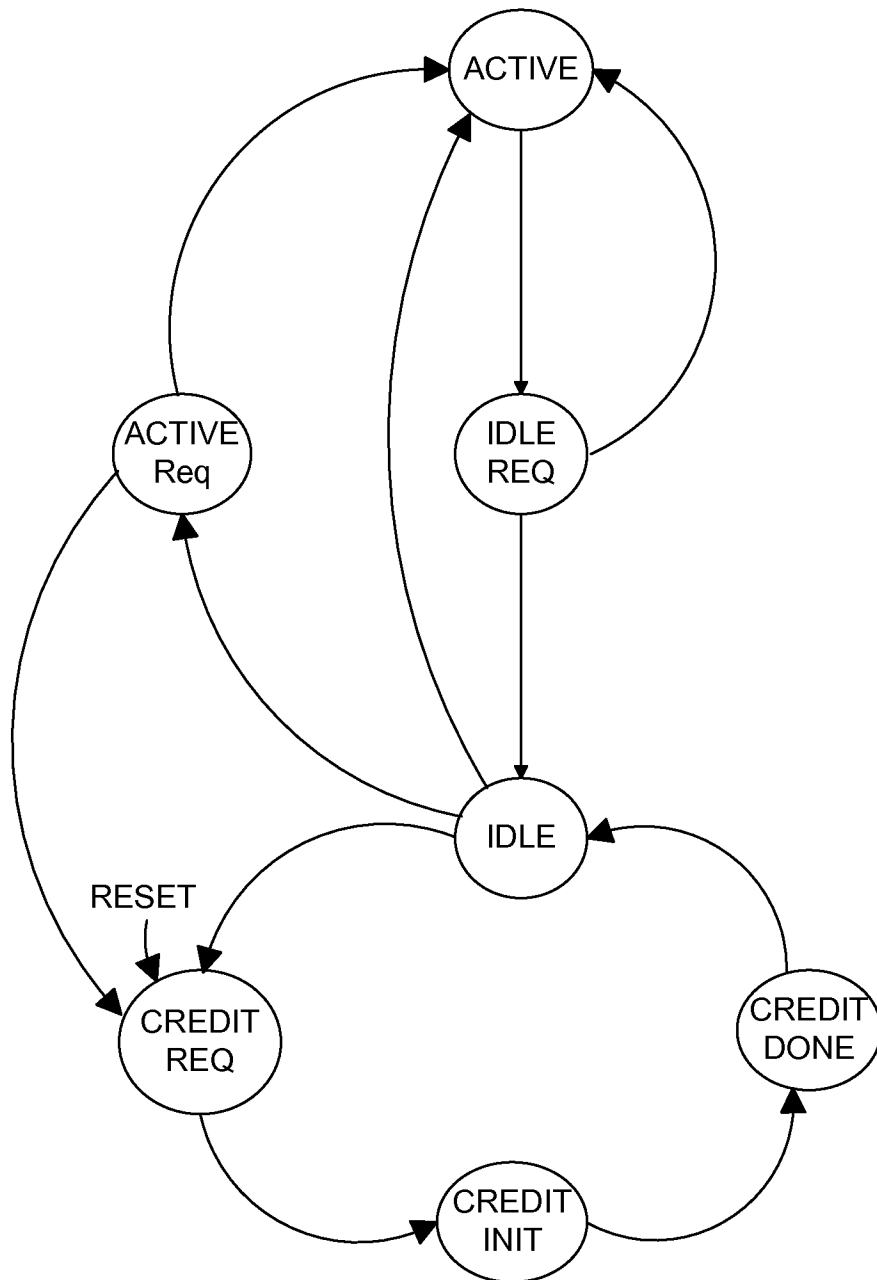


FIG. 5

