



(51) International Patent Classification:

G06F 13/10 (2006.01) G06F 13/14 (2006.01)
G06F 13/12 (2006.01)

(21) International Application Number:

PCT/US2019/024586

(22) International Filing Date:

28 March 2019 (28.03.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/650,589 30 March 2018 (30.03.2018) US
62/691,117 28 June 2018 (28.06.2018) US

(71) Applicant: **PROVINO TECHNOLOGIES, INC.**

[US/US]; 2811 Mission College Blvd., 7th Floor, Santa Clara, California 95054 (US).

(72) Inventors: **DESAI, Shailendra**; 52 Shalin Residency,

Near SP Ring Road Crossings, Ahmedabad 380058 (IN).

PEARCE, Mark; 372 28th Street, San Francisco, Califor-

nia 94131 (US). **JAIN, Amit**; A105, Shree Rang Heights, Near Vijay Park Society, Gor ka Kuva, Maninagar East, Ahmedabad 380008 (IN). **BHATT, Rutul**; 30 Foram Park Society, R.C. Technical College Road, Ghatlodiya, Ahmedabad 380061 (IN). **PATEL, Jaymin**; A/89, Arti Society, Ajay Vibhag-5, Mahadevnagar, Vastral Road, Ahmedabad 382418 (IN).

(74) Agent: **ROSE, James**; BEYER LAW GROUP LLP, P.O.

Box 51887, Palo Alto, California 94303-1887 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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.

(54) Title: PROTOCOL LEVEL CONTROL FOR SYSTEM ON A CHIP (SoC) AGENT RESET AND POWER MANAGEMENT

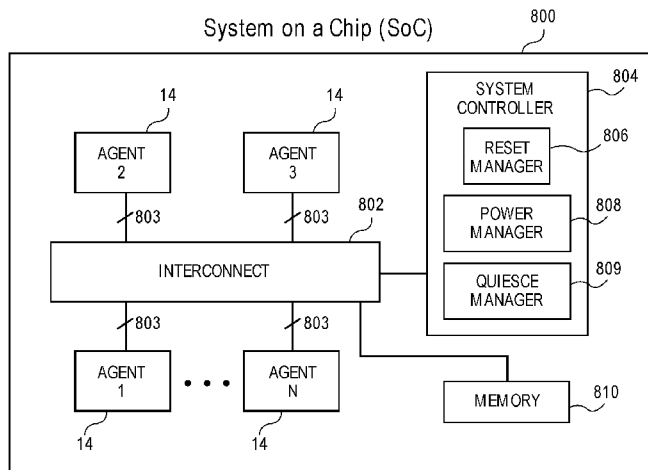


FIG. 8

(57) Abstract: A system for consistently implementing reset and power management of IP agents on a System on a Chip (SoC). When IP agents undergo a reset, an individual negotiation takes place between an interconnect and each IP agent over a link. Each IP agent can emerge from reset at its own time schedule, independently of the timing of the other IP agents. The interconnect may be configured as a proxy for any IP agent that is inoperable, including prior to reset, when in a power-down mode, or malfunctioning.



(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).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

PROTOCOL LEVEL CONTROL FOR SYSTEM ON A CHIP (SoC)
AGENT RESET AND POWER MANAGEMENT
CROSS-REFERENCE TO RELATED APPLCATIONS

5 [0001] The present application claim priority of U.S. Provisional Patent Application No. 62/650,589 (PRTIP001P) filed March 30, 2018 and U. S. Provisional Application 62/691,117 (PRT1P002P) filed June 28, 2018. Each of the above-listed priority applications are incorporated herein by reference in their entirety for all purposes.

Field of the Invention

10 [0002] The present application is directed to a System on a Chip (SoC), and more particularly, to a system and method for consistently implementing reset and/or power management functionality on SoC, which in turn, provides a more uniform system software view of SoCs, particularly across families of SoCs.

Description of Related Art

15 [0003] A System on a Chip ("SoC") is an integrated circuit that includes multiple sub-systems, often referred to as Intellectual Property ("IP") agents. IP agents are typically "reusable" blocks of circuitry designed to implement or perform a specific function. By using IP agents, the time and cost of developing complex SoCs can be significantly reduced.

20 [0004] SoCs typically include a system controller and an interconnect, such as a bus or Network on a Chip (NoC). The system controller runs system software and is provided to manage the overall operation of the SoC. The various IP agents are connected to the interconnect via one or more links and communicate with one another via the interconnect.

25 [0005] SoC developers commonly use disparate IP agents, often from multiple vendors. Each IP agent will ordinarily implement its own unique procedures for reset. From the perspective of the system controller and the interconnect on the SoC, this is problematic for several reasons.

30 [0006] A typical SoC will normally have multiple IP agents connected to the interconnect. Upon reset, each of the IP agents will likely emerge from the reset state at different times due to the unique reset procedures each uses. The different times each IP agent emerges from reset can cause significant problems. If a source IP agent

generates a transaction for a destination IP agent that is still in reset, then the (1) destination IP agent is unable to process the request and (2) source IP agent never receives a reply. As a result, the entire system may get hung up, possibly requiring a system-wide reset.

5 **[0007]** One known approach to prevent hang-ups is to design and place circuitry intermediate each link, between the interconnect and each IP agent. The purpose of this circuitry is to make sure that all the IP agents connected to the interconnect emerge from reset during the same clock cycle. This approach, however, has drawbacks for several reasons:

10 1. The design of the intermediate circuitry requires time and effort that will often delay the development of the SoC.

2. The intermediate circuitry, from one SoC to another, is typically developed by different design teams. As a result, the intermediate circuitry is usually different from one SoC to the next, or even between different sub-systems on the
15 same SoC.

3. The complexity of the circuitry normally means the number of IP agents that can be connected to a given interconnect is limited. The practical effect of this restriction is that more interconnect levels are needed to accommodate a given number of IP agents. The overall complexity of the SoC is therefore increased.

20 **[0008]** On occasion, IP agents malfunction. For example, IP agents may inject spurious transactions onto the interconnect, may fail to respond to a received transaction, generate an exception message, etc. In certain situations, the malfunctioning IP agent may need to be reset. With current SoC interconnect standards, there is no standardized IP agent reset mechanism. Either the entire SoC
25 has to be reset, or intermediary circuitry needs to be designed to perform the necessary isolation, reset, and re-introduction of the IP to the system, etc.

[0009] Power management is also not addressed with certain current SoC interconnect standards. The Advanced Microcontroller Bus Architecture (AMBA) protocol, for instance, does not address power management, and provides no method
30 for intentionally powering down or turning off IP agents. To provide this capability, power management functionality typically needs to be custom designed into the SoC

on a chip-by-chip basis, by developing for instance, additional intermediate circuitry on the links for handling power management.

5 [0010] Many companies offering multiple SoCs will share certain amounts of the system software among similar devices to reduce the time to market. However, even with SoCs that are similar, the software typically cannot simply be ported from one device to another, even in situations where the IP agents may be the same. If there are minor differences in any intermediate circuitry used for reset and/or power management, the system software may need to be modified and debugged for each device.

10 [0011] Companies that develop a large number of SoCs are thus challenged with (1) developing customized circuitry for implementing reset and possibly power management for each device and (2) modifying and debugging the system software for each device. This effort, across multiple devices is expensive, complex and time consuming, reducing the ability to quickly bring products to market.

15 [0012] A system for consistently implementing reset and power management of IP agents on SoCs, removing the need for customization, and which leads to a consistent system software view among multiple SoCs, is therefore needed.

SUMMARY

20 [0013] A system for consistently implementing reset and power management of IP agents on SoCs, removing the need for customization, and which leads to a consistent system software view among multiple SoCs, is disclosed.

[0014] In one embodiment, the system includes one or more IP agents, an interconnect and one or more links between the IP agents and the interconnect respectively. When the IP agents undergo a reset, an individual negotiation takes place between the interconnect and each IP agent over the link. With the individual negotiations, each IP agent can emerge from reset at its own time schedule, independently of the timing of the other IP agents. Upon emerging from reset, each IP agent becomes "transaction ready" and is introduced to the interconnect, becoming visible to other elements connected to the interconnect, such as the system controller.

30 [0015] In another embodiment, the interconnect may be configured as a proxy for any IP agent that is inoperable. This feature is beneficial because it prevents system wide hang ups that may otherwise occur when an IP agent is targeted with a

transaction (1) prior to being transaction ready, (2) malfunctioning and/or (3) inoperable when in a powered down state. With the interconnect acting as a proxy, an exception message can be sent to the source sending the transaction, preventing a hang up caused by the source waiting indefinitely for a response from the target IP agent.

[0016] In yet other embodiments, the ability to arrange for the interconnect to act as a proxy for an IP agent enables (1) IP agents to be individually reset and (2) IP agents to be placed in a power saving state. In various embodiments, the power saving state can include one of several modes, including a low power, operational mode, a low power inoperable mode with state information either maintained or not retained, or a power off mode.

[0017] The present invention thus solves a number of issues. It eliminates the need to create custom circuitry for (1) managing each IP agent to emerge from reset during the same time/clock cycle and (2) power management of IP agents. Instead, the present invention advantageously provides a uniform implementation for both these functions, which leads to a consistent system software view among multiple SoCs. With this consistent software view, much of the custom design and software modifications across families of SoCs is eliminated, saving development costs, reducing complexity, and providing a quicker time to market.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] The present application and the advantages thereof, may best be understood by reference to the following description taken in conjunction with the accompanying drawings in which:

[0019] FIG. 1 is a block diagram of a shared interconnect for a System on a Chip (SoC) in accordance with a non-exclusive embodiment.

[0020] FIG. 2 is an exemplary packet of a transaction in accordance with a non-exclusive embodiment.

[0021] FIG. 3A is a logic diagram illustrating an arbitration element in accordance with a first non-exclusive embodiments.

[0022] Fig. 3B is a logic diagram illustrating an arbitration element in accordance with a second non-exclusive embodiment.

- [0023] FIG. 4 is a flow diagram illustrating operational steps for arbitrating and sending portion(s) of transactions over virtual channels of the shared interconnect in accordance with a non-exclusive embodiment.
- [0024] FIG. 5 illustrates a first example of the interleaving the transmission of portions of different transactions over virtual channels of the shared interconnect in accordance with a non-exclusive embodiment.
- [0025] FIG. 6 illustrates a second example of the interleaving the transmission of portions of different transactions over virtual channels of the shared interconnect in accordance with a non-exclusive embodiment.
- [0026] FIG. 7 illustrates is a block diagram of two shared interconnects for handling traffic in two directions in accordance with another non-exclusive embodiment of the invention.
- [0027] FIG. 8 illustrates a block diagram of an SoC having reset, power management and quiesce functionality in accordance with a non-exclusive embodiment of the invention.
- [0028] FIG. 9 is a flow diagram showing an IP agent reset sequence in accordance with a non-exclusive embodiment of the invention.
- [0029] FIG. 10 is a flow diagram showing a reset sequence for a malfunctioning IP agent in accordance with a non-exclusive embodiment of the invention.
- [0030] FIG. 11 is a flow diagram illustrating a power down/up sequence for an IP agent in accordance with a non-exclusive embodiment of the invention.
- [0031] FIG. 12 is a flow diagram illustrating a power down/up sequence for an IP agent in accordance with a non-exclusive embodiment of the invention.
- [0032] FIG. 13 is a flow diagram illustrating a power down/up sequence for an IP agent in accordance with a yet another non-exclusive embodiment of the invention.
- [0033] FIG. 14 is a flow chart illustrating the steps for placing a link in a quiescent state.
- [0034] FIG. 15A-15D are flow diagrams illustrating various "wake-up" sequences for an IP agent.
- [0035] In the drawings, like reference numerals are sometimes used to designate like structural elements.

DETAILED DESCRIPTION

[0036] The present application will now be described in detail with reference to a few non-exclusive embodiments thereof as illustrated in the accompanying drawings. In the following description, numerous specific details are set forth in order to provide
5 a thorough understanding of the present disclosure. It will be apparent, however, to one skilled in the art, that the present discloser may be practiced without some or all of these specific details. In other instances, well known process steps and/or structures have not been described in detail in order to not unnecessarily obscure the present disclosure.

10 **[0037]** Many of the integrated circuits under development today are extremely complex. As a result, many chip designers have resorted to the System on a Chip or "SoC" approach, interconnecting a multiple sub-systems or IP agents on a single piece of silicon. SoCs are now available or are being developed for wide variety of applications, such as consumer devices (e.g., handheld, mobile phones, tablet
15 computers, laptop and desktop computers, media processing etc.), virtual or augmented reality (e.g., robotics, autonomous vehicles, aviation, etc.), medical instrumentation (e.g., imaging, etc.), industrial, home automation, industrial (e.g., smart appliances, home surveillance, etc.) and data center applications (e.g., network switches, attached storage devices, etc.).

20 **[0038]** The present application is broadly directed to an arbitration system and method for arbitrating access to a shared resource. Such a shared resource can be, for example, a bus interconnect, a memory resource, a processing resource, or just about any other resource that is shared among multiple vying parties. For the sake of illustration, the shared resources as described in detail below is an interconnect that is
25 shared by a plurality of sub-systems on a System on a Chip or "SoC".

[0039] With an SoC, as described in detail below, there are a plurality of sub-systems that exchange traffic with one another in the form of transactions, the shared resource is a physical interconnect, various transactions, or portions thereof, are transmitted over a multiplicity of virtual channels associated with the shared
30 interconnect and one of a number of different arbitration schemes and/or priorities may be used to arbitrate access to the shared interconnect for the transmission of transactions between the sub-functions.

Transaction Classes

[0040] Within the above-mentioned shared interconnect used for SoCs, there are at least three types or classes of transactions, including Posted (P), Non-posted (NP) and Completion (C). A brief definition of each is provided in Table I below.

5

Table I

Transaction Class	Description
Posted (P)	A transaction that requires no response (e.g., a write operation)
Non-posted (NP)	A transaction that requires a response transaction from the destination agent (e.g., a read or a write operation)
Completion (C)	A response transaction for a non-posted transaction.

[0041] A Posted transaction, such as a write, requires no response transaction. Once a source writes data to a designated destination, the transaction is finished. With a Non-posted transaction, such as either a read or a write, a response is required. However, the response is bifurcated as a separate Completion transaction. In other words with a read, a first transaction is used for the read operation, while a separate, but related, Completion transaction is used for returning the read contents. With a Non-posted write, a first transaction is used for the write, while a second related Completion transaction is required for the confirmation once the write is complete.

15

[0042] Transactions, regardless of the type, can be represented by one or more packets. In some circumstances, a transaction may be represented by a single packet. In other circumstances, multiple packets may be needed to represent the entire transaction.

20

[0043] A beat is the amount of data that can be transmitted over the shared interconnect per clock cycle. For example if the shared interconnect is physically 128 bits wide, then 128 bits can be transmitted each beat or clock cycle.

25

[0044] In some circumstances, a transaction may need to be divided into multiple portions for transmission. Consider a transaction having a single packet that has a payload that is 512 bits (64 bytes). If the shared interconnect is only 128 bits wide (16 bytes), then the transaction needs to be segmented into four portions (e.g. 4 x 128 = 512) and transmitted over four clock cycles or beats. On the other hand if a

transaction is only a single packet that is 128 bits wide or less, then the entire transaction can be sent in one clock cycle or beat. If the same transaction happens to include additional packets, then additional clock cycles or beats may be needed.

[0045] The term "portion" of a transaction is therefore the amount of data that can
5 be transferred over the shared interconnect during a given clock cycle or beat. The size of a portion may vary depending on the physical width of the shared interconnect. For instance, if the shared interconnect is physically 64 data bits wide, then the maximum number of bits that can be transferred during any one cycle or beat is 64 bits. If a given transaction has a payload of 64 bits or less, then the entire transaction
10 can be sent over the shared interconnect in a single portion. On the other hand if the payload is larger, then the packet has to be sent over the shared interconnect in multiple portions. A transaction with a payload of 128, 256 or 512 bits requires two (2), four (4) and eight (8) portions respectively. As such, the term "portion" or "portions" should therefore be broadly construed to mean either part of or an entire
15 transaction that may be sent over the share interconnect during any given clock cycle or beat.

Streams

[0046] A stream is defined as the pairing of a virtual channel and a transaction class. For instance, if there are four (4) virtual channels (e.g., VC0, VC1, VC2 and
20 VC3) and three (3) transaction classes (P, NP, C), then there are a maximum of twelve (12) different possible streams. The various combinations of virtual channels and transaction classes are detailed below in Table II.

Table II

Stream Number	VC/TC Combination
1	VC0/NP
2	VC0/P
3	VC0/C
4	VC1/NP
5	VC1/P
6	VC1/C
7	VC2/NP
8	VC2/P
9	VC2/C
10	VC3/NP
11	VC3/P
12	VC3/C

[0047] It should be noted that the number of transaction classes discussed above is merely exemplary and should not be construed as limiting. On the contrary, any number of virtual channels and/or transaction classes may be used.

Arbitration over Virtual Channels of a Shared Interconnect

[0048] Referring to Fig. 1, a block diagram of an arbitration system 10 is shown. In a non-exclusive embodiment, the arbitration system is used for arbitrating access by a number of sub-functions 14 (i.e., IP₁, IP₂ and IP₃) to a shared interconnect 12 attempting to send transactions to upstream sub-functions 14 (i.e., IP₄, IP₅ and IP₆).

[0049] The shared interconnect 12 is a physical interconnect that is N data bits wide and includes M control bits. The shared interconnect 12 is also one-directional, meaning it handles traffic only from a source (i.e., IP₁, IP₂ and IP₃) to a destination (i.e., IP₄, IP₅ and IP₆).

[0050] In various alternatives, the number of N data bits can be any integer, but typically is some power of the number 2 (e.g., 2¹, 2², 2³, 2⁴, 2⁵, 2⁶, 2⁷, 2⁸, 2⁹ etc.) or (2, 4, 6, 8, 16, 32, 64, 128, 256 etc.) bits wide respectively. With most real-world applications, the number of N bits is either 32, 64, 128, 256 or even 512. However, it

should be understood that these widths are merely illustrative and should not be construed as limiting in any manner.

[0051] The number of control bits M may also vary and be any number.

[0052] One or more logical channels (not illustrated), hereafter referred to as "virtual channels" or "VCs" are associated with the shared interconnect 12. Each virtual channel is independent. Each virtual channel may be associated with multiple independent streams. The number of virtual channels may widely vary. For example, up to thirty-two (32) or more virtual channels may be defined or associated with the shared interconnect 12.

