



US007212423B2

(12) **United States Patent**
Vogt

(10) **Patent No.:** **US 7,212,423 B2**
(45) **Date of Patent:** **May 1, 2007**

(54) **MEMORY AGENT CORE CLOCK ALIGNED TO LANE**

(75) Inventor: **Pete D. Vogt**, Boulder, CO (US)

(73) Assignee: **Intel Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 334 days.

(21) Appl. No.: **10/858,850**

(22) Filed: **May 31, 2004**

(65) **Prior Publication Data**

US 2005/0276150 A1 Dec. 15, 2005

(51) **Int. Cl.**
G11C 5/00 (2006.01)

(52) **U.S. Cl.** **365/52; 365/63; 365/233**

(58) **Field of Classification Search** 365/51,
365/52, 63, 233; 327/141, 147, 156; 711/5,
711/127, 143, 167

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,742,840 A	4/1998	Hansen et al.
5,748,872 A	5/1998	Norman
5,860,080 A	1/1999	James et al.
5,867,422 A	2/1999	John
5,898,863 A	4/1999	Ofer et al.
6,006,318 A	12/1999	Hansen et al.
6,034,878 A	3/2000	Osaka et al.
6,038,682 A	3/2000	Norman
6,092,229 A	7/2000	Boyle et al.
6,125,419 A	9/2000	Umemura et al.
6,128,750 A	10/2000	Espy et al.
6,151,648 A	11/2000	Haq
6,154,826 A	11/2000	Wulf et al.
6,154,855 A	11/2000	Norman
6,160,423 A	12/2000	Haq

6,185,644 B1	2/2001	Farmwald et al.
6,255,859 B1	7/2001	Haq
6,263,413 B1	7/2001	Motomura et al.
6,317,352 B1	11/2001	Halbert et al.
6,327,205 B1	12/2001	Haq
6,345,321 B1	2/2002	Litaize et al.
6,369,605 B1	4/2002	Bonella et al.
6,408,402 B1	6/2002	Norman
6,449,213 B1	9/2002	Dodd et al.

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 99/30240 6/1999

(Continued)

OTHER PUBLICATIONS

The Institute of Electrical and Electronics Engineers, Inc., *IEEE Standard for Scalable Coherent Interface (SCI)*, May 23, 2001, 1-243 pp.

(Continued)


Primary Examiner—Gene N. Auduong

(74) *Attorney, Agent, or Firm*—Marger Johnson McCollom, PC

(57) **ABSTRACT**

Memory apparatus and methods align a core clock for a memory agent to one of a plurality of lanes. A memory agent may have logic circuit between the lanes and a core clock generator to align the core clock to one of the lanes. A deskew circuit may be coupled to the logic circuit. Other embodiments are described and claimed.

25 Claims, 21 Drawing Sheets



U.S. PATENT DOCUMENTS

6,487,102	B1	11/2002	Halbert et al.	
6,493,250	B2	12/2002	Halbert et al.	
6,502,161	B1	12/2002	Perego et al.	
6,513,080	B1	1/2003	Haq	
6,625,687	B1	9/2003	Halbert et al.	
6,643,752	B1	11/2003	Donnelly et al.	
6,903,582	B2 *	6/2005	Gaskins et al.	327/141
7,020,757	B2 *	3/2006	Ruhovets et al.	711/167
2004/0230718	A1	11/2004	Polzin et al.	

FOREIGN PATENT DOCUMENTS

WO	WO 99/41666	8/1999
WO	WO 99/41667	8/1999

OTHER PUBLICATIONS

R. Ng (Sun Microsystems, Inc.), *Fast Computer Memories*, IEEE Spectrum, Oct. 1992, pp. 36-39.
R.H.W. Salters (Philips Research Laboratories), *Fast DRAMs for Sharper TV*, IEEE Spectrum, Oct. 1992, pp. 40-42.
F. Jones (United Memories, Inc.), *A New Era of Fast Dynamic RAMs*, IEEE Spectrum, Oct. 1992, pp. 43-49.
M. Farmwald and D. Mooring (Rambus, Inc.), *A Fast Path to One Memory*, IEEE Spectrum, Oct. 1992, pp. 50-51.

S. Gjessing (University of Oslo), D.B. Gustavson (Stanford Linear Accelerator Center), D.V. James and G. Stone (Apple Computer, Inc.) and H. Wigger (Hewlett-Packard Co.), *A RAM Link for High Speed*, IEEE Spectrum, Oct. 1992, pp. 52-53.

The Institute of Electrical and Electronics Engineers, Inc., *IEEE Standard for High-Bandwidth Memory Interface Based on Scalable Coherent Interface (SCI) Signaling Technology (RamLink)*, 1996, 1-91 pp.

U.S. Appl. No. 10/454,399, filed Jun. 3, 2003, Pete D. Vogt.

U.S. Appl. No. 10/454,400, filed Jun. 3, 2003, Pete D. Vogt and James W. Alexander.

U.S. Appl. No. 10/454,398, filed Jun. 3, 2003, Pete D. Vogt.

U.S. Appl. No. 10/456,206, filed Jun. 4, 2003, Pete D. Vogt.

U.S. Appl. No. 10/456,178, filed Jun. 4, 2003, Pete D. Vogt.

U.S. Appl. No. 10/456,174, filed Jun. 4, 2003, Pete D. Vogt.

U.S. Appl. No. 10/713,868, filed Nov. 14, 2003, Pete D. Vogt.

U.S. Appl. No. 10/714,026, filed Nov. 14, 2003, Pete D. Vogt.

U.S. Appl. No. 10/883,474, filed Jun. 30, 2004, Pete D. Vogt.

U.S. Appl. No. 10/882,999, filed Jun. 30, 2004, Pete D. Vogt.

U.S. Appl. No. 10/859,438, filed May 31, 2004, Pete D. Vogt.

U.S. Appl. No. 10/858,850, filed May 31, 2004, Pete D. Vogt.

U.S. Appl. No. 10/859,060, filed May 31, 2004, Pete D. Vogt.