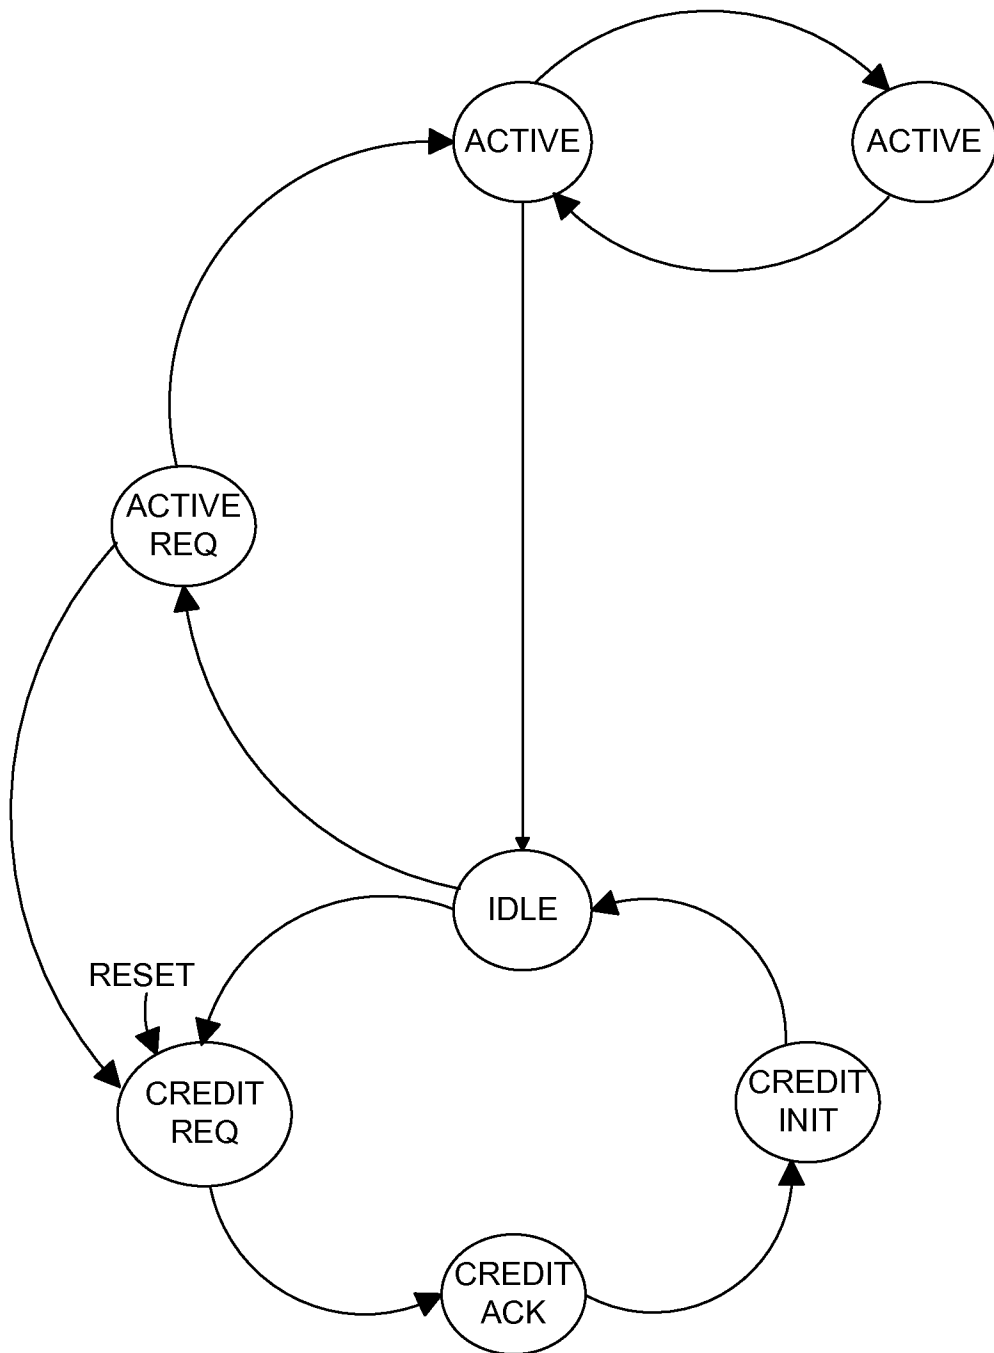


FIG. 6

700




FIG. 7