[0053] In various alternative embodiments, each virtual channel may be assigned a different priority. One or more virtual channel(s) may be assigned a higher priority, while one or more other virtual channel(s) may be assigned a lower priority. The higher priority channels are awarded or arbitrated access to the shared interconnect 12 over the lower priority virtual channels. With other embodiments, each of the virtual channels may be given the same priority, in which case, no preference is given to one virtual channel versus another when awarding or arbitrating access to shared interconnect 12. In yet other embodiments, the priority assigned to one or more of the virtual channels may also dynamically change. For instance, in a first set of circumstances, all the virtual channels may be assigned the same priority, but in a second set of circumstances, certain virtual channel(s) can be assigned a higher priority than other virtual channel(s). Thus as circumstances change, the priority scheme used among the virtual channels can be varied to best meet current operating conditions.

[0054] Each of the sub-systems 14 is typically a block of "reusable" circuitry or logic, commonly referred to as an IP core or agent. Most IP agents are designed to perform a specific function, for example, controllers for peripheral devices such as an Ethernet port, a display driver, an SDRAM interface, a USB port, etc. Such IP agents are generally used as "building blocks" that provide needed sub-system functionality within the overall design of a complex system provided on an integrated circuit (IC), such as either an Application Specific Integrated Circuit (ASIC) or a Field Programmable Gate Array (FPGA). By using a library of available IP agents, a chip designer can readily "bolt" together various logic functions in the design of a more

complex integrated circuit, reducing design time and saving development costs. Although sub-system agents 14 are described above in terms of a dedicated IP core, it should be understood that this is not a necessary requirement. On the contrary, a sub-system 14 can also be a collection of IP functions connected to or sharing a single port
5 20. Accordingly, the term "agent" should be broadly construed as any type of sub-system connected to a port 20, regardless if the sub-system performs a single function or multiple functions.

[0055] A pair of switches 16 and 18 provides access between each of the sub-system agents 14 and the shared interconnect 12 via dedicated access ports 20
10 respectively. With the exemplary embodiment shown:

(1) The sub-system agents IP₁, IP₂, and IP₃ connect with switch 16 via access Port 0, Port 1 and Port 2 respectively.

(2) The sub-system agents IP₄, IP₅, and IP₆ connect with switch 18 via Port 3, Port 4 and Port 5 respectively.

15 (3) In addition, an access port 22, via the interconnect 12, provides sub-system agents IP₄, IP₅, and IP₆ access to switch 16 in the aggregate.

[0056] The switches 16 and 18 perform multiplexing and de-multiplexing functions. Switch 16 selects up-stream traffic generated by the sub-system agents IP₁, IP₂, and/or IP₃ and sends the traffic downstream over the shared interconnect 12. At
20 the switch 18, a de-multiplexing operation is performed and the traffic is provided to a target sub-system agent (i.e., either IP₄, IP₅, or IP₆).

[0057] Each access port 20 has a unique port identifier (ID) and provides each sub-system agent 14 dedicated access to either switch 16 or 18. For instance, sub-system agents IP₁, IP₂ and IP₃ are assigned to access ports Port 0, Port 1 and Port 2
25 respectively. Similarly, the sub-system agents IP₄, IP₅ and IP₆ are assigned access ports Port 3, Port 4 and Port 5 respectively.

[0058] Besides providing ingress and egress points to/from the switches 16, 18, the unique port IDs 20 are used for addressing traffic between the sub-system agents 14. Each Port 20 has a certain amount of allocated addressable space in system
30 memory 24.

[0059] In certain non-exclusive embodiments, all or some of the access ports 20 can also be assigned a "global" port identifier as well their unique port ID.

Transactions and other traffic can be sent to all or some of the access ports assigned to the global port identifier. Accordingly, with the global identifier, transactions and other traffic can be widely disseminated or broadcast to all or some of the access ports 20, obviating the need to individually address each access port 20 using its unique
5 identifier.

[0060] The switch 16 also includes an arbitration element 26, Address Resolution Logic (ARL) 28 and an address resolution Look Up Table (LUT) 30.

[0061] During operation, the sub-system agents IP_1 , IP_2 and IP_3 generate transactions. As each transaction is generated, it is packetized by the originating sub-
10 system agent 14 and then the packetized transaction is injected via the corresponding port 20 into a local switch 16. For instance, portions of transactions generated by IP_1 , IP_2 and IP_3 are provided to switch 16 by via ports Port 0, Port 1 and Port 2 respectively.

[0062] The ports 20 each include a number of first-in, first-out buffers (not
15 illustrated) for each of the virtual channels associated with the interconnect channel 12 respectively. In a non-exclusive embodiment, there are four (4) virtual channels. In which case, each port 20 includes four buffers, one for each virtual channel. Again, it should be understood that the number of virtual channels and buffers contained in the ports 20 may vary and is not limited to four. On the contrary, the number of virtual
20 channels and buffers may be more or less than four.

[0063] If a given transaction is represented by two (or more) portions, those portions are maintained in the same buffer. For instance, if interconnect 12 is 128 data bits wide and a transaction is represented by a packet containing a payload of 512 bits, then the transaction needs to be segmented into four (4) portions that are transmitted
25 over four clock cycles or beats. On the other hand if the transaction can be represented by a single packet having a payload of 64 bits, then the single portion can be transmitted in one clock cycle or beat. By maintaining all the portion(s) of given transaction in the same buffer, the virtual channels remain logically independent. In other words, all the traffic associated with a given transaction is always sent over the
30 same virtual channel as a stream and is not bifurcated over multiple virtual channels.

[0064] The arbitration element 26 is responsible for arbitrating among the competing buffered portions of transactions maintained by the various access ports

20. In a non-exclusive embodiment, the arbitration element 26 performs an arbitration every clock cycle, provided multiple competing transactions are available. The arbitration winner per cycle yields a portion of a transaction, from one of the sub-systems IP₁, IP₂ and IP₃, that is granted access to and is transmitted over the interconnect 12.

[0065] When generating transactions, the source sub-system IP₁, IP₂ and IP₃ ordinarily knows the address in the address space for the possible destination sub-system agents IP₄, IP₅ and IP₆, but does not know the information (e.g., the Port IDs 20 and/or 22) needed to route the transactions to their destinations. In one embodiment, the local Address Resolution Logic (ARL) 28 is used for resolving the known destination address into the needed routing information. In other words, a source sub-agent 14 may simply know that it wishes to access a given address in system memory 24. The ARL 28 is therefore tasked to access the LUT 30 and performs an address look up of the port(s) 20/22 along the delivery path to the final destination corresponding to the specified address. Once the ports 20/22 is/are known, this information is inserted in a destination field in the packet(s) of the transaction. As a result, the packet(s) is/are delivered to the ports 20/22 along the delivery path. As a general rule, downstream nodes along the delivery path do not have to perform additional look up(s) since the required delivery information is already known and included in the destination field of the packet(s). With other types of transactions, referred to as Source Based Routing (SBR) as described in more detail below, the source IP agent knows the destination port address. As a result, the lookup performed by the ARL 28 typically does not need to be performed.

[0066] In an alternative embodiment, not all the nodes within the interconnect require an ARL 28 and LUT 30. For nodes that do not have these elements, transactions without needed routing information can be forwarded to a default node. At the default node, an ARL 28 and LUT 30 are accessed and the needed routing information can then be inserted into the headers of the packet(s) of transactions. The default node is typically upstream from the node without the ARL 28 and LUT 30. However, this is by no means a requirement. The default node, or nodes, can be located anywhere on the SoC. By eliminating ARLs 28 and LUTs 30 from certain nodes, their complexity can be reduced.

[0067] The ARL 28 may also be referred to as an "ordering point" because, besides decoding the forwarding destination for winning portion(s) of transactions, it defines a sequence order for the winning portion(s) of transactions within each virtual channel. As each arbitration is resolved, regardless of whether or not the ARL 28 is used to perform an address port lookup, the winning portions of transactions are inserted into a first-in, first out queue provided for each virtual channel. The winning portions of transactions then await their turn for transmission over the interconnect 12 in the buffer.

[0068] The ARL 28 is also used for defining "upstream" and downstream" traffic. In other words any transactions generated by the IP agents 14 associated with switch 16 (i.e., IP₁, IP₂ and IP₃) is considered upstream with respect to the ARL 28. All transaction post the ARL 28 (i.e., transmitted to IP₄, IP₅ and IP₆) is considered downstream traffic.

[0069] The IP agents 14 associated with switch 16 (i.e., IP₁, IP₂ and IP₃) may communicate and send transactions to one another, either directly or indirectly. With direct communication, often referred to as Source Based Routing (SBR), the IP agents 14 can send transactions to one another in a peer-to-peer model. With this model, the source IP agent knows the unique Port ID of its peer IP agents 14, bypassing the need to use the ARL 28 to access the LUT 30. Alternatively, the transactions between the IP agents associated with the switch 16 can be routed using the ARL 28. With this model, similar to that described above, the source IP agent only knows the address of the destination IP agent 14, but not the information needed for routing. The ARL 28 is then used to access the LUT 30, find the corresponding Port ID, which is then inserted into the destination field of the packet(s) of the transaction.

Packet Format

[0070] The IP agents 14 create and process transactions over virtual channels associated with the interconnect 12. Each transaction typically is made up of one or more packets. Each Packet typically has a fixed header size and format. In some instances, each packet may have a fixed sized payload. In other instances, packet payloads may vary in size, from large too small, or even with no payload at all.

[0071] Referring to Fig. 2, an exemplary packet 32 is shown. The packet 32 includes a header 34 and a payload 36. In this particular embodiment, the header 34 is

sixteen (16) Bytes in size. It should be understood that this size is exemplary and either a larger size (e.g., more Bytes) or smaller size (e.g., fewer Bytes) packets may be used. It should also be understood that headers 34 of packets 32 do not necessarily have to all be the same size. In alternative embodiments, the size of packet headers in a SoC may be variable.

[0072] The header 34 includes a number of fields including a destination identifier (DST_ID), a source identifier (SRC_ID), a payload size indicator (PLD_SZ), a reserved field (RSVD), a command field (CMD), a TAG field, a status (STS), a transaction ID field (TAG), an address or ADDR field, a USDR/Compact payload field, a transaction Class or TC field, a format FMT field, and a Byte Enable (BE) field. The various fields of the header 34 are briefly described in Table III below.

Table III

Name of Field	Description
DST	Specifies the corresponding Port ID for the targeted IP agent 14.
SRC	Specifies the Port ID for the sending IP agent 14.
PLD_SZ	Specifies the size of the payload of the packet.
CMD	Specifies the type of transaction/command the packet contains. Exemplary commands may include Incremental Read, Incremental Write, Compact Read, Compact Write, Write to FIFO, Destructive Incremental Read, Wrap, etc.
TAG	Specifies a transaction ID for the packet. Transactions IDs are used for matching Non-posted transactions and their corresponding Completion transactions. When a Completion transaction including a matching transaction ID in the TAG field is received, the corresponding Non-posted read or write operation is complete.
ADDR	Specifies the physical address in system memory 24 of the request
USRD/Compact payload	If the payload of the packet is sufficiently small, it can be transported in this field in the header, not the payload. This field can also be used to transport a private or secure message between the source and the destination IP ports.
STS	This field is valid only with Completion packets. Specifies the status of the corresponding Non-posted transaction, (i.e., either a successful completion or a failed transaction).

RSVD	This is a reserved field that can also be used to transport a private or secure message between the source and the destination IP ports.
FMT	Specifies the format specification if multiple header formats are defined and used.
BE	Byte Enable, indicates which bytes in the payload are valid.

[0073] The payload 36 contains the contents of the packet. The size of the payload may vary. In some instances, the payload may be large. In other instances, it could be small. In yet other instances, if the content is very small or "compact", it can be transported in the USRD field of the header 34.

[0074] The type of transaction will often dictate whether or not the packet(s) used to represent the transaction has/have payloads or not. For example with either a Posted or Non-posted read, the packet(s) will designate the location address to be accessed, but will typically have no payload. The packets for the related Completion transaction, however, will include payload(s) containing the read content. With both Posted and Non-posted write transactions, the packet(s) will include a payload containing the data to be written to the destination. With Non-posted versions of a write, the packets for the Completion transaction will ordinarily not defined a payload. However, in some situations, a Completion transaction will define a payload.

[0075] The exemplary packet and above description covers many of basic fields that may be included in a packet. It should be understood that additional fields may be deleted or added. For instance, a private signaling field may be used so a source and a destination may share private messages.

Arbitration

[0076] Referring to Fig. 3A, a logic diagram illustrating the arbitration logic performed by the arbitration element 26 with Peripheral Component Interconnect (PCI) ordering is shown.

[0077] With PCI ordering, each Port 20 includes separate buffers for each virtual channel and transaction class (P, NP and C) combination. For instance, with four virtual channels (VC0, VC1, VC2 and VC3), the Ports 0, Port 1 and Port 2 each have twelve first-in, first-out buffers. In other words for each Port 20, a buffer is provided for each transaction class (P, NP, and C) and virtual channel (VC0, VC1, VC2 and VC30 combination.

[0078] As each IP agent 14 (e.g., IP₁, IP₂ and IP₃) generates transactions, the resulting packets are placed in the appropriate buffer, based on transaction type, in the corresponding port (e.g., Port 0, Port 1 and Port 2) respectively. For instance, Posted (P), Non-posted (NP) and Completion (C) transactions generated by IP₁ are each placed in the Posted, Non-posted and Completion buffers for the assigned virtual channel in Port 0 respectively. Transactions generated by IP₂ and IP₃ are similarly placed in the Posted, Non-posted and Completion buffers for the assigned virtual channels in Ports 1 and Port 2 in a similar manner.

[0079] If a given transaction is represented by multiple packets, all of the packets of that transaction are inserted into the same buffer. As a result, all of the packets of the transaction are eventually transmitted over the same virtual channel. With this policy, the virtual channels remain independent, meaning different virtual channels are not used for transmission of multiple packets associated with the same transaction.

[0080] Within each port 20, packets can be assigned to a given virtual channel in a number of different ways. For instance, the assignment can be arbitrary. Alternatively, the assignment can be based on workload and the amount of outstanding traffic for each of the virtual channels. If one channel is very busy and the other not, then the port 20 will often attempt to balance the load and assign newly generated transaction traffic to under-utilized virtual channels. As a result, routing efficiency is improved. In yet other alternatives, transaction traffic can be assigned to a particular virtual channel based on urgency, security, or even a combination of both. If a certain virtual channel is given a higher priority and/or security than others, then high priority and/or secure traffic is assigned to the higher priority virtual channel. In yet other embodiments, a port 20 can be hard-coded, meaning the port 20 has only one virtual channel and all traffic generated by that port 20 is transmitted over the one virtual channel. In yet other embodiments, the assignment can be based on the route chosen to reach the destination port 20

[0081] In yet other embodiments, the assignment of virtual channels can be implemented by the source IP agents 14, either alone or in cooperation with its corresponding port 20. For example, a source IP agent 14 can generate a control signal to the corresponding port 20 requesting that packet(s) of a given transaction be assigned to a particular virtual channel. IP agents 14 can also make assignment

decisions that are arbitrary, hard coded, based on balanced usage across all the virtual channels, security, urgency, etc., as discussed above.

[0082] In selecting an arbitration winner, the arbitration element 26 performs multiple arbitration steps per cycle. These arbitration steps include:

- 5 (1) Selecting a port;
- (2) Selecting a virtual channel; and
- (3) Selecting a transaction class.

[0083] The above order (1), (2) and (3) is not fixed. On the contrary, the above three steps may be completed in any order. Regardless of which order is used, a
10 single arbitration winner is selected each cycle. The winning transaction is then transmitted over the corresponding virtual channel associated with the interconnect 12.

[0084] For each arbitration (1), (2) and (3) performed by arbitration element 26, a number of arbitration schemes or rule sets may be used. Such arbitration schemes may
15 include strict or absolute priority, a weighed priority where each of the four virtual channels is assigned a certain percentage of transaction traffic or a round-robin scheme where transactions are assigned to virtual channels in a predefined sequence order. In additional embodiments, other priority scheme such may be used. Also, it should be understood that the arbitration element 26 may dynamically switch among
20 the different arbitration schemes from time-to-time and/or use the same or different arbitration schemes for each of the (1), (2) and (3) arbitrations respectively.

[0085] In an optional embodiment, availability of the destination ports 20 defined by the outstanding transaction(s) considered during a given arbitration cycle are considered. If a buffer in a destination port 20 does not have the resources available
25 to process a given transaction, then the corresponding virtual channel is not available. As a result, the transaction in question does not compete in the arbitration, but rather, waits until a subsequent arbitration cycle when the target resource becomes available. On the other hand, when target resource(s) is/are available, the corresponding transaction(s) are arbitrated and compete for access to the interconnect 12.

30 **[0086]** The availability of the destination ports 20 may be checked at different times with respect to the multiple arbitration steps (1), (2) and (3) noted above. For instance, the availability check can be performed prior to the arbitration cycle (i.e.,

prior to completion of any of steps (1), (2) and (3)). As a result, only transactions that define available destination resources is/are considered during the subsequent arbitration. Alternatively, the availability check can be performed intermediate any of the three arbitration steps (1), (2) and (3), regardless of the order in which they are
5 implemented.

[0087] There are advantages and disadvantages in performing the destination resource availability check early or late in the arbitration process. By performing the check early, possible competing portions of transactions can potentially be eliminated from the competition if their destinations are not available. However, early notice of
10 availability may create a significant amount of overhead on system resources. As a result, depending on circumstances, it may be more practical to perform the availability check later in a given arbitration cycle.

[0088] For the arbitration step involving the selection of a transaction class, a number of rules are defined to arbitrate among competing portions of N, NP and C
15 transactions. These rules include:

For Posted (P) Transactions:

- A Posted transaction portion may not pass another Posted transaction portion;
- A Posted transaction portion must be able to pass a Non-posted transaction
20 portion to avoid deadlock;
- A Posted transaction portion must be able to pass a Completion if both are in a strong order mode. In other words in the strong mode, the transaction need to be performed strictly in accordance with the rules and the rules cannot be relaxed; and
- A Posted request is permitted to pass a Completion, but passage is not
25 mandatory, if any transaction portion has its Relaxed Order (RO) bit set. With relaxed order, the rules are generally followed, however exceptions can be made.