* cited by examiner




FIG. 1
(PRIOR ART)




FIG. 2
(PRIOR ART)




FIG. 3




FIG. 4




FIG. 5




FIG. 6





FIG.7






FIG. 9




FIG. 10




FIG. 11




FIG.12




FIG.13




FIG.14




FIG.15




FIG.16




FIG.17




FIG.18




FIG.19

BIT TRANSFER	12	11	10	9	8	7	6	5	4	3	2	1	0
	0	0	0	0	0	0	0	0	0	0	0	0	1
1	0	0	0	0	0	0	0	0	0	0	0	0	1
2	0	0	0	0	0	0	0	0	0	0	0	0	1
3	0	0	0	0	0	0	0	0	0	0	0	0	1
4	0	0	0	0	0	0	0	0	0	0	0	0	1
5	0	0	0	0	0	0	0	0	0	0	0	0	1
6	0	0	0	0	0	0	0	0	0	0	0	0	1
7	0	0	0	0	0	0	0	0	0	0	0	0	1
8	0	0	0	0	0	0	0	0	0	0	0	0	1
9	0	0	0	0	0	0	0	0	0	0	0	0	1
10	0	0	0	0	0	0	0	0	0	0	0	0	1
11	0	0	0	0	0	0	0	0	0	0	0	0	1

FIG.20

TRANSFER \ BIT	12	11	10	9	8	7	6	5	4	3	2	1	0
0	1	1	0	0	0	0	0	0	0	0	0	0	0
1	1	1	0	0	0	0	0	0	0	0	0	0	0
2	1	1	0	0	0	0	0	0	0	0	0	0	0
3	1	1	0	0	0	0	0	0	0	0	0	0	0
4	1	1	0	0	0	0	0	0	0	0	0	0	0
5	1	1	0	0	0	0	0	0	0	0	0	0	0
6	1	1	0	0	0	0	0	0	0	0	0	0	0
7	1	1	0	0	0	0	0	0	0	0	0	0	0
8	1	1	0	0	0	0	0	0	0	0	0	0	0
9	1	1	0	0	0	0	0	0	0	0	0	0	0
10	1	1	0	0	0	0	0	0	0	0	0	0	0
11	1	1	0	0	0	0	0	0	0	0	0	0	0

FIG.21

TRANSFER \ BIT	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	1	0	0	0	0	0	0	0	0	0	0
1	0	0	1	0	0	0	0	0	0	0	0	0	0
2	0	0	1	0	0	0	0	0	0	0	0	0	0
3	0	0	1	0	0	0	0	0	0	0	0	0	0
4	0	0	1	0	0	0	0	0	0	0	0	0	0
5	0	0	1	0	0	0	0	0	0	0	0	0	0
6	0	0	1	0	0	0	0	0	0	0	0	0	0
7	0	0	1	0	0	0	0	0	0	0	0	0	0
8	0	0	1	0	0	0	0	0	0	0	0	0	0
9	0	0	1	0	0	0	0	0	0	0	0	0	0
10	0	0	1	0	0	0	0	0	0	0	0	0	0
11	0	0	1	0	0	0	0	0	0	0	0	0	0

FIG.22

TRANSFER \ BIT	12	11	10	9	8	7	6	5	4	3	2	1	0
0	0	0	0	1	0	0	0	0	0	0	0	0	0
1	0	0	0	1	0	0	0	0	0	0	0	0	0
2	0	0	0	1	0	0	0	0	0	0	0	0	0
3	0	0	0	1	0	0	0	0	0	0	0	0	0
4	0	0	0	1	0	0	0	0	0	0	0	0	0
5	0	0	0	1	0	0	0	0	0	0	0	0	0
6	0	0	0	1	0	0	0	0	0	0	0	0	0
7	0	0	0	1	0	0	0	0	0	0	0	0	0
8	0	0	0	1	0	0	0	0	0	0	0	0	0
9	0	0	0	1	0	0	0	0	0	0	0	0	0
10	0	0	0	1	0	0	0	0	0	0	0	0	0
11	0	0	0	1	0	0	0	0	0	0	0	0	0

FIG.23







FIG.24

**FIG.25**






FIG.28

TRANSFER \ BIT	9	8	7	6	5	4	3	2	1	0
0										
1										
2										
3										
4										
5										
6										
7										
8										
9										
10										
11										

FIG.29

TRANSFER \ BIT	9	8	7	6	5	4	3	2	1	0
0	aE0	aE7	aE8	F0	aC20	aC16	aC12	aC8	aC4	aC0
1	aE1	aE6	aE9	F1	aC21	aC17	aC13	aC9	aC5	aC1
2	aE2	aE5	aE10	aE13	aC22	aC18	aC14	aC10	aC6	aC2
3	aE3	aE4	aE11	aE12	aC23	aC19	aC15	aC11	aC7	aC3
4	FE21	B8	B7	B6	B5	B4	B3	B2	B1	B0
5	FE20	B9	B10	B11	B12	B13	B14	B15	B16	B17
6	FE19	B26	B25	B24	B23	B22	B21	B20	B19	B18
7	FE18	B27	B28	B29	B30	B31	B32	B33	B34	B35
8	FE17	B44	B43	B42	B41	B40	B39	B38	B37	B36
9	FE16	B45	B46	B47	B48	B49	B50	B51	B52	B53
10	FE15	B62	B61	B60	B59	B58	B57	B56	B55	B54
11	FE14	B63	B64	B65	B66	B67	B68	B69	B70	B71

FE0	FE7	FE8	
FE1	FE6	FE9	
FE2	FE5	FE10	FE13
FE3	FE4	FE11	FE12

FIG.31

TRANSFER \ BIT	9	8	7	6	5	4	3	2	1	0
0										
1										
2										
3										
4										
5										
6										
7										
8										
9										
10										
11										

FRAME N-1

TRANSFER \ BIT	9	8	7	6	5	4	3	2	1	0
0										
1										
2										
3										
4										
5										
6										
7										
8										
9										
10										
11										

FRAME N

FIG.30




FIG.32

TRANSFER BIT LANE	13	12	11	10	9	8	7	6	5	4	3	2	1	0	MAPPING "A"
RECEIVE BIT LANE	3	2	1	0	4	3	2	1	0	4	3	2	1	0	
RECEIVE BIT LANE	8	7	6	5	9	8	7	6	5	9	8	7	6	5	MAPPING "B"

FIG.33




FIG.34

MEMORY AGENT CORE CLOCK ALIGNED TO LANE

BACKGROUND

FIG. 1 illustrates a prior art memory system known informally as RamLink, which was proposed as a standard by the Institute of Electrical and Electronics Engineers (IEEE). The standard was designated as IEEE Std 1596.4-1996 and is known formally as *IEEE Standard for High-Bandwidth Memory Interface Based on Scalable Coherent Interface (SCI) Signaling Technology (RamLink)*. The system of FIG. 1 includes a memory controller 10 and one or more memory modules 12. The memory controller 10 is typically either built into a processor or fabricated on a companion chipset for a processor. Each memory module 12 has a slave interface 14 that has one link input and one link output. The components are arranged in a RamLink signaling topology known as RingLink with unidirectional links 16 between components. A control interface 18 on each module interfaces the slave interface 14 with memory devices 20. In the system shown in FIG. 1, another RamLink signaling topology known as SyncLink is used between the slave interfaces and memory devices.

The purpose of the RamLink system is to provide a processor with high-speed access to the memory devices. Data is transferred between the memory controller and modules in packets that circulate along the RingLink. The controller is responsible for generating all request packets and scheduling the return of slave response packets.