800

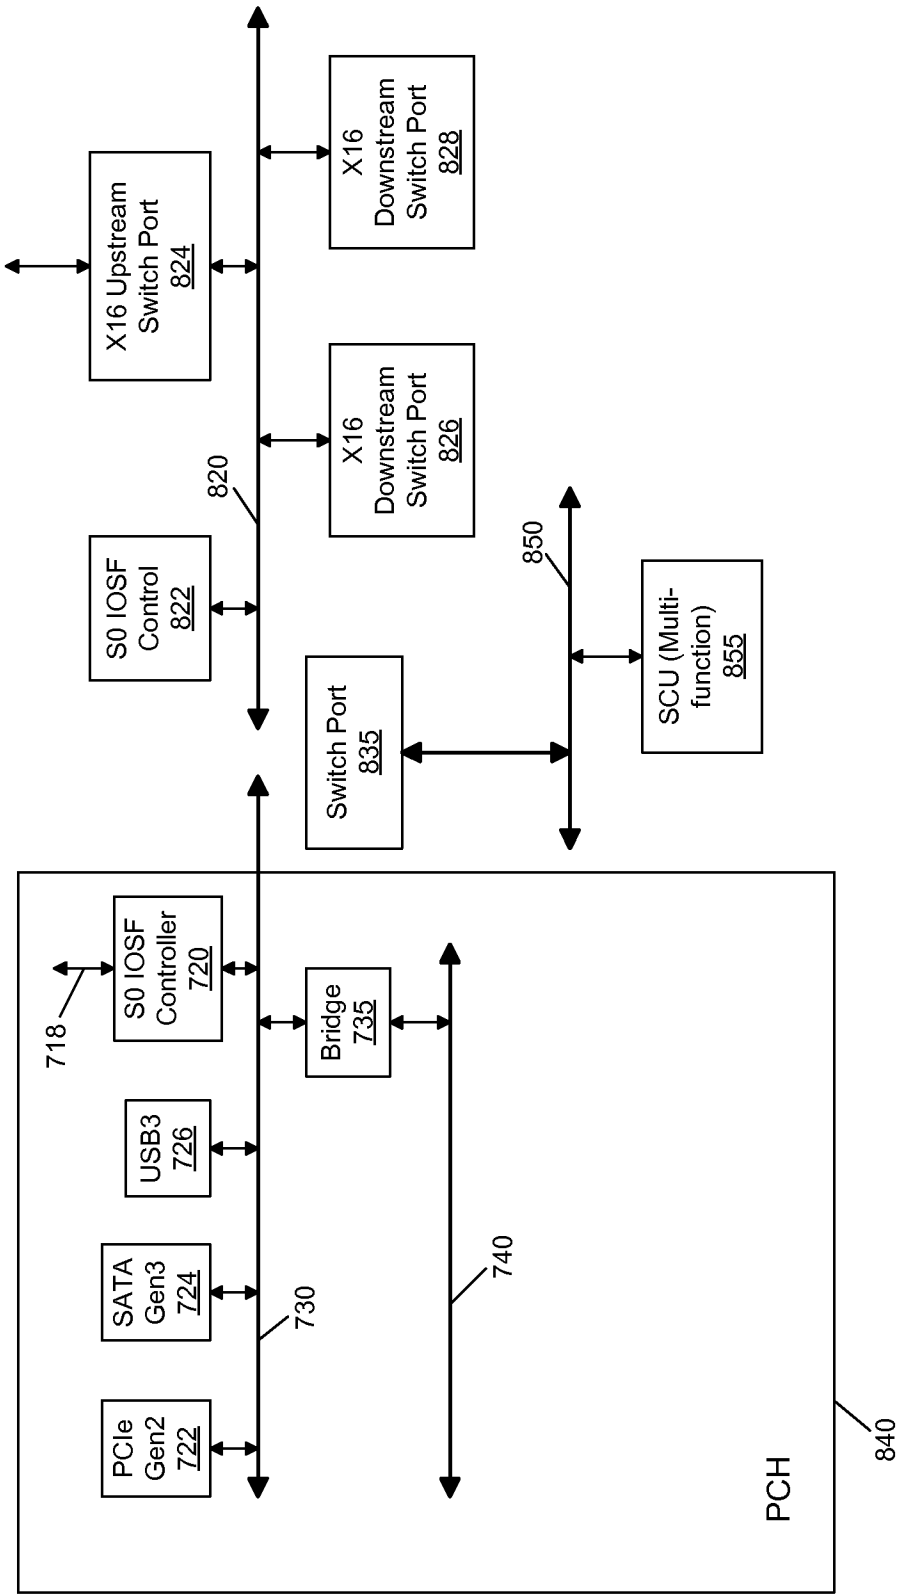


FIG. 8