For Non-posted (NP) Transactions:

- A Non-posted transaction portion must not pass a Posted transaction portion;
- A Non-posted transaction portion must not pass another Non-posted
30 transaction portion;
- A Non-posted transaction portion must not pass a Completion if both are in the strong order mode; and

- A Non-posted transaction portion is permitted to pass a Completion, but is not mandatory, if any transaction portion has its RO bit set.

For Completion (C) Transactions:

5 - A Completion must not pass a Posted transaction portion if both are in the strong order mode;

- A Completion is permitted to pass a Posted transaction portion, but is not mandatory, if any transaction portion has its RO bit set;

- A Completion must not pass a Non-posted transaction portion if both are in the strong order mode;

10 - A Completion is permitted to pass a Non-posted transaction portion, but is not mandatory, if any transaction portion has its RO bit set; and

- A Completion is not permitted to pass another Completion.

[0089] Table IV below provides a summary of the PCI ordering rules. In the boxes with no (a) and (b) options, then the strict ordering rules need to be followed. In the boxes of the Table having (a) and (b) options, either strict order (a) or relaxed order (b) rules may be applied, depending on if the RO bit is reset or set respectively. In various alternative embodiments, the RO bit can be set or reset either globally or on individually on the packet level.

Table IV

Row Pass Column?	Posted Request (Column 2)	Non-posted Request (Column 3)	Completion (Column 4)
Posted Request Row A	No	Yes	(a) Yes (b) Y/N
Non-posted Request Row B	No	No	(a) No (b) Y/N
Completion Row C	(a) No (b) Y/N	(a) Yes (b) Y/N	No

20

[0090] The arbitration element 26 selects an ultimate winning transaction portion by performing, in no particular order, arbitrations among competing Ports 20, virtual channels and transactions classes respectively. The winning portion per cycle gains access to the shared interconnect 12 and is transmitted over the corresponding virtual channel.

25

[0091] Referring to Fig. 3B, a logic diagram illustrating the arbitration logic performed by the arbitration element 26 with Device ordering is shown. The arbitration process, and possibly the consideration of available destination resources, is essentially the same as described above, except for two distinctions.

5 [0092] First, with Device ordering, there are only two transaction classes defined, including (a) Non-posted read or write transactions where a response for every request is required and (b) Completion transactions, which defined the required responses. Since there are only two transaction classes, there are only two (2) buffers per virtual channel in each Port 20. For instance, with four (4) virtual channels (VCO, VC1, VC2
10 and VC3), each Port 20 (e.g., Port 0, Port 1 and Port 2) has a total of eight (8) buffers.

[0093] Second, the rules for selecting a Transaction for Device ordering are also different than PCI ordering. With Device ordering, there are no strict rules governing the selection of one class over the other class. On the contrary, either transaction class can be arbitrarily selected. However, common practice typically calls for favoring
15 Completion transactions to free up resources that may not be available until a Completion transaction is resolved.

[0094] Otherwise, the arbitration process for Device order is essentially the same as described above. In other words for each arbitration cycle, the arbitration steps (1), (2) and (3) are performed, in any particular order, to select an arbitration winner.
20 When the transaction class arbitration is performed, Device order rather than PCI order rules are used. In addition, the availability of destination resources and/or virtual channels may also be considered either prior to or intermediate any of the arbitration steps (1), (2) and (3).

Operational Flow Diagram

25 [0095] As previously noted, the above-described arbitration scheme can be used for sharing access to any shared resource and is not limited to use with just a shared interconnect. Such other shared resources may include the ARL 28, a processing resource, a memory resource such as the LUT 30, or just about any other type of resource that is shared among multiple parties vying for access.

30 [0096] Referring to FIG. 4, a flow diagram 40 illustrating operational steps for arbitrating access to a shared resource is shown.

[0097] In step 42, the various source sub-system agents 14 generate transactions. The transactions can be any of the three classes, including Posted (P), Non-posted (NP) and Completion (C).

5 [0098] In step 44, each of the transactions generated by the source sub-system agents 14 are packetized. As previously noted, packetization of a given transaction may result in one or multiple packets. The packets may also vary in size, with some packets having large payloads and others having small or no payloads. In situations where a transaction is represented by a single packet having a data payload 36 that is smaller than the width of the interconnect 12, the transaction can be represented by a
10 single portion. In situations where a transaction is represented by multiple packets, or a single packet with a data payload 36 that is larger than the access width of the shared resource, then multiple portions are needed to represent the transaction.

[0099] In step 46, the portion(s) of the packetized transactions generated by each of the sub-system agents 14 are injected into the local switch 16 via its corresponding
15 port 20. Within the port 20, the packet(s) of each transaction are assigned to a virtual channel. As previously noted, the assignment can be arbitrary, hard coded, based on balanced usage across all the virtual channels, security, urgency, etc.

[0100] In step 48, the portion(s) of the packetized transactions generated by each of the sub-system agents 14 are stored in the appropriate, first-in, first-out, buffer by
20 both transaction class and by their assigned virtual channel (e.g., VC0, VC1, VC2 and VC3) respectively. As previously noted, virtual channels may be assigned by one of a number of different priority schemes, including strict or absolute priority, round-robin, weighted priority, least recently serviced, etc. If a given transaction has multiple portions, each portion will be stored in the same buffer. As a result, the
25 multiple portions of a given transaction are transmitted over the same virtual channel associated with the interconnect 12. As transaction portions are injected, the corresponding a counter for tracking the number content items in each buffer is decremented. If a particular buffer is filled, its counter is decremented to zero, meaning the buffer can no longer receive additional contents.

30 [0101] In steps 50, 52 and 54, first, second and third level arbitrations are performed. As previously noted, the selection of a Port 20, a virtual channel and a transaction class can be performed in any order.

[0102] Element 56 may be used to maintain the rules used to perform the first, second and third levels of arbitration. In each case, the element 56 is used as needed in resolving each of the arbitration levels. For instance, element 56 may maintain PCI and/or Device ordering rules. Element 56 may also contain rules for implementing
5 several priority schemes, such as strict or absolute priority, weighted priority, round robin, etc., and the logic or intelligence for deciding which to use in a given arbitration cycle.

[0103] In step 58, a winner of the arbitration is determined. In step 60, the winning portion is placed in a buffer used for accessing the shared resource and a
10 counter associated with the buffer is decremented.

[0104] In step 62, the buffer associated with the winning portion is incremented since the winning portion is no longer in the buffer.

[0105] In step 64, the winning portion gains access to the shared resource. Once the access is complete, the buffer for the shared resource is incremented

15 [0106] The steps 42 through 64 are continually repeated during successive clock cycles respectively. As different winning portions, each gains access to the shared resource.

Interleaving - Example One

[0107] Transactions can be transmitted over the interconnect 12 in one of several
20 modes.

[0108] In one mode, referred to as the "header in-line" mode the header 34 of packet(s) 32 of a transaction are always transmitted first ahead of the payload 36 in separate portions or beats respectively. The header in-line mode may or may not be wasteful of the bits available on the interconnect 12, depending the relative size of the
25 header 34 and/or the payload 36 with respect to the number of data bits N of the interconnect 12. For instance, consider an interconnect 12 that is 512 bits wide (N = 512) and a packet having a header that is 128 bits and a payload of 256 bits. With this scenario, the 128 bits of the header are transmitted in a first portion or beat, while the remaining 384 bits of bandwidth of the interconnect 12 are not used. In a second
30 portion or beat, the 256 bits of the payload 36 are transmitted, while the remaining 256 bits of the interconnect 12 are not used. In this example, a significant percentage of the bandwidth of the interconnect is not used during the two beats. On the other

hand if the majority of the packets of transactions are the same size or larger than the interconnect, than the degree of wasted bandwidth is reduced or possibly eliminated. For example with headers and/or payloads that are 384 or 512 bits, the amount of waste is either significantly reduced (e.g., with 384 bits) or eliminated altogether (e.g.,
5 with 512 bits).

[0109] In another mode, referred to as "header on side-band", the header 34 of a packet is transmitted "on the side" of the data, meaning using the control bits M, while the payload 36 is transmitted over the N data bits of the interconnect 12. With the header on side band mode, the number of bits or size of the payload 36 of a packet 32
10 determines the number of beats needed to transmit the packet over a given interconnect 12. For instance, with a packet 32 having a payload 36 of 64, 128, 256 or 512 bits and an interconnect 12 having 128 data bits ($N=128$), the packet requires 1, 1, 2 and 4 beats respectively. With the transmission of each of the beat(s), the header information is transmitted over the control bits M along with or "on the side"
15 of the data of the payload over the N data bits of the interconnect 12.

[0110] In yet another mode, the header 34 of packets 32 are transmitted in line with the payload, but there is no requirement that the header 34 and the payload 36 must be transmitted in separate portions or beats. If a packet 32 has a header 34 that is 128 bits and a payload 36 that is 128 bits, then the total size is 256 bits ($128 + 128$).
20 If the N data bits of interconnect 12 is 64, 128, 256 or 512 bits wide, then a packet of 256 bits is transmitted in 4, 2, 1 and 1 beats respectively. In another example, a packet 32 has a header of 128 bits and a payload 36 of 256 bits, or a total packet size of 384 bits ($128 + 256$). With the same interconnect 12 of N data bits of 64, 128, 256 or 512 wide, the packet is transmitted in 6, 3, 2, or 1 beats respectively. This mode will
25 always be as least as efficient or more efficient as the header in-line mode described above.

[0111] Referring to Fig. 5, a first example of the interleaving of portions of different transactions over multiple virtual channels is illustrated. In this example, for the sake of simplicity, only two transactions are defined. The two transactions are
30 competing for access to shared interconnect 12, which is 128 data bits wide ($N = 128$) in this example. The details of the two transactions include:

(1) Transaction 1 (T1), which was generated at Time T1 and which is assigned to virtual channel VC2. The size of T1 is four beats, designated as T1A, T1B, T1C and T1D; and

(2) Transaction 2 (T2), which was generated at Time T2 (after Time T1) is assigned to virtual channel VC0. The size of T2 is a single portion or beat.

[0112] In this example, VCO is assigned absolute or strict priority. Over the course of multiple cycles, the portions of the two transactions T1 and T2 are transmitted over the shared interconnect, as depicted in Fig. 5, as follows:

Cycle 1: Beat T1A of T1 is transmitted over VC2 because it is the only available transaction;

Cycle 2: Beat T1B of T1 and the single portion of T2 are competing for access to the interconnect 12. Since VCO has strict priority, T2 automatically wins. Accordingly, the beat of T2 is transmitted over VC0.

Cycle 3: Since there are no competing transactions, beat T1B of T1 is transmitted over VC2.

Cycle 4: Since there are no competing transactions, beat T1C of T1 is transmitted over VC2.

Cycle 5: Since there are no competing transactions, beat T1D of T1 is transmitted over VC2.

[0113] This example illustrates (1) with a virtual channel with absolute priority, access to the shared interconnect 12 is immediately awarded whenever traffic becomes available, regardless of whether or not other traffic has been previously waiting and (2) the winning portions or beats of different transactions are interleaved and transmitted over different virtual channels associated with the interconnect 12. In this example, virtual channel VCO was given absolute priority. It should be understood that with absolute or strict priority schemes, any of the virtual channels may be assigned the highest priority.

Interleaving - Example Two

[0114] Referring to Fig. 6, a second example of the interleaving of portions of different transactions over multiple virtual channels is illustrated.

[0115] In this example, the priority scheme for access to the interconnect 12 is weighted, meaning VCO is awarded access (40%) of the time and VC1-VC3 are each

awarded access (20%) of the time respectively. Also, the interconnect is 128 bits wide.

[0116] Further in this example, there are four competing transactions, T1, T2, T3 and T4:

5 - T1 is assigned to VC0 and includes four (4) portions or beats T1A, T1B, T1C and T1D;

- T2 is assigned to VC1 and includes two (2) portions or beats T2A and T2B;

- T3 is assigned to VC2 and includes two (2) portions or beats T3A and T3B;

and

10 - T4 is assigned to VC3 and includes two (2) portions or beats T4A and T4B.

[0117] With this example the priority scheme is weighed. As a result, each virtual channel will win according to its weight ratio. In other words over the course of ten cycles, VC0 will win four times and VC1, VC2 and VC3 will each win two times. For instance, as illustrated in Fig 6:

15 - The four portions or beats T1A, T1B, T1C and T1D of T1 are transmitted over VC0 in four (40%) of the ten (10) cycles (i.e., cycles 1, 4 7 and 10);

- The two portions or beats of T2A and T2B of T2 are transmitted over VC1 in two (20%) of the ten (10) cycles (i.e., cycle 2 and cycle 6);

20 - The two portions or beats of T3A and T3B of T3 are transmitted over VC2 in two (20%) of the ten (10) cycles (i.e., cycle 5 and cycle 9); and

- The two portions or beats of T4A and T4B of T4 are transmitted over VC3 in two (20%) of the ten (10) cycles (i.e., cycle 3 and cycle 8);

[0118] This example thus illustrates: (1) a weighted priority scheme where each virtual channel is awarded access to the interconnect 12 based on a predetermined ratio and (2) another illustration of the winning portions of different transactions being interleaved and transmitted over different the virtual channels associated with the interconnect 12.

[0119] It should be understood with this weighted example there is sufficient traffic to allocate portions of transactions to the various virtual channels in accordance with the weighted ratios. If the amount of traffic on the other hand is insufficient, then the weighted ratios can be either strictly or not strictly enforced. For example, if there is a large degree of traffic on virtual channel VC3 and limited to no traffic on the

other virtual channels VC0, VC1 and VC2, then VC3 will carry all or a bulk of the traffic if the weighted ratio is strictly enforced. As a result, however, the interconnect 12 may be under-utilized as portions of transactions may not be sent every clock cycle or beat. On the other hand if the weighted ratio is not strictly enforced, then it is possible for the transaction traffic to be reallocated to increase the utilization of the interconnect (e.g., traffic is sent over a higher number of cycles or beats).

5 [0120] The above two examples are applicable regardless which of the above-described transmission modes are used. Once transaction(s) is/are divided into portions or beats, they can be interleaved and transmitted over the shared interconnect 10 12 using any of the arbitration schemes as defined herein.

[0121] The above-described arbitration schemes represent just a few examples. In other examples, low jitter, weighted, strict, round-robin or just about any other arbitration scheme may be used. The arbitration schemes listed or described herein should therefore be considered as exemplary and not limiting in any manner.

15 **Multiple Simultaneous Arbitrations**

[0122] Up to now, for the sake of simplicity, only a single arbitration has been described. It should be understood, however, that in real-world applications, such as on a SoC, multiple arbitrations may occur simultaneously.

[0123] Referring to Fig. 7, a block diagram of two shared interconnects 12 and 20 12Z for handling traffic in two directions between switches 16, 18 is illustrated. As previously described, the switch 16 is responsible for directing transaction traffic from source sub-functions 14 (i.e., IP₁, IP₂ and IP₃) to destination sub-functions 14 (i.e., IP₄, IP₅ and IP₆) over the shared interconnect 12. To handle transactional traffic in the opposite direction, switch 18 includes arbitration element 26Z and optionally ARL 25 28Z. During operation, elements 26Z and ARL 28Z operate in the complement of that described above, meaning transaction traffic generated by source IP agents 14 (i.e., IP₄, IP₅ and IP₆) is arbitrated and sent over shared interconnect 12Z to destination IP agents (i.e., IP₁, IP₂ and IP₃). Alternatively, the arbitration can be performed without the ARL 28Z, meaning the arbitration simply decides among competing ports 20 (e.g., 30 Port 3, port 3 or Port 5) and the portion of the transaction associated with the winning port is transmitted over the interconnect 12, regardless of the final destination of the

portion. As elements 12Z, 26Z and 28Z have previously been described, a detailed explanation is not provided herein for the sake of brevity.

[0124] In a SoC, there can be multiple levels of sub-functions 14 and multiple shared interconnects 12. With each, the above described arbitration scheme can be used to arbitrate among transactions sent over the interconnects 12 between the various sub-functions simultaneously.

IP Agent Reset and Power Management

[0125] Referring to Fig. 8, a block diagram of an SoC 800 having reset and power management functionality is illustrated. The SoC 800 includes an interconnect 802, a plurality of IP agents 14 (e.g., Agent 1 through Agent N), one or more links 803 connecting or coupling the IP agents 14 to the interconnect 802, and a system controller 804. Although not illustrated, each IP agent 14 may also include one or more dedicated "hard-wire" inputs for receiving reset input instructions. Such instructions may come from a number of sources, including from off the SoC, the system controller 804, or another IP agent 14, etc.

[0126] In various embodiments, the IP agents 14 may be disparate and may implement a wide variety of different functions.

[0127] The interconnect 802 can be a wide variety of different types of interconnects, such a Network on a Chip (NoC), a bus, a switching network, etc.

[0128] In various embodiments, the links 803 may each be a dedicated link or a bus between each IP agent 14 and the interconnect 802. Alternatively, access to the interconnect 802 can be shared among multiple IP agents 14 using one link 803 and an arbitration scheme is used to select among the competing IP agents 14. In yet another embodiment, a number of virtual channels may be associated with the one or more links 803, such as the virtual channels associated with the shared link as previously described.

[0129] The system controller 804 and the managers 806, 808 and 809 may also be implemented in a number of different ways. For instance, as a CPU or microcontroller, as programmable logic, a complex state machine for handling all or most system control functions on the SoC 800, a simple state machine for handling a few exception situations, or any combination thereof. The system controller 804 may reside on the SoC 800 as shown or, alternatively, located off the SoC 800 (not

illustrated). Where a state machine is used, the states and the transitions between the states is typically hard-coded into the SoC 800.

[0130] In yet other embodiments, one or more of the reset, power and/or quiesce managers 806, 808 and 809 can each be centralized within the system controller 804
5 as shown. Alternatively, each manager 806, 808 and/or 809 can be decentralized and distributed throughout various locations on the SoC 800 or even off the SoC. Each of the reset manager 806, the power manager 808 and the quiesce manager 809 can be implemented in software, hardware, programmable logic, a state machine or any other suitable means.

[0131] The reset manager 806 is responsible for managing the emergence of the various IP agents 14 on the SoC 800 from reset in an organized manner. A reset of an IP agent 14 may be required or desired under a number of circumstances. For instance, a "cold reset" occurs following removal or disruption of power provided to the SoC 800 or a system wide reset of the SoC 800. Alternatively, a "warm reset"
15 occurs when one, a group or even all the IP agents 14 (similar to a cold reset) are reset, but power is not removed or disrupted from the SoC 800. A warm reset can be implemented via signaling that originates either on the SoC 800 or externally. Regardless of how a reset is initiated, the reset manager 806 is responsible for managing the emergence of the IP agent 14 or IP agents 14 from reset in an organized
20 manner.