A write transaction is initiated when the controller sends a request packet including command, address, time, and data to a particular module. The packet is passed from module to module until it reaches the intended slave, which then passes the data to one of the memory devices for storage. The slave then sends a response packet, which is passed from module to module until it reaches the controller to confirm that the write transaction was completed.

A read transaction is initiated when the controller sends a request packet including command, address, and time to a module. The slave on that module retrieves the requested data from one of the memory devices and returns it to the controller in a response packet, which is again passed from module to module until it reaches the controller.

FIG. 2 illustrates a prior art RamLink slave interface circuit. In the circuit of FIG. 2, source-synchronous strobing is used to clock the incoming data signals. That is, a strobe signal, which accompanies the incoming data signals, is used to sample the incoming data. The circuit of FIG. 2 uses a phase-locked loop (PLL) to generate a stable local clock signal from a reference clock signal that is distributed to other slave interface circuits. The local clock signal is used to reclock the outgoing data signal so as to avoid cumulative jitter as the data is passed along downstream.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a prior art RamLink memory system.

FIG. 2 illustrates a prior art RamLink slave interface circuit.

FIG. 3 illustrates an embodiment of a memory interface system according to the inventive principles of this disclosure.

FIG. 4 illustrates an embodiment of a memory module according to the inventive principles of this disclosure.

FIG. 5 illustrates another embodiment of a memory module and an embodiment of a memory buffer according to the inventive principles of this disclosure.

FIG. 6 illustrates additional example embodiments of a memory system, memory modules, and memory buffers according to the inventive principles of this disclosure.

FIG. 7 illustrates another example embodiment of a memory buffer according to the inventive principles of this disclosure.

FIG. 8 illustrates an embodiment of a redrive circuit according to the inventive principles of this disclosure.

FIG. 9 illustrates an embodiment of an I/O cell according to the inventive principles of this disclosure.

FIG. 10 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure.

FIG. 11 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure.

FIG. 12 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure.

FIG. 13 illustrates an embodiment of a fail-over circuit according to the inventive principles of this disclosure.

FIG. 14 illustrates another embodiment of a fail-over circuit operating in a normal mode according to the inventive principles of this disclosure.

FIG. 15 illustrates another embodiment of a fail-over circuit operating in a fail-over mode according to the inventive principles of this disclosure.

FIG. 16 illustrates an embodiment of a memory buffer with bit lane fail-over capability according to the inventive principles of this disclosure.

FIG. 17 illustrates an embodiment of a memory controller with bit lane fail-over capability according to the inventive principles of this disclosure.

FIG. 18 illustrates an embodiment of a method for implementing permuting status patterns according to the inventive principles of this disclosure.

FIG. 19 illustrates an embodiment of a permuting pattern generator in accordance with the inventive principles of this disclosure.

FIGS. 20 through 23 illustrate example embodiments of status patterns in accordance with the inventive principles of this disclosure.

FIG. 24 illustrates an embodiment of a memory agent according to the inventive principles of this disclosure.

FIG. 25 illustrates an embodiment of a polling operation according to the inventive principles of this disclosure.

FIG. 26 illustrates an embodiment of a memory module utilizing data accumulation according to the inventive principles of this disclosure.

FIG. 27 illustrates another embodiment of a memory module and an embodiment of a memory buffer utilizing data accumulation according to the inventive principles of this disclosure.

FIG. 28 illustrates another example embodiment of a memory buffer utilizing data accumulation according to the inventive principles of this disclosure.

FIG. 29 illustrates an example embodiment of a frame according to the inventive principles of this disclosure.

FIG. 30 illustrates an example embodiment of a scheme for delivering a check code across multiple frames according to the inventive principles of this disclosure.

FIG. 31 illustrates an embodiment of a frame transfer scheme incorporating early check code delivery and distributed check codes according to various inventive principles of this disclosure.

3

FIG. 32 illustrates an embodiment of a memory agent, in this example assumed to be a memory module or buffer, according to the inventive principles of this disclosure.

FIG. 33 illustrates example embodiments of two possible bit lane mappings according to the inventive principles of this disclosure.

FIG. 34 illustrates an example embodiment of a memory agent according to the inventive principles of this disclosure.

DETAILED DESCRIPTION

This disclosure encompasses numerous inventions that have independent utility. In some cases, additional benefits are realized when some of the principles are utilized in various combinations with one another, thus giving rise to additional inventions. These principles may be realized in myriad embodiments. Although some specific details are shown for purposes of illustrating the inventive principles, numerous other arrangements may be devised in accordance with the inventive principles of this disclosure. Thus, the inventive principles are not limited to the specific details disclosed herein.

FIG. 3 illustrates an embodiment of a memory interface system according to the inventive principles of this disclosure. The system of FIG. 3 includes a memory controller 50 and one or more memory modules 52 that communicate through a channel made up of unidirectional links. The channel has an outbound path that includes one or more outbound links 54, and an inbound path that includes one or more inbound links 56. Each module may be capable of redriving signals from link to link on the outbound path and from link to link on the inbound path. Each module may also be capable of selectively disabling any redrive features, for example, if the module detects that it is the outermost module, or responsive to a command from the memory controller.

Each module includes one or more memory devices 58 arranged to transfer data to and/or from one or more of the paths. For example, the module may be arranged such that data from the outbound path is transferred to a memory device, and data from the memory device is transferred to the inbound path. One or more buffers may be disposed between one or more memory devices and one or more of the paths. The modules and controller are not limited to any particular mechanical arrangement. For example, the modules may be fabricated on substrates separate from the rest of the system, they may be fabricated on a common substrate with the controller and links, or they may be realized in any other mechanical arrangement. The modules are also not limited to any particular type of memory device, e.g., read only memory (ROM), dynamic random access memory (DRAM), flash memory, etc.

FIG. 4 illustrates an embodiment of a memory module according to the inventive principles of this disclosure. The module of FIG. 4 includes two redrive circuits 60 and 62 to receive signals on unidirectional links 54A and 56A, and redrive the signals on unidirectional links 54B and 56B, respectively. One or more memory devices 58 are arranged to transfer data to and/or from one or more of the redrive circuits.

The module of FIG. 4 is not limited to any particular arrangement of unidirectional links or any particular arrangement for transferring data to and/or from the redrive circuits. If the module of FIG. 4 is to be used in a memory system such as that shown in FIG. 3, then redrive circuit 60 might be designated as an outbound redrive circuit and arranged to receive and redrive signals on an outbound path

4

including links 54A and 54B, and the other redrive circuit 62 might be designated as an inbound redrive circuit and arranged to receive and redrive signals on an inbound path including links 56A and 56B. In this example, one or more memory devices 58 may be arranged so that data is transferred from the outbound redrive circuit 60 to the memory device(s) and from the memory device(s) to the inbound redrive circuit 62.