[0132] If an IP agent 14 is malfunctioning for some reason, it may have to be reset. Examples of malfunctioning IP agents 14 include situations where the IP agent 14 is non-responsive, is in an error state, or actively generating erroneous transactions. In yet other examples, an IP agent 14 may have to undergo a reset
25 operation upon exiting a lower power state, such as one of several power saving modes as described below.

[0133] The power manager 808 manages the process of placing the various IP agents 14 into a lower power state, typically one of several power saving modes. Depending on the mode, the power manager 808 may operate in cooperation with the
30 reset manager 806 to reset an IP agent 14 if necessary.

[0134] The quiesce manager 809 operates in cooperation with the system controller 804, reset manager 806, power manager 808 and the interconnect 802 to (1)

transition an operational or malfunctioning IP agent 14 into either a reset or a power savings mode where the IP agent becomes inoperable, (2) places the link 803 between the interconnect and the IP agent 802 into a quiescent state and (3) directs the interconnect to operate as a proxy for the IP agent while inoperable.

5 **[0135]** The memory 810 may include both volatile and non-volatile types of memory. In addition, the memory 810 may be centralized on the SoC 800 or may be widely distributed among the system controller 804, the interconnect 802, the links 803, and any of the managers 806, 808 and/or 809. In yet other embodiments, portions or all of the memory 810 may be provided off the SoC 800.

10 **[0136]** The volatile portions of the memory 810 are typically used for system memory, where the current data generated by the system controller 804, managers 806, 808, 809, interconnect 802, IP agents 14, etc., are stored. Such memory may include various caches, SRAM, DRAM, etc.

[0137] The non-volatile or persistent portions of memory 810 is typically used for
15 storing "boot-up" code for the SoC 800. The boot code enables the system controller 804, including the managers 806, 808, 809, the interconnect 802 and the IP agents 14, to each load their operating systems and/or other system software as needed to initiate operation after powering on. The reboot process typically includes a number of self-tests, which when completed, allow the entire system, including each of the IP agents
20 14, to perform their normal operations. The non-volatile or persistent portions may be implemented using NVRAM (non-volatile random-access memory), EEPROM (electrically erasable programmable read only memory), a hard drive, CD ROM, etc.

[0138] The reset manager 806 is responsible for coordinating the emergence from reset of any of the IP agents 14 in an organized manner. As noted herein, a reset of a
25 given IP agent 14 may occur for any number of reasons, including (1) when the entire SoC 800 emerges from reset following an external reset, a re-start command or a power-on event or (2) or an individual IP agent 14 reset during operation of the SoC 800 due to malfunction, following a power down or sleep mode, etc. Regardless of the reason, a given IP agent 14 is ready to be introduced to the interconnect 802 once its
30 internal reset sequence is complete. Upon emergence from reset, a negotiation is then coordinated between the IP agent 14 and its IP port 20 on the interconnect 802 over the link 803.

[0139] Referring to Fig. 9 is a flow diagram showing an exemplary IP agent reset negotiation sequence between an IP agent 14 and the interconnect 802.

[0140] In the initial step 902, a determination is made if an IP agent 14 has emerged from reset and is ready to be introduced to the interconnect 802 or not. When
5 emergence occurs, the subsequent steps 904 through 912 are followed to reintroduce the IP agent 14 to the interconnect 802.

[0141] In step 904, the interconnect 802 generates inquiries for the IP agent 14 at periodic intervals. With each inquiry, the interconnect 802 essentially asks the IP agent 14 if it is "awake" (i.e., is it transaction ready, meaning is it capable of sending
10 or processing received transactions).

[0142] In decision 906, the interconnect determines if it has received a positive response to the inquiry(s) from the IP agent 14. If not, then the interconnect 802 continues to send the inquiries. If yes, then it signifies to the interconnect 802 that the IP agent 14 has partially completed its reset routine and is ready for the next phase of
15 the negotiation.

[0143] In step 908, the interconnect 802 and the IP agent 14 continue their negotiation by exchanging their credit information respectively. The interconnect 802 and the IP agent 14 each exchange with the other the available number of beats (i.e., the amount of data that can be transmitted over the link 803 per clock cycle. Each
20 partner on opposing sides of the link 803, after the exchange, knows the available number of credits the other has as a result of this negotiation.

[0144] In an optional step 910, interconnect 802 and the IP agent 14 continue their negotiation by exchanging other useful information such as security credentials, an agreed upon number of virtual channels that may be associated with the link 803
25 coupling the interconnect 802 and the IP agent 14, etc.

[0145] In the last step 912, when the negotiation is complete, the IP agent 14 is declared "transaction ready". In other words, the IP agent is ready to either process incoming transactions received from the interconnect 802 or to send outgoing transactions over the interconnect 802 to another destination. Once the IP agent 14 is
30 transaction ready, it becomes visible to both the interconnect 802, the system controller 804 and any other element connected or otherwise coupled to the

interconnect 802, either directly or indirectly through intermediate circuitry, logic or other element.

[0146] The reset manager 806 is also responsible for coordinating the reset of malfunctioning IP agents 14. During operation of the SoC 800, an IP agent 14 may misbehave (e.g., become non-responsive, enter an error state, erroneously generate transactions, or otherwise malfunction). For instance, the IP agent may be unable to process a received transaction. As a result, the originating IP agent that sent the transaction may get hung up waiting for a response. Depending on the severity of the problem, the hang up can be limited to just the originating IP agent 14, the destination IP agent 14, or in a worst case scenario, other portions or even the entire SoC 800 may be adversely affected. Accordingly, in certain circumstances, the misbehaving IP agent may need to be reset to correct the issue.

[0147] Referring to Fig. 10, a flow diagram 1000 showing a reset sequence for a malfunctioning IP agent is shown.

[0148] In step 1002, the various IP agents 14 on the SoC 800 operate as normal by generating transmitted transactions and/or processing received transactions.

[0149] In decision step 1004, the system controller 804 monitors the operation of the IP agents. If no problems are detected, then the IP agents 14 continue their normal operation. On the other hand if an IP agent malfunctions, for any reason, then the reset manager 806 flags it as a malfunctioning IP agent 14.

[0150] In step 1005, the system controller 804 and interconnect 802 further cooperate to initiate a number of processes that help the remainder of the SoC 800 operate without further issues or problems. These additional processes may include:

1. The system controller 804 requests that the interconnect 802 disallow any further transactions from being generated by the malfunctioning IP agent 14;
2. Keeping track of outstanding transactions targeting the malfunctioning IP agent 14;
3. The interconnect 802 may act as a proxy and respond to any transactions targeted for the malfunctioning IP agent 14 while undergoing the reset negotiation process. For example, the interconnect 802 may generate an exception message in response to the non-processed transaction. By acting as a proxy, potentially much larger system wide issues are avoided, including the entire system

getting hung up because the sender of the transaction never received a response from the malfunctioning IP agent 14. In various embodiments, the exception message may be a number of different types, such as the IP agent 14 is not available, the IP agent is in a low power mode, etc. In general, a wide variety of different types of exception
5 messages may be used, each indicative of the condition or error that has occurred.

[0151] In step 1006, the reset manager 806 generates a reset instruction for the malfunctioning IP agent 14.

[0152] In 1007, the link 803 between the IP agent 14 to be reset and the interconnect 802 is placed in a quiescent state. This process is further described with
10 regard to Fig. 14.

[0153] In step 1008, the malfunctioning IP agent 14 initiates its reset routine in response to the instruction received over the interconnect 802 or which may be received via a dedicated reset wire. This process involves the IP agent 14 (1) executing its own reset protocol or routine and (2) negotiating with the interconnect
15 802, as described above with regard to Fig. 9.

[0154] In decision step 1012, it is determined if the reset negotiation of the IP agent 14 is complete. When complete, control returns to step 1002 and operation of the IP agents 14 and the SoC 800 resume as normal. As noted above, the reset IP agent 14 becomes visible to the interconnect 802 and the system controller after
20 emerging from the reset and becomes transaction ready. Finally, in step 1014, the link 803 between the now reset IP agent 14 and the interconnect 802 exits the quiescent mode. At this point, the interconnect 802 no longer needs to act as a proxy for the IP agent 14.

[0155] The power manager 808 is responsible for intelligently and selectively
25 placing IP agents 14 into a lower power state, by placing the IP agents 14 in one of several power down modes. The powering down or placing of IP agents 14 into a powered down mode can be performed for a variety of reason.

[0156] For example, if the SoC 800 is used in a battery powered device, the power manager 808 may place IP agents into a power down mode to preserve limited battery
30 power. Alternatively, even in non-battery powered devices, the power manager 808 may place non-critical IP agents 14 into a low power mode to prevent overheating. These are just a few of the possible reasons for implementing power management.

Other reasons may include placing one or more IP agents 14 in a power down mode if they are not being used. In various alternative embodiments, the power down modes include:

1. Low Power Mode, Operational: In one alternative, the clock frequency
5 for the IP agent 14 is slowed down if applicable. Alternatively, the supply voltage may be reduced if applicable. In yet another embodiment, both the clock frequency and supply voltage may be reduced if applicable further reducing power consumption. It should be understood that reducing the clock frequency and/or supply voltage is done only when applicable, meaning not all IP agents 14 have the ability to operate at
10 either a reduced clock frequency, a reduced supply voltage, or both. In yet other embodiments, the commands for reducing the clock and/or supply voltage, when applicable, can be derived from the system controller 804 or the IP agent 14 itself, provided the IP agent 14 has a low power, operational mode.

Since the IP agent remains functional, the interconnect 802 may not play a
15 significant role in this mode, meaning it may not have to act as a proxy for the IP agent 14 and generate responses for incoming transactions since the IP agent 14 can generate the response itself. However, the system controller 804 and/or interconnect 802 may reconfigure the link 803 settings for the IP agent 14 since its performance capability may be reduced while operating at the lower clock frequency. The
20 setting(s) that may possibly be changed include the arbitration settings for the IP agent 14 or the possible a reduction in the count of permitted outstanding transactions. When the IP agent exits this Low Power Mode, the voltage is first increased (if decreased) followed by an increase in the clock frequency (if decreased) and any changes to the link 803 settings reverted back to normal operational mode (if
25 reconfigured).

2. Low Power, Inoperable Mode, State Information Maintained: In this mode, the clock is shut off and the power supply is reduced, but may not be turned off completely. As a result, state information maintained in memory in the IP agent 14 is retained. Prior to entering this mode, the interconnect 802 "drains" the transactions
30 that the IP agent has already issued by preventing new transactions from being initiated and waiting for outstanding transactions to complete. Once all the transactions are drained, the interconnect 802 may act as a proxy and perform similar

processes (1), (2) and (3) as described above with regard to the resetting of a malfunctioning IP agent 14. When the IP agent is returned to normal and exits this mode, the voltage is first increased followed by an increase in the clock frequency.

3. Low Power, Inoperable, Mode - No State Information Retained: This mode is similar to mode 2 described immediately above, except the power is reduced to a point where state information maintained in the IP agent is lost. The interconnect 802 operates as a proxy as discussed above in this mode as well. When powered back up, the IP agent is required to go through a reset negotiation process, similar to that as described above with regard to Fig. 9.

10 4. Power Off Mode: In this mode, the clock is turned off and power is completely removed. The interconnect 802 operates as a proxy as discussed above. Upon power up, the supply voltage is first ramped up followed by the reset negotiation process as described above with regard to Fig. 9.

[0157] FIG. 11 is a flow diagram 1100 illustrating a sequence for placing an IP agent 14 in and out of the Low Power, Operational Mode.

[0158] In the initial step 1102, the IP agent 14 on the SoC 800 operates in its normal mode, meaning the standard clock frequency and voltage are used.

[0159] In decision step 1104, conditions within the SoC 800 are monitored by the system controller 804. If operating conditions are relatively normal or no event occurs triggering a power down of the IP agent 14, then the SoC and IP agent 14 continues to operate in its normal mode per step 1102. However, if a trigger condition is met (e.g., a reduced battery supply, overheating, etc.), then the power manager 808 may elect to place the IP agent 14 into the low, power operational mode.

[0160] In an optional step 1106, the interconnect 802 may elect to reconfigure the link 803. The reconfiguration may include changing the arbitration settings for the IP agent 14 or reducing the count of possible outstanding transactions to take into account the lower processing capability of the IP agent when operating at the lower power mode.

[0161] In step 1108, the operating clock frequency of the IP agent 14 is reduced if applicable. With the reduced clock frequency, the IP agent consumes less power.

[0162] In step 1110, the voltage supplied to the IP agent is reduced if applicable. By reducing the voltage, further power savings can be realized.

[0163] With the clock frequency and/or the voltage reduced, the IP agent 14 remains operational. As a result, it is capable of processing transactions, although possibly at a slower rate when operating at its standard clock frequency and/or supply voltage. In optional embodiments, the interconnect 802 can act as a proxy as
5 described above or can be adjusted or reconfigured to take into account and support the lower rate of performance of the IP agent 14 in the low power mode. Since these alternatives are optional, they do not necessarily have to be implemented.

[0164] In decision step 1112, the IP agent 14 operates in the low power mode until a decision is made to resume normal operation. In which case, the IP agent 14
10 undergoes a sequence to resume normal operation.

[0165] In optional step 1114, the voltage is increase to the standard operating voltage if applicable (i.e., if the voltage was previously decreased).

[0166] In step 1116, the clock frequency is increased if applicable (i.e., provided the clock was previously decreased. In step 1117, the IP agent returns to normal
15 operation.

[0167] Finally, in optional step 1118, the interconnect returns any reconfigured interconnect setting to normal. At this point, the IP agent is ready to resume normal operation, as provided in step 1102.

[0168] Referring to Fig. 12, a flow diagram 1200 illustrating a sequence for powering down/up an IP agent 14 in the Low Power, Inoperable, State Information
20 maintained mode is illustrated.

[0169] In step 1202, the IP agent 14 operates in its normal mode.

[0170] In step 1204, a decision is made to operate the IP agent 14 in low power, inoperable, state information maintained mode.

[0171] In step 1206, the link 803 is placed in the quiescent state and the interconnect 802 is configured to operate as a proxy for the IP agent 14. This typically involves (1) disallowing any new transactions from being generated by the IP agent
25 14, (2) waiting for any outstanding transactions to complete and then (3) acting as a proxy by responding to any transactions targeted for the IP agent 14. For example, the interconnect 802 may send an exception message to the source of the non-processed
30 transaction, possibly preventing a hang up situation from occurring because the sender of the transaction never received a response from the IP agent 14.

[0172] In step 1208, the clock frequency of the IP agent 14 is reduced if applicable.

[0173] In step 1210, the operating voltage of the IP agent 14 is reduced if applicable. However, the voltage remains adequate so that memory or storage
5 elements in the IP agent 14 maintain their state information.

[0174] In decision 1212, the IP agent 14 remains in the lower power state until a decision has been made to resume normal operation. The system controller 804, an event external to the SoC (e.g., a signal received from a sensor, signal received an external source, etc.), a timer, the IP agent itself or another IP agent, can all trigger the
10 wake-up. When this decision is made, the IP agent undergoes a sequence to resume normal operation.

[0175] In steps 1214 and 1216, the voltage and clock frequency provided to the IP agent 14 are each increased if applicable. Since the state information has been retained, the IP agent 14 resumes normal operation in step 1217.

[0176] In step 1218, link 803 exists the quiescent mode and the IP agent becomes
15 transaction ready and the interconnect 802 is notified that it longer has to act as a proxy.

[0177] Referring to Fig. 13, a flow diagram 1300 illustrating a sequence for the Low Power, Inoperable, Mode is illustrated. In this sequence, steps 1202, 1206 and
20 1212 are the same as described above with regard to Fig. 12. As such, a discussion of these steps are not repeated herein.

[0178] In steps 1302, a decision is made to power down the IP agent 14. Thereafter, the interconnect is configured as a proxy (step 1206) and the clock for the IP agent 14 is turned off completely (if applicable) and/or the voltage is significantly
25 reduced (if applicable) to the point where state information is lost in step 1304. Without state information, when a decision is made to resume normal operation per step 1212, the voltage is ramped up (if applicable) and clock turned on (if applicable) in step 1306. Thereafter, the IP agent 14 undergoes a reset operation, as previously described with regard to Fig. 9. Once the reset is complete, the IP agent 14 becomes
30 transaction ready. The system then waits for the link to exit the quiescent mode in step 1310. Once the exit occurs, the IP agent is visible on the interconnect 802.

Thereafter, in step 1312, the interconnect 802 no longer acts as a proxy for the IP agent 14.

[0179] Finally, for the Power Off Mode, the sequence is the same as Fig. 13, except the power is turned off completely, as opposed to simply reduced. Otherwise
5 the Power Off Mode sequence is the same. In this mode, the IP agent 14 consumes virtually no power, is inoperable, and the interconnect 802 may act as a proxy on behalf of the IP agent.

[0180] Referring to Fig 14, a flow chart 1400 illustrating the steps for placing a link 803 in the quiescent state is illustrated.

10 [0181] In the initial step 1402, the system controller 804 makes a decision that an IP agent 14 should be either reset or placed in one of the inoperable power saving modes.

[0182] In step 1404, the IP agent 14 is instructed to stop generating transactions.

[0183] In decision 1406, the system determines if all outstanding transactions are
15 complete. For all outstanding Non-posted transactions, a Completion transaction must be received (i.e., with read transactions, the accessed data must be returned, with non-posted write transactions, an acknowledgement must be received). With Posted transactions, no response transaction is required. Posted transactions are therefore considered "complete" once they are sent by the IP agent.

20 [0184] In step 1408, the link 803 is placed in the quiescent state when all the outstanding transactions are complete. Thereafter, the interconnect 802 is configured as a proxy for the IP agent 14.

[0185] In step 1410, the IP agent is ready to placed in either reset or the desired inoperable low power mode.

25 [0186] Figs 15A-15D show flow diagrams of various for IP agent "wake-up" sequences.

[0187] Referring to Fig. 15A a flow diagram 1500 illustrating an agent-initiated "wake-up" sequence is illustrated. In this embodiment, the wake up sequence is initiated by the IP agent, but implemented through the system controller 804.

30 [0188] In step 1502, an IP agent 14 in an inoperable state detects a wake-up trigger event. Although an IP agent may be powered down or "off", it may remain at least partially functional in the sense that it maintains the ability to detect when a

wake-up trigger occurs. The wake-up trigger may include a number of different types of events. For example, it could be an internal timer that causes the IP agent 14 to wake-up after a predetermined period of time, or it can be an event external to the SoC 800, such as another device that wishes to communicate with the IP agent 14. In
5 step 1504, the IP agent sends a "wake-up" communication over its link 803 to the interconnect 802. Again, although the link is in the quiescent state when its corresponding IP agent 14 is in an inoperable state, it is capable of transmitting the wake-up signal to the interconnect 802.

[0189] In step 1506, the interconnect 802 is configured to "listen" for a wake-up
10 signal from an inoperable IP agent. If the signal is detected, the interconnect 802 notifies the system controller 804.

[0190] In step 1508, the system controller 804 may send command(s) over the interconnect 802 for the IP agent 14 to initiate its wake-up sequence.

[0191] In step 1510, the IP agent initiates its wake-up sequence in response to the
15 command(s).

[0192] With the embodiment described above, the IP agent 14 asks the system controller to initiate the wake-up sequence. In response to a wake-up command from the system controller, the IP agent initiates its own wake-up sequence. The system controller is therefore aware of the status of the IP agent as it emerges from its
20 inoperable state and becomes visible on the interconnect 802.

[0193] Fig. 15B shows the sequence when the system controller 804 initiates a wake-up of an IP agent 14. With this sequence, the system controller 804 sends wake up command(s) to the IP agent in step 1508, and in response, the IP agent initiates its own wake up sequence in step 1510. In a variation of this embodiment (not
25 illustrated), the wake up may be initiated off the SoC 800 via the system controller 804. When the system controller 804 receives the command(s), the above described process is initiated.

[0194] Fig. 15C shows the sequence when the wake up command for an IP agent 14 that originates off the SoC 800 and is implemented through the system controller
30 804. With this sequence, the system controller 804 receives the command in step 1512. In response, the system controller sends a wake up command to the IP agent in step 1508, and in response, the IP agent initiates its own wake up sequence in step

1510. With direct wake up from off the SoC 800, the command is provided directly to the IP agent 14 via its hard-wire input. In response, the IP agent initiates its own wake up sequence.

[0195] Referring to Fig. 15D, a flow diagram 1520 illustrating an IP agent
5 initiated and implemented wake up sequence is illustrated. In this embodiment, a wake up condition, such as any of those noted above, occurs in step 1522. In response, the IP agent initiates its own wake up sequence in step 1524. In step 1526, the wake up sequence completes. Thereafter, in step 1528, the IP agent notifies the interconnect 802 and the system controller 804, either directly or through the
10 interconnect 802, of its awoken status.

[0196] In the above examples, sequences for transitioning a single IP agent into one of the above-described low power modes was described for the sake of simplicity. In actual embodiments, multiple IP agents 14 on an SoC may be powered down concurrently. If two or more are powered down at or around the same time, each
15 would independently undergo one the above described sequences, depending on the mode.

[0197] Although only a few embodiments have been described in detail, it should be appreciated that the present application may be implemented in many other forms without departing from the spirit or scope of the disclosure provided herein.
20 Therefore, the present embodiments should be considered illustrative and not restrictive and is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

What is claimed is:

1. A System on a Chip (SoC), comprising:
an interconnect;
an IP agent provided on the SoC and arranged to communicate with the
5 interconnect; and
a reset manager for initiating a reset of the IP agent, the reset involving a negotiation between the interconnect and the IP agent that results in the IP agent becoming transaction ready.
- 10 2. The SoC of claim 1, wherein the negotiation involves the interconnect generating at least one inquiry for the IP agent, the at least one inquiry requesting an "awake" status of the IP agent.
3. The SoC of claim 1 or 2, wherein the negotiation involves the interconnect generating a plurality of inquiries for the IP agent each requesting an "awake" status of the IP agent, the plurality of inquiries sent at periodic intervals respectively.
- 15 4. The SoC of any of claims 1-3, wherein the negotiation involves the IP agent generating a positive response in reply to an inquiry from the interconnect requesting an "awake" status of the IP agent, the positive response generated when the IP agent awakens and is transaction ready.
- 20 5. The SoC of any of claims 1-4, wherein the negotiation involves an exchange of credits available at the interconnect and at the IP agent respectively.
6. The SoC of any of claims 1-5, wherein the negotiation involves the interconnect and the IP agent sharing security information.
7. The SoC of any of claims 1-6, wherein the negotiation involves the interconnect sharing a number of virtual channels associated with a link coupling the
25 interconnect and the IP agent.
8. The SoC of any of claims 1-7, wherein, after the negotiation is complete, availability of the IP agent is made known to elements coupled to the interconnect including a system controller.
- 30 9. The SoC of any of claims 1-8, further comprising a link coupled between the interconnect and the IP agent, the negotiation conducted between the interconnect and the IP agent over the link.

10. The SoC of any of claims 1-9, wherein a plurality of virtual channels are associated with the link.
11. The SoC of any of claims 1-10, wherein the negotiation is conducted using link-level information that is transmitted over a link between the interconnect and the
5 IP agent.
12. The SoC of claim 11, wherein the link-level information is conveyed in one or more control packets that is/are transmitted over the link between the interconnect and the IP.
13. The SoC of claim 12, wherein the one or more control packets are transmitted
10 over the link between the interconnect and the IP agent and are not forwarded to another destination.
14. The SoC of any of claims 1-13, wherein the reset manager initiates the reset of the IP agent in response to a reset trigger event.
15. The SoC of claim 14, wherein the reset trigger event includes one of the
15 following:
- power up of the SoC;
 - the IP agent has malfunctioned; or
 - the IP agent is emerging from a low power mode.
16. The SoC of any of claims 1-15, further comprising a quiesce manager for
20 placing a link between the interconnect and the IP agent into a quiescent state.
17. The SoC of claim 16, wherein the quiesce manager is configured to place the link in the quiescent state after all outstanding transactions generated by the IP agent are complete.
18. The SoC of claim 17, wherein the quiesce manager is configured to prevent
25 the IP agent from generating new transactions while waiting for the outstanding transactions to complete.
19. The SoC of claim 17, wherein the quiesce manager is configured to place the link in the quiescent state before the reset manager initiates the reset of the IP agent.
20. A System on a Chip (SoC), comprising:
30 an interconnect;
a plurality of IP agents; and

- a reset manager arranged to implement a reset of the plurality of IP agents, the reset of each of the plurality of IP agents involving individual negotiations between the interconnect and each of the plurality of IP agents respectively, the individual negotiations resulting in each of the IP agents becoming transaction ready
5 respectively.
21. The SoC of claim 20, wherein each of the plurality of IP agents uses its own reset routine.
22. The SoC of claims 20 or 21, wherein each of the plurality of IP agents becomes transaction ready on its own time schedule respectively.
- 10 23. The SoC of any of claims 20-22, wherein each of the plurality of IP agents becomes transaction ready independent of the other IP agents and all the IP agents are not required to become transaction ready at the same time.
24. The SoC of any of claims 20-23, wherein the individual negotiations for each of the plurality of IP functions are the same regardless if the IP agents are different.
- 15 25. The SoC any of claims 20-24, wherein the interconnect is further configured to operate as a proxy on behalf of at least one of the plurality of IP agents prior to the at least one IP agent becoming transaction ready.
26. The SoC of claim 25, wherein the interconnect operates as the proxy for the at least one IP agent prior to becoming transaction ready by generating an exception
20 message to a source attempting to send a transaction to the at least one IP agent.
27. A system on a Chip (SoC), comprising:
a plurality of IP agents; and
an interconnect in communication with the plurality of IP agents, the interconnect arranged to operate as a proxy on behalf of at least one of the plurality of
25 IP agents when the at least one IP agent is not transaction ready.
28. The SoC of claim 27, wherein the at least one IP agent is not transaction ready when undergoing a reset.
29. The SoC of claim 27 or 28, wherein the at least one IP agent is not transaction ready while in a power down mode.
- 30 30. The SoC any of claims 27-29, wherein the at least one IP agent is not transaction ready when malfunctioning.

31. The SoC any of claims 27-30, wherein the interconnect acts as the proxy on behalf of the at least one IP agent by:
- ascertaining when the at least one IP function is not transaction ready;
 - ascertaining if a source sends a transaction to the at least one IP agent while
5 not transaction ready; and
 - sending an exception message to the source notifying the source that the at least one IP agent is not available.
32. The SoC any of claims 27-31, wherein the source, in response to the exception message, undergoes its own response procedure for transactions sent but not delivered
10 to a destination.
33. The SoC any of claims 27-32, further comprising a reset manager for initiating a reset of the plurality of IP agents respectively.
34. The SoC any of claims 27-33, wherein the interconnect is further configured to engage in individual negotiations with each of the plurality of IP agents when each
15 of the plurality of IP agents is reset respectively.
35. The SoC any of claims 27-34, wherein upon a reset of two or more of the plurality of IP agents, each of the two or more IP agents is arranged to individually reset and become transaction ready independent of the other IP agents respectively.
36. The SoC of claim 35, wherein each of the two or more IP agents
20 independently become transaction ready on their own time-schedule with no requirement that the two or more IP agents become transaction ready at the same time.
37. The SoC any of claims 27-36, further comprising a system controller, at least one IP agent becoming visible to the system controller when the at least one IP agent
25 becomes transaction ready.
38. A System on a Chip (SoC), comprising:
- a plurality of IP agents;
 - an interconnect;
 - one or more links connecting the plurality of IP agents with the interconnect;
- 30 and

a power manager for at least partially managing power consumption by the SoC by selectively placing at least one of the plurality of IP agents in a low power state,

5 wherein the interconnect acts as a proxy on behalf of the at least one IP agent when (a) in the low power state and (b) is unavailable for processing a transaction, the interconnect acting as the proxy by providing an exception message to a source attempting to send the transaction to the at least one IP agent.

39. The SoC of claim 38, wherein the at least one IP agent implements its own low power sequence when entering the low power state.

10 40. The SoC of any of claims 38 or 39, wherein the low power state includes a low power mode where the at least one IP agent remains transaction ready.

41. The SoC of claim 40, wherein the interconnect does not act as a proxy when the at least one IP agent is in the low power state and remains transaction ready, but adjusts settings on one of the one or more links to accommodate a lowered clock
15 frequency and/or voltage for operating the at least one IP agent.

42. The SoC of any of claims 38-41, wherein the low power state includes a low power mode where the at least one IP agent is no longer transaction ready.

43. The SoC of claim 42, wherein the at least one IP agent retains state information when in the low power mode and is no longer transaction ready.

20 44. The SoC of claim 42, wherein the at least one IP agent does not retain state information when in the low power mode and is no longer transaction ready.

45. The SoC of any of claims 38-44, wherein the low power state includes a power off mode where minimal power is supplied to the at least one IP agent.

46. The SoC of any of claims 38-45, wherein the low power state is implemented
25 by reducing an operating clock frequency for the at least one IP agent.

47. The SoC any of claims 38-46, wherein the low power state is implemented by reducing a voltage used by the at least one IP agent.

48. The SoC any of claims 38-47, wherein the low power state is implemented by:
reducing an operating clock frequency for the at least one IP agent; and
30 reducing a supply voltage used by the at least one IP agent after reducing the operating clock frequency.

49. A System on a Chip (SoC), comprising:
at least one IP agent;
an interconnect;
a link connecting the IP agent and the interconnect;
- 5 a power manager arranged to place the IP agent into an inoperable state; and
a quiescent manager arranged to place the link in a quiescent state by:
- (a) instructing the IP agent to stop generating transactions;
(b) placing the link in the quiescent state once all transactions generated
by the IP agent are complete; and
- 10 (c) causing the IP agent to initiate a sequence to enter the inoperable state.
50. The SoC of claim 49, further comprising configuring the interconnect as a
proxy on behalf of the IP agent when in the inoperable state.
51. The SoC of claim 50, wherein the interconnect acts as the proxy on behalf of
the IP agent by:
- 15 ascertaining if a source sends a transaction to the IP agent while in the
inoperable state; and
sending an exception message to the source notifying the source that the IP
agent is not available.
52. The SoC of any of claims 49-51, wherein the IP agent is further configured to
20 initiate a "wake-up" sequence when in the inoperable state in response to a wake-up
trigger event.
53. The SoC of claim 52, wherein the wake-up trigger event includes one of the
following:
- (a) the IP agent receiving a communication from a source external to the
25 SoC;
- (b) a predetermined time period has expired;
(c) when instructed by a system controller; or
(d) a valid transaction targeting the IP agent in the inoperable state.
54. The SoC of claim 52, wherein the wake-up sequence involves the IP agent
30 sending a wake-up communication over the link to the interconnect in response to the
wake-up trigger event.

55. The SoC of claim 54, wherein the wake-up sequence involves the interconnect forwarding the wake-up communication to a system controller.
56. The SoC of claim 55, wherein the wake-up sequence further involves the system controller sending a wake-up command to the IP agent in response to
5 receiving the wake-up communication forwarded by the interconnect.
57. The SoC of claim 56, wherein the wake-up sequence further involves the IP agent initiating its own wake-up protocol in response to the wake-up command.