The module may be capable of detecting if it is the outermost module on a channel and selectively disabling any redrive features accordingly. For example, if the module of FIG. 4 is to be used in a memory system such as that shown in FIG. 3, and the module detects that it is the outermost module, the outbound redrive circuit receives incoming signals on link 54A but does not redrive them. Likewise, the inbound redrive circuit only drives link 56B with signals corresponding to data received from the memory device(s), and/or other signals that may be generated internally by the module. Alternatively, even if the module detects that it is not the outermost module, it may be constructed so that it is capable of operating as though it is the outermost module (for example, in response to a command from a memory controller), in which case it may ignore signals received on inbound link 56A, and it may not redrive signals onto outbound link 54B.

FIG. 5 illustrates another embodiment of a memory module and an embodiment of a memory buffer according to the inventive principles of this disclosure. The module of FIG. 5 includes a memory buffer 64 having two redrive circuits 60 and 62 to receive signals on unidirectional links 54A and 56A, and redrive the signals on unidirectional links 54B and 56B, respectively. The memory buffer also includes a memory interface 66 arranged to transfer data to and from one or more memory devices 58. The buffer may be capable of detecting if it is the last agent on a channel and selectively disabling any redrive features accordingly. The buffer may be capable of operating as though it is the last agent on a channel, even if it is not, for example, responsive to a command from a memory controller. As used herein, agent refers to any memory controller (also called a host), module, buffer, etc. that is interfaced to the channel.

The module and buffer of FIG. 5 are not limited to any particular arrangement of unidirectional links or any particular arrangement for transferring between the memory interface and the redrive circuits. If the module of FIG. 5 is to be used in a memory system such as that shown in FIG. 3, then redrive circuit 60 might be designated as an outbound redrive circuit and arranged to receive and redrive signals on an outbound path including links 54A and 54B, and the other redrive circuit 62 might be designated as an inbound redrive circuit and arranged to receive and redrive signals on an inbound path including links 56A and 56B. In this example, the memory interface may be configured to receive data from the outbound redrive circuit 60 and to send data to the inbound redrive circuit 62.

Various mechanical arrangements may be used to implement the memory modules and/or buffer of FIGS. 4 and 5. For example, the memory devices 58, redrive circuits 60 and 62, and buffer 64 may all be realized as separate integrated circuits mounted on a common circuit board or on separate circuit boards. Various combinations of the components may be fabricated together on a common integrated circuit, or they all might be fabricated on a single integrated circuit. The circuit board or boards, if any, may be capable of being plugged into sockets on a motherboard, fabricated integral with a motherboard, or arranged in any other way. There may not be a circuit board, for example, if the components

5

are fabricated as part of a multi-chip module. A memory buffer according to the inventive principles of this disclosure may be used to interface devices other than memory devices to a channel. For example a memory buffer according to the inventive principles of this disclosure may be used to interface an I/O controller or a bridge to a channel.

Additional embodiments of apparatus according to the inventive principles of this disclosure are described with reference to “inbound” and “outbound” paths, links, redrive circuits, etc. to facilitate an understanding of how the apparatus may be utilized in a memory system such as the embodiment shown in FIG. 3. These apparatus, however, are not limited to any particular arrangement of unidirectional links, to the particular arrangements shown for transferring data between the links and other circuitry, or to any of the implementation details shown.

FIG. 6 illustrates additional example embodiments of a memory system, memory modules, and memory buffers according to the inventive principles of this disclosure. Referring to FIG. 6, one or more memory modules 52 are based on printed circuit boards having contact fingers along both sides of one edge to create a dual inline memory module (DIMM) that may be plugged into a connector on another circuit board that holds other components of the system. An existing form-factor may be adopted for the module, for example the DIMM form-factor used for Double Data Rate II (DDR2) dynamic random access memory (DRAM) modules.

The modules are populated with memory devices 58, for example, commodity-type DRAM such as DDR2 DRAM. A memory buffer 64 on each module isolates the memory devices from a channel that interfaces the modules to the memory controller 50, which is also referred to as a host. The channel has point-to-point interconnects with an outbound path that includes outbound links 54, and an inbound path that includes inbound links 56. The links may be implemented with parallel unidirectional bit lanes using low-voltage differential signals.

In the embodiments of FIG. 6, no additional signal lines are used for functions such as command, reset, initialization, and the like. Instead, these functions are encoded directly in the data sent over the channel. Alternatively, however, any number of additional signal lines may be used to implement such functions.

A reference clock signal REF CLK is generated by a clock synthesizer 76 distributed to the host and modules, maybe through a clock buffer 78. This facilitates a quasi-asynchronous clocking scheme in which locally generated clock signals are used to sample and redrive incoming data. Because a common reference clock is available at each agent, data signals may be clocked without any frequency tracking. Alternatively, a local clock signal may be generated independently of any reference clock. As another alternative, a synchronous clocking scheme such as source synchronous strobing may be used.

In one possible embodiment, the host initiates data transfers by sending data, maybe in packets or frames (terms used interchangeably here), to the innermost module on the outbound path. The innermost module receives and redrives the data to the next module on the outbound path. Each module receives and redrives the outbound data until it reaches the outermost module. Although the outermost module could attempt to redrive the data to a “nonexistent” outbound link, each module may be capable of detecting (or being instructed) that it is the outermost module and disabling any redrive circuitry to reduce unnecessary power consumption, noise, etc. In this embodiment, data transfers

6

in the direction of the host, i.e., inbound, are initiated by the outermost module. Each module receives and redrives inbound data along the inbound path until it reaches the host.

Any suitable communication protocol may be used over the physical channel. For example, the host may be designated to initiate and schedule all inbound and outbound data transfers. Alternatively, any agent may be allowed to initiate data transfers. Frames of data may be configured to carry commands, read data, write data, status information, error information, initialization data, idle patterns, etc., or any combination thereof. A protocol may be implemented such that, when the host sends a command frame to a target module along the outbound path, the target module responds by immediately sending a response frame back to the host along the inbound path. In such an embodiment, the target module does not redrive the command frame on the outbound path.

In an alternative embodiment, the target module receives the command frame and then redrives the command frame on the outbound path. When the outermost module receives the command frame, it initiates a response frame (maybe nothing more than an idle frame) on the inbound path. The target module waits until the response frame reaches its inbound receiver. The target module then merges its response into the inbound data stream, e.g., by replacing the response frame sent by the outermost module with the target module’s true response frame.

FIG. 7 illustrates another example embodiment of a memory buffer according to the inventive principles of this disclosure. The memory buffer of FIG. 7 includes an outbound redrive circuit 60 to receive and redrive signals on an outbound path including links 54A and 54B, and an inbound redrive circuit 62 to receive and redrive signals on an inbound path including links 56A and 56B. A memory interface 66 interfaces the buffer to one or more memory devices, which may be through a memory bus 68. The memory interface may include read and/or write buffers such as FIFO buffers. Data from the outbound path is coupled to the memory interface, which may be through a deskew circuit 70 which eliminates skew between bits of data if the outbound path has more than one bit lane. A pattern generator 72 may be used to generate status patterns to transmit onto the inbound path, for example, if the buffer happens to be the outermost agent on a channel, in which case, there may be no signals being received on incoming inbound link 56A. A multiplexer 74 selectively couples data from the memory interface or pattern generator to the inbound redrive circuit.

The memory interface is not limited to any particular arrangement, and it may be compatible with standard memory devices, particularly commodity memory devices such as DDR2 DRAM. The entire memory buffer may be integrated on a single integrated circuit, it may be integrated into one or more memory devices, its constituent elements may be integrated onto separate components, or any other mechanical arrangement may be employed. The embodiment shown in FIG. 7 is exemplary only, and other embodiments are possible in accordance with the inventive principles of this disclosure. For example, the embodiment of FIG. 7 is shown with unidirectional data flowing from the outbound redrive circuit to the memory interface and from the memory interface to the inbound redrive circuit. This data flow, however, may be bi-directional, and other arrangements are contemplated. Even if the embodiment of FIG. 7 is to be used in a channel system in which data for the memory interface only needs to flow as shown in FIG. 7, it may still be realized with redrive circuits having full

7

bi-directional data access as this may facilitate, for example, implementation of built-in self-test (BIST) functions, in which case a second deskew circuit for deskewing data from the inbound path may be helpful.

FIG. 8 illustrates an embodiment of a redrive circuit according to the inventive principles of this disclosure. The circuit of FIG. 8 includes one or more input/output (I/O) cells 74, each of which receives an input data signal RX that it may redrive as an output data signal TX. Alternatively, an I/O cell may substitute or merge a read data signal RDX into the output data signal. A write data signal WDX may be taken from the input data signal, either before or after it is redriven as the output data signal.

The "X" in any of the above signal names indicates that it might be one of multiple similar signals depending on the number of I/O cells in the redrive circuit. For example, a redrive circuit having nine bit lanes would have nine I/O cells with input data signals named R0, R1 . . . R8. In a redrive circuit with only a single I/O cell, the data input signal would be R0 or simply R. The term RX is used to refer generically to any or all of the input data signals.

The term "write data" is used for convenience to indicate any data being taken from the data stream traveling through the I/O cell. This does not imply, however, that write data must be directed to a memory interface or memory device. Likewise, "read data" refers to any data that is input to the I/O cell, but read data may come from any source, not just a memory device or memory interface.

Referring again to FIG. 8, a clock generator 80 generates a number of phase clock signals PCX and a transmit clock signal TC in response to a reference clock signal REF CLK. The clock generator includes a phase locked loop (PLL) 82 that generates the transmit clock TC as a multiple of the reference clock signal REF CLK, and a phase clock generator 84. In one possible embodiment, there are four phase clock signals PC0, PC1, PC2 and PC3 spaced 90 degrees apart and derived from the transmit clock TC. Each of the I/O cells may use one or more of the TC and PCX clock signals to sample and/or redrive data signals, and/or to generate additional local clock signals. In this embodiment, the phase clock and transmit clock signals are stable signals in the sense that they are not adjusted in response to the phase of any of the input data signals RX.

FIG. 9 illustrates an embodiment of an I/O cell according to the inventive principles of this disclosure. A receiver 86 is arranged to receive a data signal RX and redrive it as data signal TX in response to a sampling clock signal SC. The sampling clock signal is generated by a sampling clock generator 88, which is capable of adjusting the sampling clock signal in response to the data signal RX. A write data signal WDX may be taken from the input or the output of receiver 86. If taken from the output of the receiver as shown in FIG. 9, the sampling clock signal SC may be used as, or to derive, a strobe signal for the write data. The input to the sampling clock generator may be taken from points other than the input of the receiver as shown in FIG. 9. For example, it may be taken from the output of the receiver as well.

FIG. 10 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure. In the embodiment of FIG. 10, the sampling clock generator 88 is implemented with an interpolator 90 and a receiver tracking unit (RTU) 92. The interpolator generates the sampling clock signal by interpolating between a number of phase clock signals PCX (in this case four signals that are 90 degrees out of phase) in response to a tracking signal from the receiver tracking unit. The receiver tracking unit

8

observes the data signal RX and adjusts the tracking signal so that the sampling clock signal causes the receiver to sample and redrive the data signal at an appropriate time. Thus, the sampling clock signal may dynamically track the data signal.

In one possible embodiment, the receiver tracking unit observes transitions in the data signal RX by over sampling the data signal and adjusting the sampling clock signal to sample and redrive the data signal at the center of the data eye, i.e., at the midway point between transitions in the data signal. The sampling clock generator 88 may include a loop filter that measures several bit cells and may eventually determine that it should adjust the phase of the sampling clock signal to capture the data closer to the center of the data eye location. The input to the sampling clock generator may be taken from points other than the input of the receiver as shown in FIG. 10. For example, it may be taken from the output of the receiver as well.

An embodiment of an I/O cell according to the inventive principles of this disclosure may be used with a scheme that trains the I/O cells to dynamically track the data signal. For example, if the I/O cell of FIG. 10 is used as one of the memory modules shown in FIG. 3, the host may periodically send training frames onto the outbound path. These training frames have an edge density that is adequate to assure that the receiver tracking unit observes enough transitions in the data signal to be able to adjust the sampling clock signal. Likewise, the outermost module in FIG. 3 may periodically send training frames onto the inbound path.

FIG. 11 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure. The embodiment of FIG. 11 is similar to that of FIG. 9, but a buffer 94 is added in the data signal path. The buffer 94 may be a jitter avoidance or drift compensation buffer that compensates for voltage and temperature induced effects. The buffer resynchronizes the data signal TX to a transmit clock signal TC. The transmit clock signal is stable in the sense that its phase is not adjusted in response to the data signal the way the sample clock signal SC is.