FIG. 1




FIG. 2

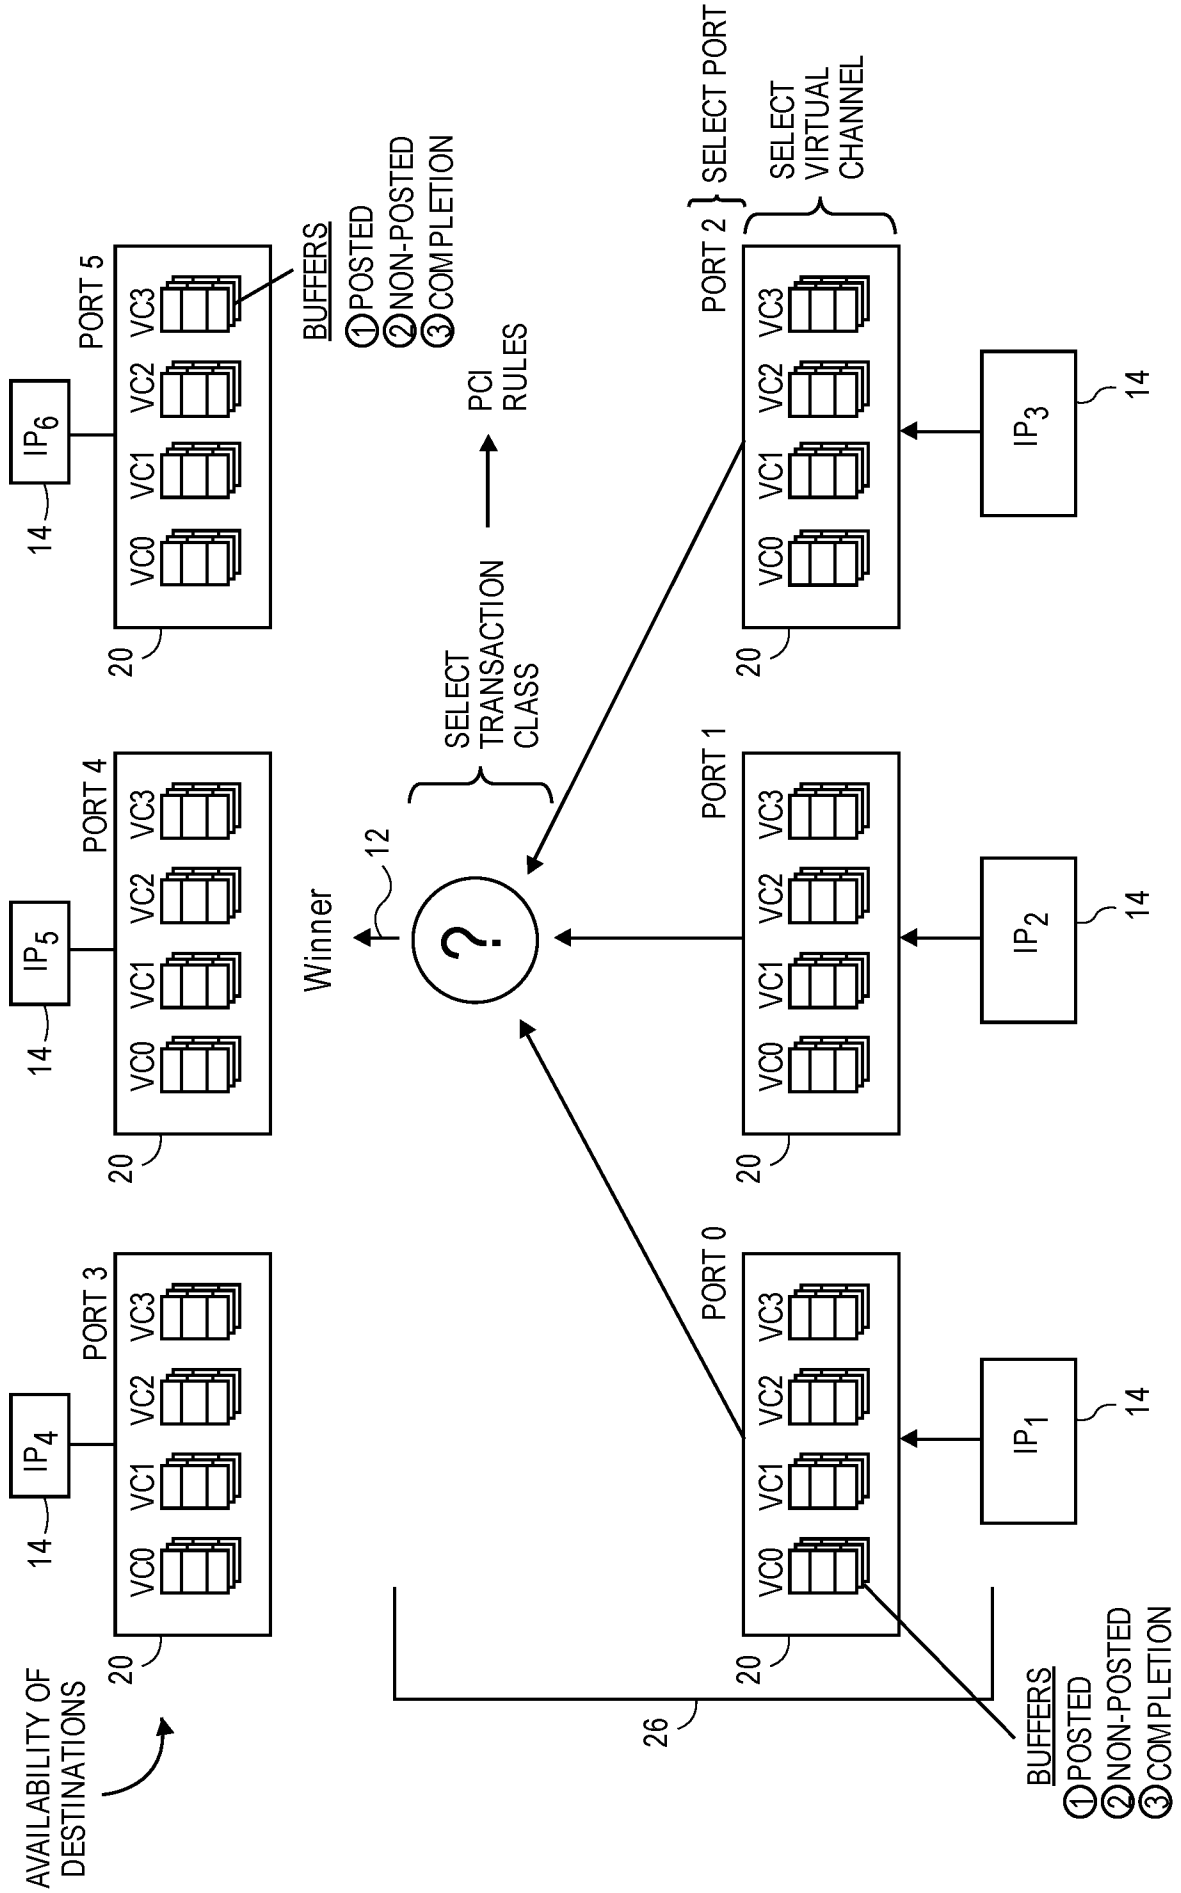


FIG. 3A


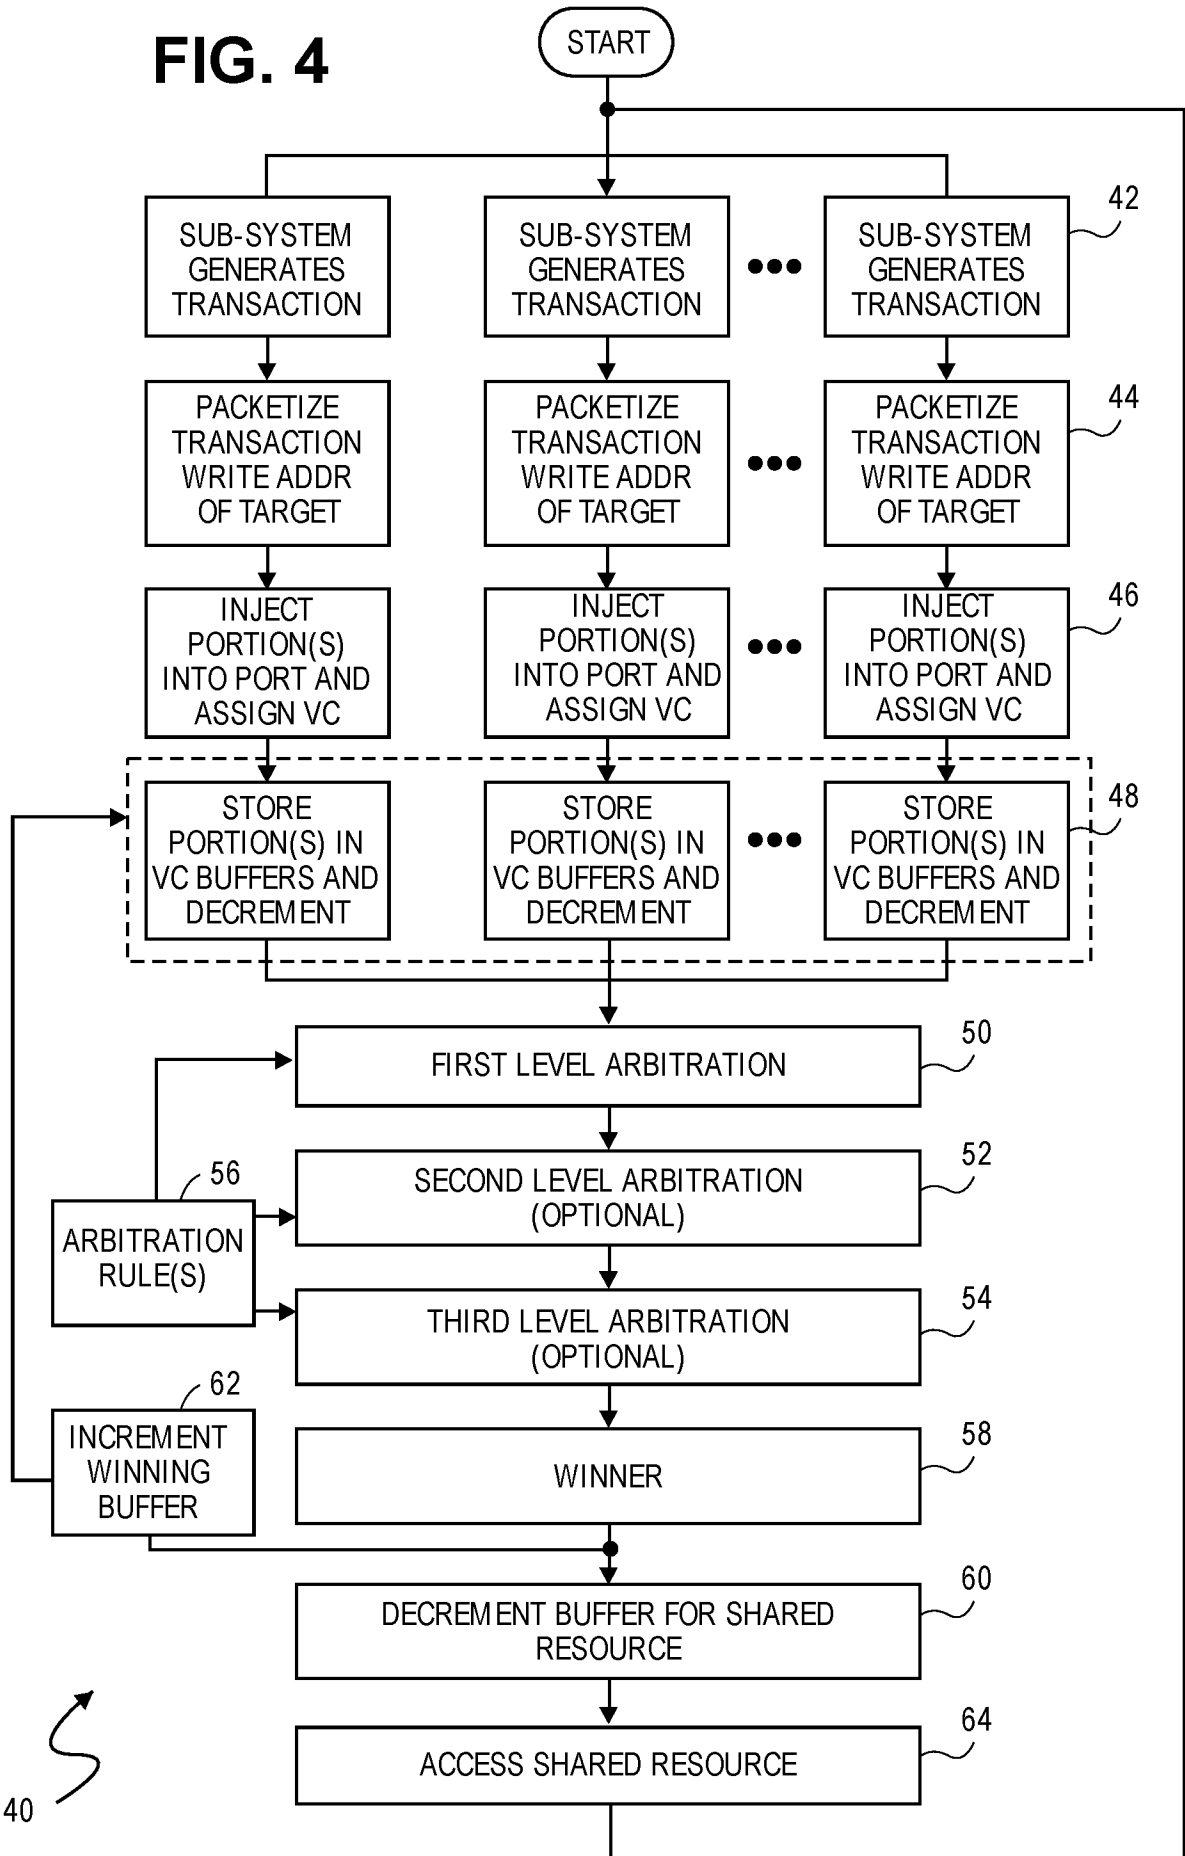


FIG. 4



Example 1
Strict Priority VCO, VC1, VC2, VC3

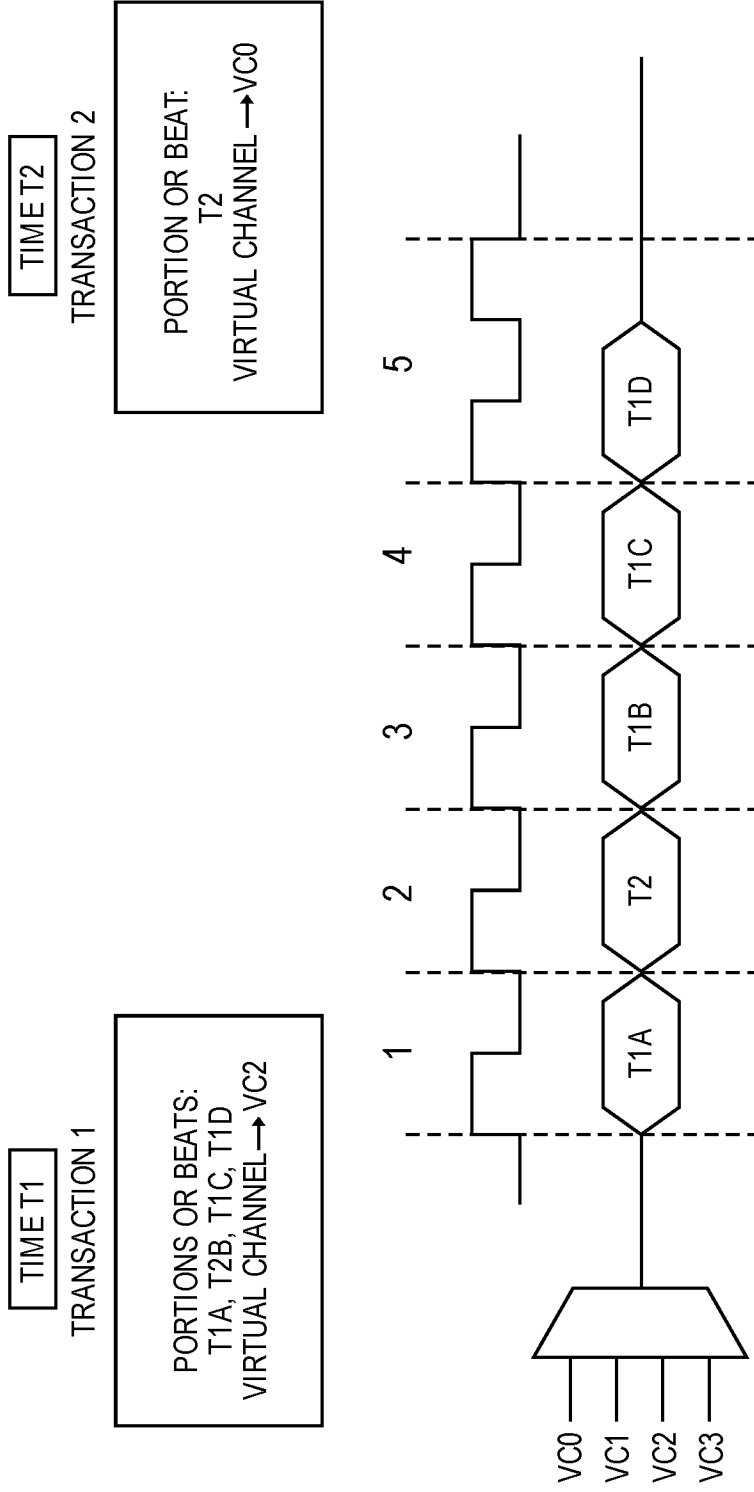


FIG. 5

Example 2
Weighted Priority: VC0=40%, VC1=20%, VC2=20%, VC3=20%

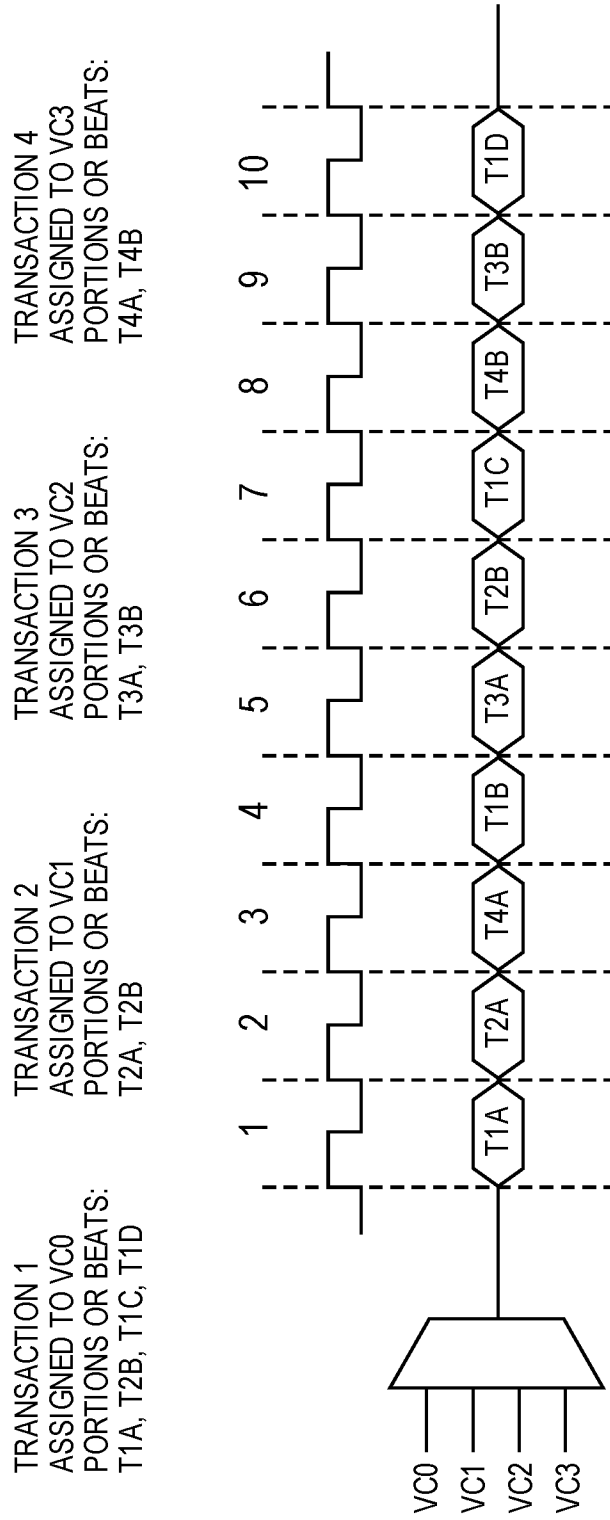


FIG. 6




FIG. 7




FIG. 8

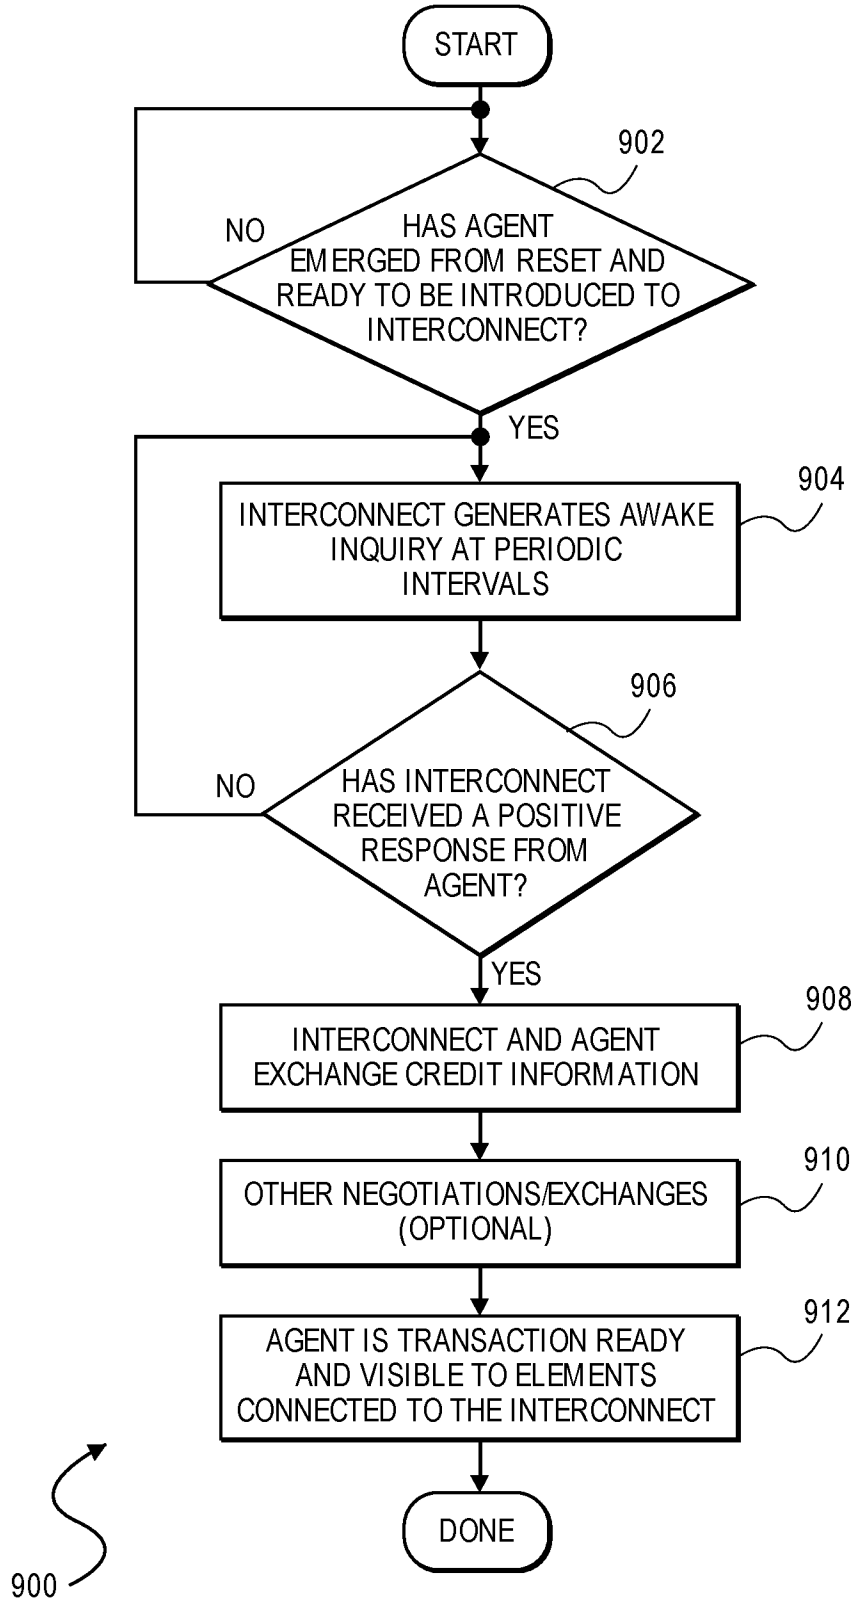



FIG. 9



1000

FIG. 10

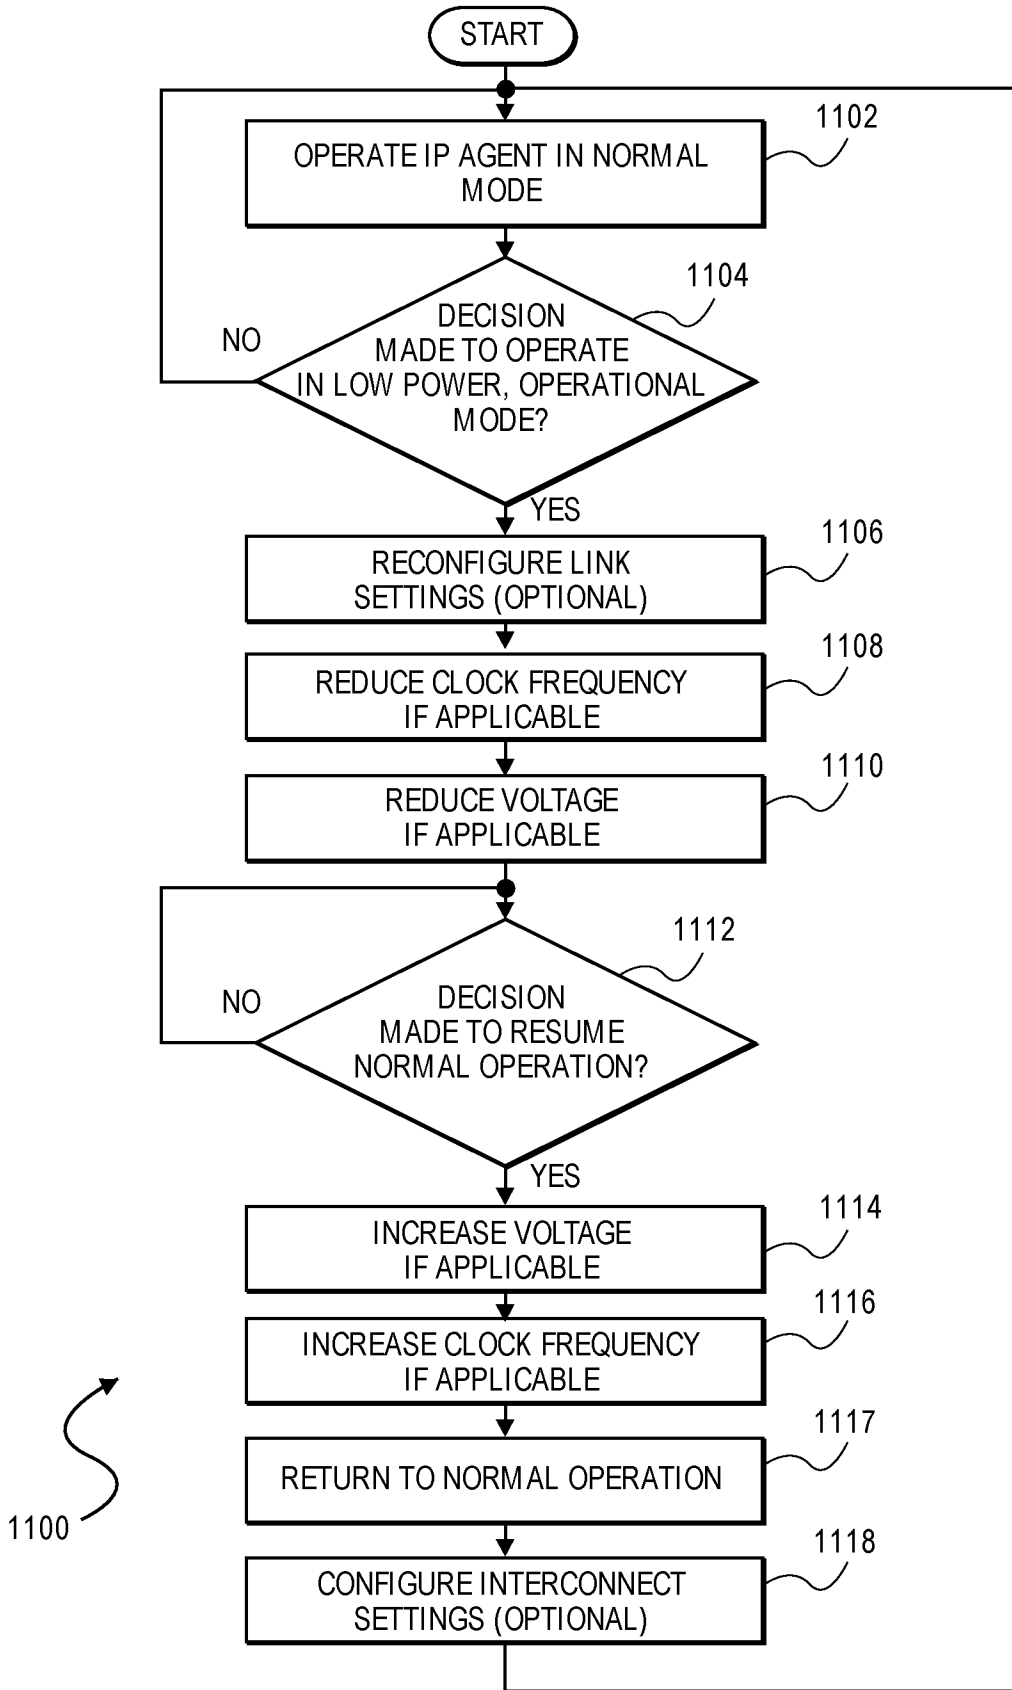


FIG. 11

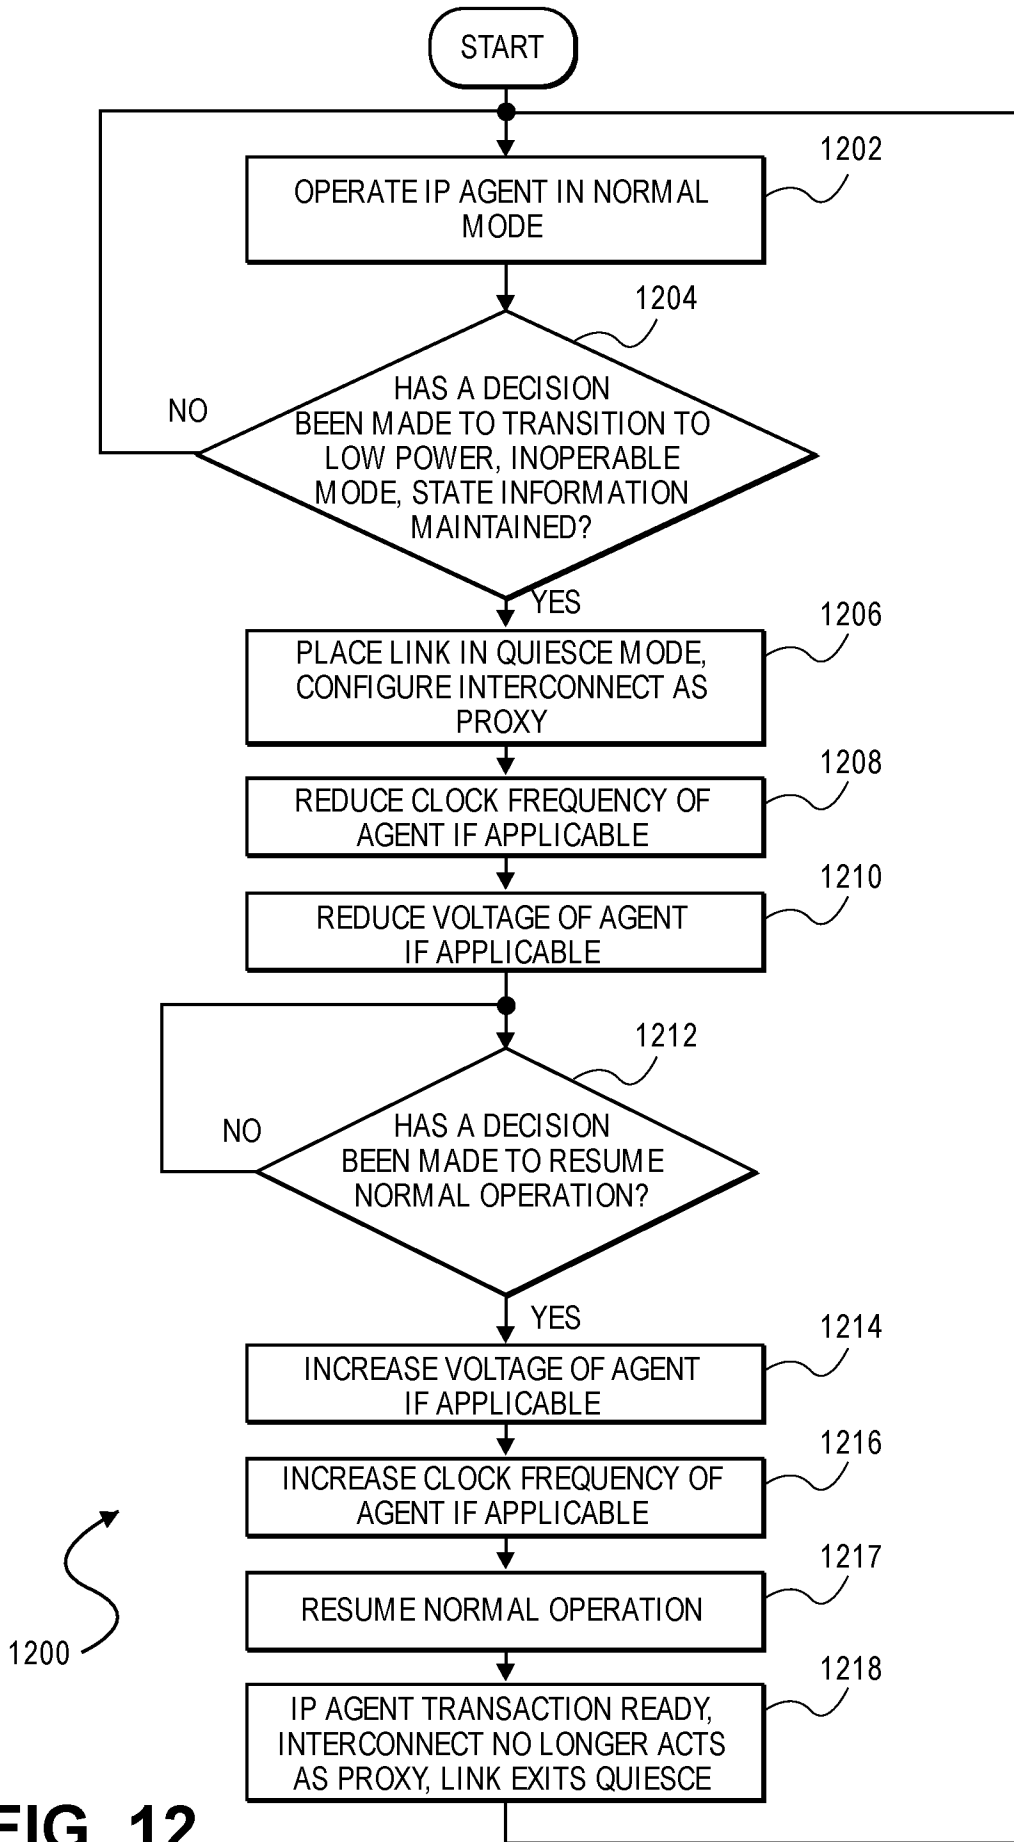


FIG. 12

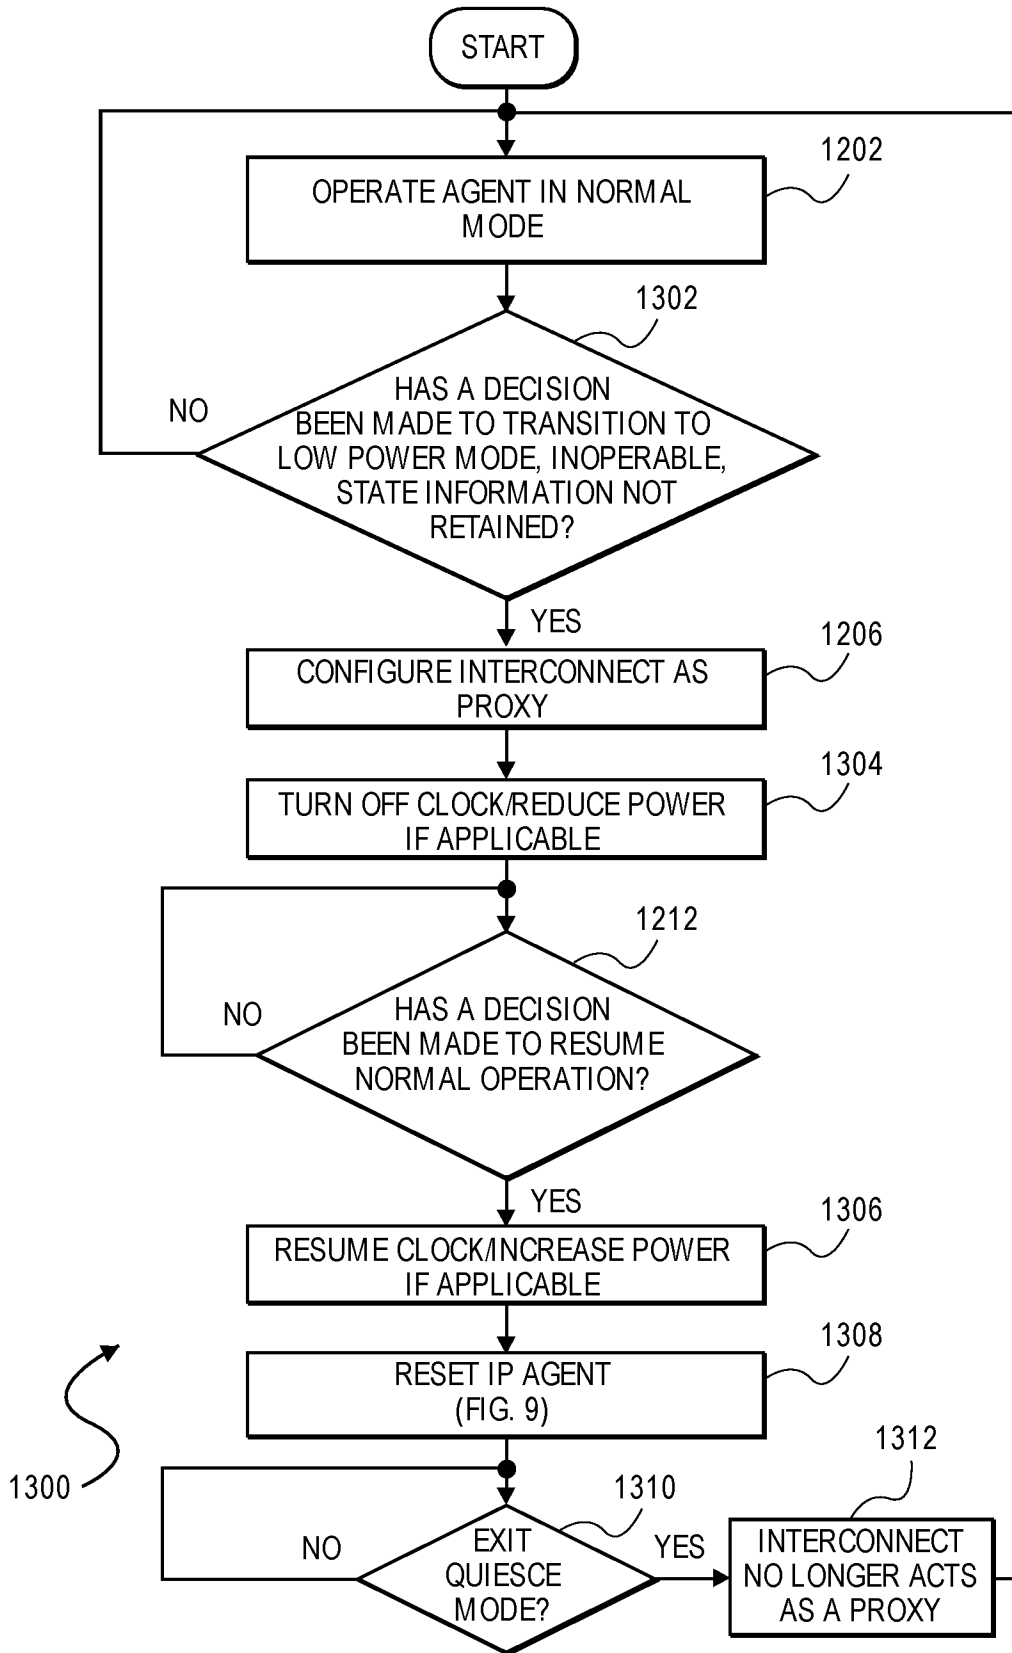


FIG. 13

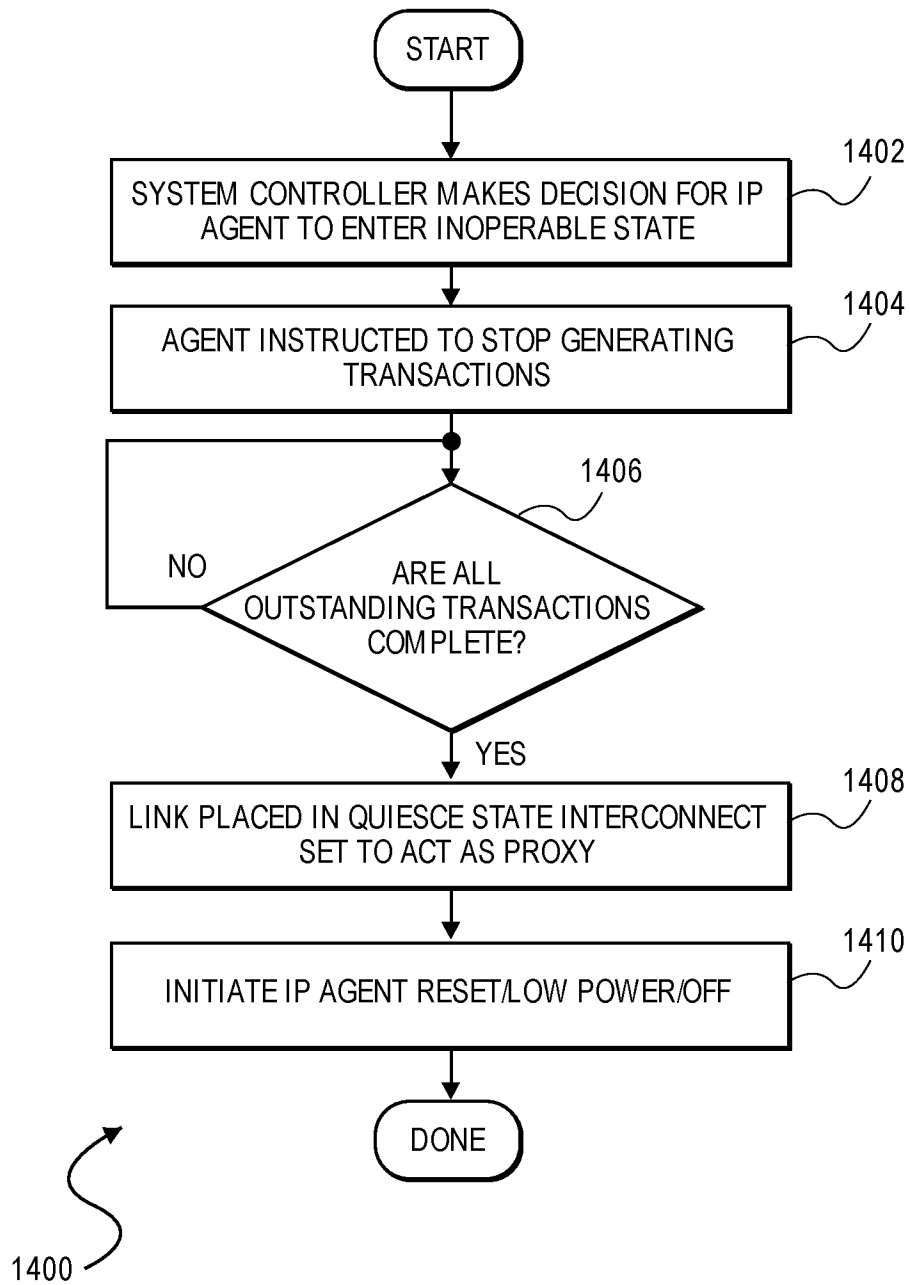
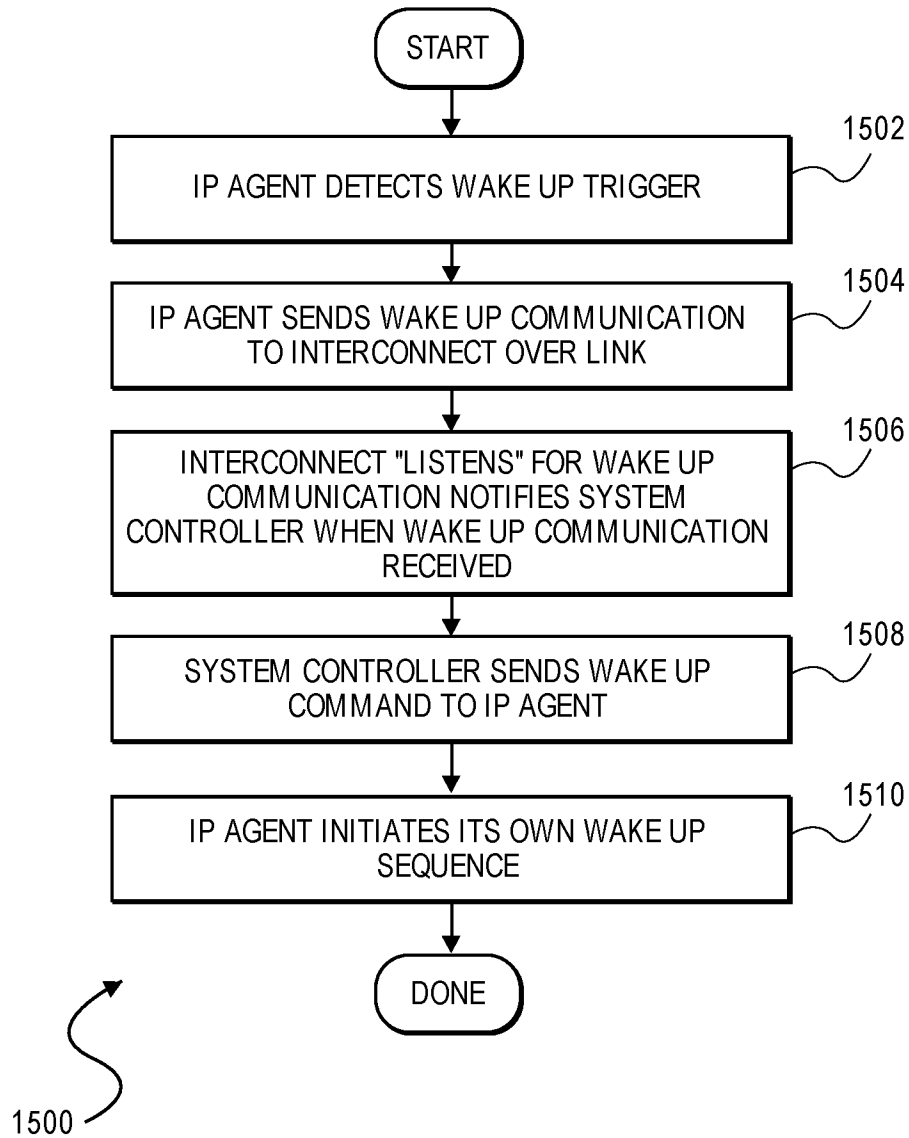


FIG. 14

**FIG. 15A**

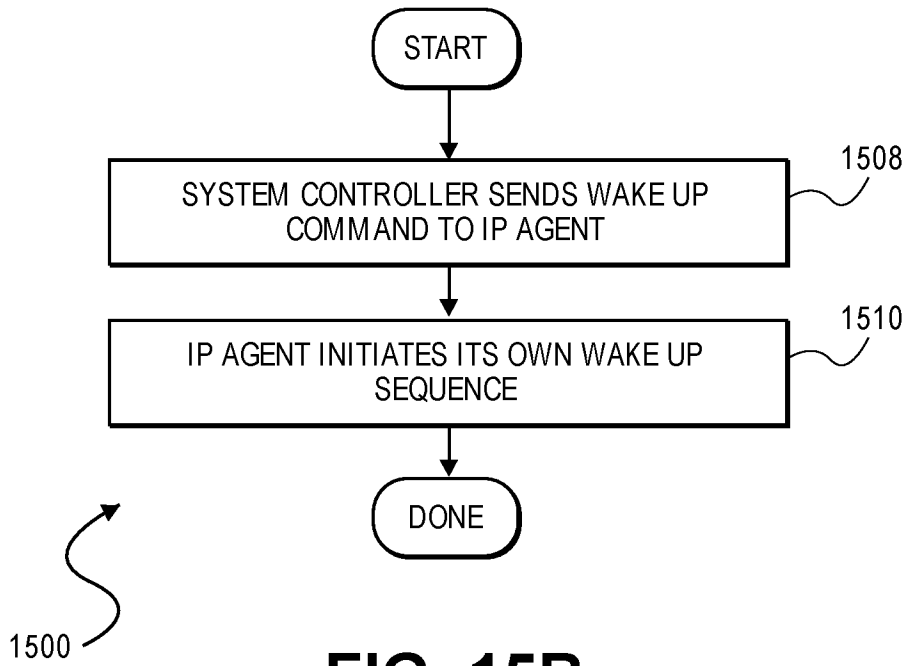


FIG. 15B

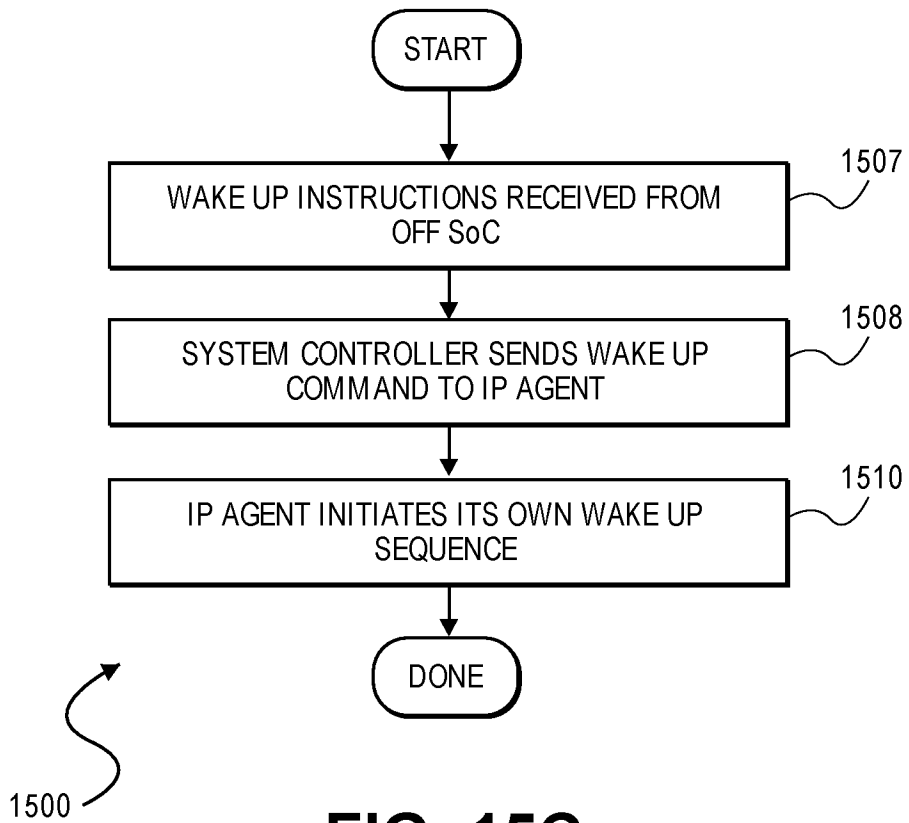


FIG. 15C

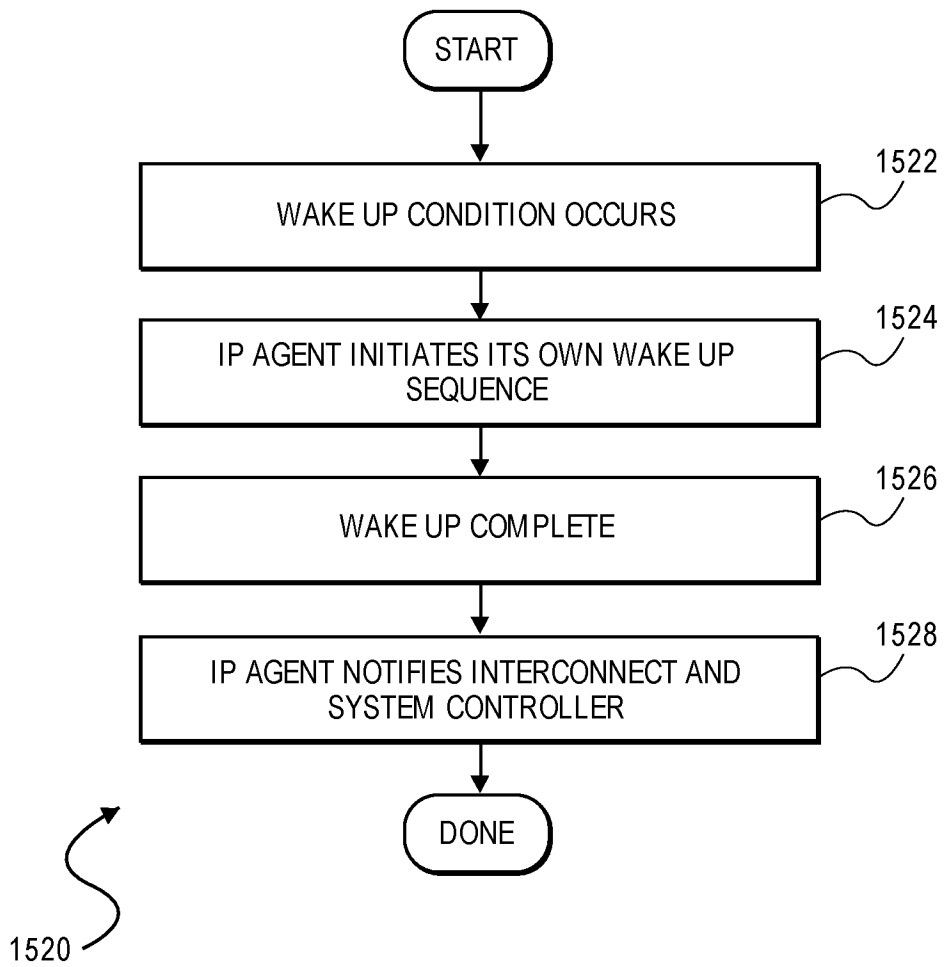


FIG. 15D

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US19/24586

A. CLASSIFICATION OF SUBJECT MATTER
 IPC - G06F 13/10, 13/12, 13/14 (2019.01)
 CPC - G06F 13/102, 13/122, 13/126, 13/14

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)
 See Search History document

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2013/0073878 A1 (JAYASIMHA, D et al.) 21 March 2013; abstract; figures 1, 9; paragraphs [0007], [0015], [0023], [0038]	1, 2, 20-22 --- 3
Y	US 9781058 B1 (JUNIPER NETWORKS, INC.) 03 October 2017; column 2, lines 20-24	3
A	US 7577857 B1 (HENDERSON, N et al.) 18 August 2009; entire document	1-3, 20-22

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"&" document member of the same patent family
"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	

Date of the actual completion of the international search
 30 May 2019 (30.05.2019)

Date of mailing of the international search report

22 JUL 2019

Name and mailing address of the ISA/
 Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, Virginia 22313-1450
 Facsimile No. 571-273-8300

Authorized officer
 Shane Thomas
 PCT Helpdesk: 571-272-4300
 PCT OSP: 571-272-7774

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US19/24586

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.: 4-19, 23-26, 30-37, 42-48
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Group I: Claims 1-3, 20-22; Group II: Claims 27-29, 38-41; Group III: Claims 49-57

-Continued in extra sheet-

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
Claims 1-3, 20-22

Remark on Protest

- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

-***-Continued from Box No. III - Observations where unity of invention is lacking-***-

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fee must be paid.

Group I: Claims 1-3, 20-22 are directed towards a reset manager to reset IP agents involving individual negotiations.

Group II: Claims 27-29, 38-41 are directed towards an interconnect acting as a proxy providing an exception message to a source attempting transactions.

Group III: Claims 49-57 are directed towards a quiescent manager placing a link in a quiescent state once all transactions are complete.

The inventions listed as Groups I-III do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

The special technical features of Group I include at least a reset manager arranged to implement a reset of the plurality of IP agents, the reset of each of the plurality of IP agents involving individual negotiations between the interconnect and each of the plurality of IP agents respectively, the individual negotiations resulting in each of the IP agents becoming transaction ready respectively, which are not present in Groups II-III.

The special technical features of Group II include at least wherein the interconnect acts as a proxy on behalf of the at least one IP agent when (a) in the low power state and (b) is unavailable for processing a transaction, the interconnect acting as the proxy by providing an exception message to a source attempting to send the transaction to the at least one IP agent, which are not present in Groups I and III.