In the embodiment of FIG. 11, the buffer is capable of operating in either a pass-through mode, or a redrive mode in response to a mode signal. In pass-through mode, the signal passes through without being sampled and redriven. In redrive mode, the signal is sampled and redriven in response to the clock signal. This enables the I/O cell to operate in different redrive modes. In one possible embodiment, the buffer operates in pass through mode if the mode signal is asserted. This is referred to as resample mode and may result in a shorter latency because the data signal is being redriven by the same clock that is used to sample the data. When the mode signal is not asserted, the buffer operates in redrive mode, so the data is resynchronized to the transmit clock. This is referred to as resync mode and may result in a longer latency but may reduce jitter. The I/O cell may be designed into a memory buffer or module that has an input for receiving the mode signal. If the memory buffer or module is to be used on a system in which there is a relatively short signal path to the next agent, the input may be asserted (or not asserted depending on polarity) to cause the I/O cell to operate in resample mode because more jitter may be tolerated on a short signal path. On the other hand, if the memory buffer or module is to be used on a system in which there is a relatively long signal path to the next agent, the input may be de-asserted to cause the I/O cell to operate in resync mode because this reduces jitter, albeit at the possible expense of longer latency. Alternatively, a regis-

tered flag may be used on the memory buffer or module, or in a redrive circuit, or in the I/O cell itself to control the mode signal.

FIG. 12 illustrates another embodiment of an I/O cell according to the inventive principles of this disclosure. In the embodiment of FIG. 12, the received and transmitted data signals RX and TX are differential signals and are shown traversing the edge of an integrated circuit die on which the I/O cell may be fabricated. The receiver 86 includes a sampling unit 96 and a termination unit 98. The sampling unit samples the incoming data signal in response to a sampling clock signal SC which is generated by interpolator 90 in response to phase clock signals from the sampling clock generator. The termination unit provides differential termination and converts the differential data signal into a single-ended signal. A jitter avoidance or drift compensation buffer 94 clocks data in response to either the sampling clock signal SC or a stable transmit clock signal TC. A multiplexer 100 selectively couples data signals from either the buffer 94 or a serializer 102 to a transmit latch 104. Read data signals RDX[0 . . . n] are received at the I/O cell at serializer 102. Another multiplexer may be disposed between buffer 94 and transmit latch 104 with one input connected to the buffer and another input connected to an output of the interpolator.

When the I/O cell needs to merge read data into the data stream, the multiplexer selects its input that is coupled to the serializer so that the transmit latch clocks the read data out of the I/O cell in response to the transmit clock signal TC. Otherwise, the multiplexer selects the data signal from the buffer which is then redriven by the transmit latch. The transmit data signal is converted back to a differential signal by transmitter 106 before being driven onto the next unidirectional link. Write data is taken from the output of the transmit latch, collected in a deserializer 108, and then routed to a deskew circuit, bit lane fail-over mechanism, or other circuitry. The deserializer may also provide a bit line clock signal BLC, which may be derived from the sample clock signal, to indicate when the write data WDX[0 . . . n] is valid.

Another inventive principle of this disclosure relates to deskewing signals separately from redrive paths. A redrive path is defined by one or more components through which a signal propagates as it is received and redriven. For example, in the embodiments of FIGS. 9 and 10, the redrive path includes receiver 86. In the embodiment of FIG. 11, the redrive path includes receiver 86 and buffer 94. In the embodiment of FIG. 12, the redrive path includes sampling unit 96, termination unit 98, buffer 94, multiplexer 100, transmit latch 104, and transmitter 106.

According to some of the inventive principles of this disclosure, a deskew circuit may be integrated into a redrive circuit such that the individual bit lanes of the deskew circuit are included in the redrive paths. Thus, the signals on the bit lanes may be deskewed in each redrive circuit as it is redriven along a path. Alternatively, however, a deskew circuit according to the inventive principles of this disclosure may be separate from the redrive paths. For example, in the embodiment of FIG. 7, a deskew circuit is shown separate not only from the redrive paths in redrive circuit 60, but also from the entire redrive circuit. Alternatively, a deskew circuit according to the inventive principles of this disclosure may be integrated into the redrive circuit, but still be separate from the redrive paths. For example, in the embodiment of FIG. 12, one or more deskew latches may be located at the output of serializer 102 and/or the input of deserializer 108.

The embodiments of methods and apparatus for deskewing signals separately from redrive paths as described above are exemplary only and are not limited to these specific examples. Moreover, the principles relating to deskewing signals separately from redrive paths according to this disclosure are independent of other inventive principles of this disclosure. For example, just as the embodiments of redrive circuits illustrated in FIGS. 9–12 are not limited to use in memory systems having separate outbound and inbound paths, so too may the principles relating to deskewing signals separately from redrive paths according to this disclosure may be used with other types of memory architectures utilizing unidirectional links, e.g., an architecture that utilizes a ring-type arrangement of links such as Ram-Link.

Another inventive principle of this disclosure relates to coping with failed bit lanes. For example, any of the unidirectional links between any of the agents shown in the embodiments of FIGS. 3, 4, 5, 6 or 7 may have more than one bit lane. According to the inventive principles of this disclosure, one or more signals may be redirected on the bit lanes to avoid a bad bit lane. Any agent such as a memory controller (host), module, buffer, etc. may be capable of redirecting one or more signals on a number of bit lanes. A signal may be redirected at either or both ends of a link. Any agent may be capable of detecting a failed bit lane either automatically or with assistance from another agent, and any agent may be capable of redirecting signals responsive to a command from another agent.

FIG. 13 illustrates an embodiment of a fail-over circuit according to the inventive principles of this disclosure. The fail-over circuit 110 of FIG. 13 is shown along with an embodiment of a redrive circuit 112 for purposes of illustration only, but the inventive principles are not limited to use with any particular redrive circuit, nor is the fail-over circuit limited to the specific details shown in FIG. 13. Redrive circuit 112 includes a number of bit lanes arranged to receive and redrive signals on unidirectional links. Each bit lane is embodied as an I/O cell 114 having a receiver 116 and a transmitter 118.

A fail-over circuit refers to a circuit that is capable of redirecting one or more signals to or from a plurality of bit lanes. In the embodiment of FIG. 13, the fail-over circuit is implemented as a multiplexer having one or more multiplexer switches 120. Each switch has a first input coupled to one bit lane and a second input coupled to an adjacent bit lane so that it may redirect signals from either bit lane to its output. The embodiment shown in FIG. 13 is shown with six switches to service six bit lanes, but any number switches and bit lanes may be used, and the switches may be arranged in various configurations other than the adjacent bit lane configuration as shown.

During a normal mode of operation, each of the switches directs the signal from its first input to its output as shown in FIG. 14 so that write data signals WD0, WD1, WD2, WD3, WD4, and WD5 are directed to outputs OUT0, OUT1, OUT2, OUT3, OUT4, and OUT5, respectively. In such an embodiment, one of the bit lanes, for example, the bit lane corresponding to WD5, may be used for error checking the data on the other bit lanes.

If a bad bit lane is detected, the multiplexer may operate in a fail-over mode in which one or more of the switches are manipulated to map out the bad bit lane. For example, if the bit lane associated with WD3 does not operate properly, the multiplexer switches may redirect write data signals WD4 and WD5 to outputs OUT3 and OUT4, respectively as shown in FIG. 15. In this mode, one bit lane worth of signal

11

capacity is lost. If one of the bit lanes had been designated for error checking, signals originally intended for the bad bit lane may be rerouted over the error checking lane, and the error checking function may be disabled.