The special technical features of Group III include at least a quiescent manager arranged to place the link in a quiescent state by: (a) instructing the IP agent to stop generating transactions; (b) placing the link in the quiescent state once all transactions generated by the IP agent are complete; and (c) causing the IP agent to initiate a sequence, which are not present in Groups I-II.

The common technical features shared by Groups I-III are a System on a Chip "SoC" comprising an interconnect; one or more links connecting the plurality of IP agents with the interconnect; and a power manager managing power consumption by the SoC by placing at least one of the plurality of IP agents in an inoperable state.

However, these common features are previously disclosed by US 2013/0073878 A1 to JAYASIMHA, D et al. (hereinafter "JAYASIMHA"). JAYASIMHA discloses a System on a Chip "SoC" comprising an interconnect; one or more links connecting the plurality of IP agents with the interconnect (System-on-a-Chip having multiple initiator IP cores, each having its own agent to interface (link) to the remainder of the interconnect network 118; figure 1; paragraphs [0007], [0023]); and a power manager managing power consumption by the SoC (interconnect-power-manager 106 cooperating with the integrated-circuit-system-power-manager (managing power consumption by the SoC); figure 1; paragraphs [0023], [0024]) by placing at least one of the plurality of IP agents in an inoperable state (sending signals to quiesce (an inoperable state) the interconnect network components including initiator and target agents (IP agents); paragraphs [0024], [0028]).

Since the common technical features are previously disclosed by the JAYASIMHA reference, these common features are not special and so Groups I-III lack unity.