The outputs of the fail-over circuit may be coupled to a memory interface, to a memory device, or to other circuitry. In the embodiment of FIG. 13, the fail-over circuit is shown separate from the redrive circuit, but it may also be integrated into the redrive circuit. A fail-over circuit according to the inventive principles of this disclosure may be realized with simple multiplexers as shown, but other arrangements such as a full crossbar switch are also possible.

The embodiment of a fail-over circuit shown in FIG. 13 is arranged to couple write data from the bit lanes to its outputs. Alternatively, an embodiment of a fail-over circuit according to the inventive principles of this disclosure may be arranged to transfer data in the opposite direction, in which case the outputs OUTX would become inputs that receive read data, the multiplexer switches may be referred to as demultiplexer switches, and each of the I/O cells may have a multiplexer between the receiver and transmitter to merge the read data from the fail-over circuit into the bit lane. Thus a multiplexer refers to both a multiplexer and a demultiplexer. As another alternative, an embodiment of a fail-over circuit according to the inventive principles of this disclosure may be arranged for bi-directional data flow between the bit lanes and memory device, memory interface, or other circuitry.

A memory buffer, memory module, memory controller (host), or other agent having bit lane fail-over capability may also have various capabilities for detecting failed bit lanes, redirecting signals, mapping out bad bit lanes, and the like according to the inventive principles of this disclosure. For example, an agent having the embodiment of a fail-over circuit shown in FIG. 13 may be designed so that it can detect a failed bit lane, e.g., by observing a test data pattern sent by another agent, and redirecting signals to map-out the failed bit lane. Alternatively, the agent may be designed so that it may map out a failed bit lane in response to a command from another agent, for example, a memory controller that instructs one or more agents on a memory channel. Alternatively, the agent may have both capabilities.

FIG. 16 illustrates an embodiment of a memory buffer with bit lane fail-over capability according to the inventive principles of this disclosure. The embodiment of FIG. 16 is similar to that of FIG. 7 but also includes a fail-over circuit 122 coupled between the deskew circuit 70 and the memory interface 66. Alternative embodiments are possible. For example, the fail-over circuit may be disposed between the redrive circuit 60 and the deskew circuit, or it may be integrated into the redrive circuit. The embodiment of FIG. 16 also includes another fail-over circuit 124 which is shown coupled between the multiplexer 74 and redrive circuit 62, but which may also be integrated into the redrive circuit or arranged in other ways. The memory buffer of FIG. 16 may alternatively be embodied as a memory module, in which case the memory interface is replaced by a memory device.

FIG. 17 illustrates an embodiment of a memory controller with bit lane fail-over capability according to the inventive principles of this disclosure. The controller of FIG. 17 includes outbound and inbound unidirectional link interfaces 126 and 128 having a plurality of bits lanes which, in this embodiment, include a number of transmitters and a number of receivers, respectively. Fail-over circuits 130 and 132 are coupled to the bit lanes in the interfaces 126 and 128, respectively. In the embodiment of FIG. 17, the fail-over circuits are shown separate from the link interfaces, but they

12

may alternatively be integral with the interfaces. The controller may be capable of detecting a failed bit lane, in which case the fail-over circuits may map out the failed bit lane. Additionally or alternatively, the controller may be capable of issuing a command that directs an agent to map out a failed bit lane.

Additional fail-over methods and apparatus according to the inventive principles of this disclosure will now be described in the context of an exemplary embodiment of a complete memory channel system including additional embodiments of a memory controller (host), memory modules, and memory buffers according to the inventive principles of this disclosure. None of the components, however, are limited to this exemplary system or any of the details described therein.

The exemplary system includes an embodiment of a host having fail-over capabilities such as those described with reference to FIG. 17 and embodiments of one or more memory modules having buffers with fail-over capabilities such as those described with reference to FIG. 16. In this example, the host and modules are arranged in a channel configuration having outbound and inbound paths such as that shown in FIG. 7, although the system may only include one module.

In this example, the host and modules are interconnected with a system management bus known as "SMBus", which is a serial bus system used to manage components in a system. However, the use of SMBus is not necessary to the inventive principles of this disclosure, and other forms of communication between components may be used, including the memory channel paths themselves.

An embodiment of a method according to the inventive principles of this disclosure for detecting and mapping out a failed bit lane in the exemplary system proceeds as follows. The host transmits a test pattern on each bit lane of the outbound path. The test pattern is received and redriven along the outbound path by the buffer on each module until it reaches the outermost module. The outermost module then transmits a test pattern on each bit lane of the inbound path. The test pattern is received and redriven along the inbound path by the buffer on each module until it reaches the host. The host and the buffers on the modules observe the test pattern on each bit lane of the inbound and/or outbound paths to check for proper bit lane operation. The bit lanes in the inbound and outbound paths may be tested concurrently.

Failed bit lanes are reported by sending results to the host through the SMBus and/or by transmitting a results frame over the channel to the host. Such a results frame may be initiated on the inbound path by the outermost module, and the other modules, if any, may merge their results information into the data in the inbound path. If the results from each module are transmitted redundantly on more than one bit lane, a failed bit lane is unlikely to interfere with reporting the results.

Once the host receives the results, it may issue a configuration command to the modules, through the SMBus, over the channel, or through any other form of communication. The configuration command instructs the modules which, if any, bit lanes are bad and should be mapped out. The modules respond to the configuration command by manipulating one or more fail-over circuits to redirect signals around bad bit lanes, if any, and reconfiguring any internal functionality to accommodate the loss of a bit lane. For example, if one bit lane was designated for error checking data, the buffer or module may disable error checking functions.

13

The embodiments of fail-over methods and apparatus described above are exemplary only, and the inventive principles of this disclosure are not limited to these specific examples. The principles of fail-over methods and apparatus according to this disclosure have been described with reference to a memory system having separate inbound and outbound paths such as the embodiment of FIG. 3, but the principles may also be applied to any memory architecture utilizing unidirectional links, for example an architecture that utilizes a ring-type arrangement of links such as Ram-Link.

Another inventive principle of this disclosure relates to permuting status patterns. In memory systems such as those described above with reference to FIGS. 1 and 3 where memory read and write data is transferred between memory agents, it may also be useful to send status information such as idle patterns, alert patterns, and other status information between memory agents. This may be accomplished by sending data patterns and status patterns on the same link or links that connect the memory agents. According to the inventive principles of this disclosure, the status patterns may be permuted over time.

For example, referring to FIG. 3, the memory controller 50 may send frames having data patterns such as read commands to one or more of the modules 52 which respond by sending frames having data patterns such as read data back to the controller. It may be useful for the one or more of the modules to send a frame having an idle pattern back to the memory controller, for example, if the module was not able to retrieve read data from a memory device 58 fast enough. A predetermined data pattern may be designated as an idle pattern so that, if the memory controller receives the idle pattern, it knows it is not receiving read data. This may cause problems, however, if the actual read data pattern happens to match the designated idle pattern.

According to the inventive principles of this disclosure, the memory controller and one or more modules may both be capable of permuting the idle pattern in a predictable manner so that the idle pattern changes over time. For example, the memory controller and modules may change the idle pattern according to a predetermined sequence each time an idle frame is sent and/or received. An embodiment of such a method according to the inventive principles of this disclosure is illustrated in FIG. 18. Thus, if the memory controller sends a read command frame (158) and receives a response frame (160) having the current idle pattern (162), it may resend the same read command (164). If the second response frame (166) contains the same pattern as the first (168), it interprets the pattern as actual read data (170). If, however, the pattern in the second response frame matches the permuted idle pattern (168), the memory controller knows that the first response frame was an idle frame (172).

According to the inventive principles of this disclosure, the status information sent in status patterns may be idle patterns, alert patterns, and other status information such as command error information from a module, thermal overload information from a module, and information that indicates that a module has detected the presence of another module on the outbound path of memory channel. Some types of status patterns may be implemented as complementary patterns. For example, an alert pattern (which may be used to notify an agent of an error condition) may be implemented as the logical complement of an idle pattern. This may simplify the implementation by, for example, allowing a memory agent to use the same pattern generator

14

for idle and alert patterns. The use of complementary status patterns may be beneficial even if permuting patterns are not used.

A memory agent according to the inventive principles of this disclosure may also be capable of intentionally generating an error such as a check code error in a status pattern. Such a technique may be useful as an alternative or supplemental way to distinguish a data pattern from a status pattern. For example, in some memory systems, each frame is sent along with a check code that used to check the integrity of the data in the frame. According to the inventive principles of this disclosure, a memory agent may intentionally send the wrong check code with frame that contains a status pattern. The receiving agent may then interpret the frame as a status frame rather than a data frame. Some memory systems may utilize a path or paths having an extra bit lane to carry check code data. If such a system is capable of operating in a fail-over mode, the agent or agents may only utilize an intentional check code error if not operating in fail-over mode. A cyclical redundancy check code (CRC) code is one example of a type of check code suitable for use as a check code according to the inventive principles of this disclosure, but check codes of any other type of error checking scheme may also be used to verify the integrity of a frame or pattern according to the inventive principles of this disclosure.

Although the principles of status pattern permuting and handling according to the inventive principles of this disclosure are applicable to any type of memory agent, and are independent of other inventive principles of this disclosure, some additional aspects will be described with respect to a memory buffer such as the embodiment shown in FIG. 7 and in the context of a system such as the embodiment shown in FIG. 6. Referring to FIG. 6, if the memory buffer 64 is the outermost agent on a memory channel, it may be capable of constantly transmitting permuting idle status frames on the inbound link 56B whenever it is not sending data that the host has requested from any memory devices attached to the memory interface 66.

FIG. 19 illustrates an embodiment of a permuting pattern generator in accordance with the inventive principles of this disclosure. The embodiment of FIG. 19 is a 12-bit linear-feedback shift register (LFSR) with a polynomial of $x^{12} + x^7 + x^4 + x^3 + 1$. The initial state may be set to 000000000001, and the LFSR cycles through $2^{12}-1$ states (4095 frames) before the pattern is repeated. Each bit of the LFSR may be mapped to a bit lane in a link on a data path, and each bit may be used for all of the transfers that occur on the corresponding bit lane during an entire frame. For example, in a system having a data path with 12 bit lanes in each link, the output from each stage of the LFSR may be mapped to one of the bit lanes. Additional lanes, for example, a 13th bit lane, may be accommodated by utilizing the value from the least significant bit of the LFSR delayed by one frame.

FIG. 20 illustrates an example of the first status pattern generated by the permuting pattern generator of FIG. 19. In this example, a frame is 12 transfers long. FIGS. 21–22 illustrate the second, third and fourth status patterns, respectively. By using the same value on each bit lane during an entire frame, electromagnetic interference (EMI or noise) may be reduced.

The 13 bit lane by 12 bit transfer frame illustrated here is by way of example, and the inventive principles of this disclosure are not limited to these details, nor to the specific embodiment of a permuting pattern generator described above. For example, a permuting pattern generator according to the inventive principles of this disclosure need not be

15

implemented with dedicated logic circuitry such as the LFSR described above. Alternatively it may be implemented with programmable logic, or as an algorithm in a processor or other programmable state machine that may be used to oversee and/or implement the logic in the memory interface or other functionality of a buffer or other memory agent that utilizes permuting status patterns.

Another inventive principle of this disclosure relates to utilizing more than one bit lane to detect the presence of a memory agent on a memory link. For example, in the embodiment of a memory buffer shown in FIG. 7, the buffer may be capable of detecting whether there is another memory agent coupled to the outbound link 54B. This may be accomplished by utilizing a single bit lane in the link to test for the presence of another memory agent. If there is more than one bit lane in the link, however, more than one of the bit lanes may be used to detect the presence of another memory agent according to the inventive principles of this disclosure. This may prevent the existence of a bad bit lane from interfering with the presence detect operation.

For convenience, the inventive principles of this disclosure relating to utilizing more than one bit lane to detect the presence of a memory agent will be referred to individually and collectively as redundant presence detect. Redundant presence detect may be applied to any type of memory agent having a link interface with a plurality of bit lanes. For example, any two or more of the transmitters 118 shown in the embodiment of FIG. 13 may be considered a link interface, in this case a transmit link interface. Likewise, any two or more of the receivers 116 shown in FIG. 13 may be considered a link interface, in this case a receive link interface. Redundant presence detect may be applied to either of these link interfaces, as well as either of the link interfaces 126 and 128 shown in the embodiment of FIG. 17.

Returning to the embodiment of FIG. 7 as an example again, the memory buffer may drive three bit lanes on its inbound transmit link 56B to a predetermined presence detect logic level, e.g., one, to signal its presence to another buffer when a presence detect event such as a reset occurs. Also during a presence detect event, a second such memory buffer located inbound from the first buffer on a channel may configure the corresponding three bit lanes on its inbound receive link 56A to detect the presence of the first buffer. In this example, the first memory buffer will be referred to as an outer agent, and the second buffer will be referred to as an inner agent.

An example of a technique for configuring a bit lane to detect the presence of another agent is to have the receiver for that bit lane try to place a bias current on the bit lane so as to force the bit lane to the opposite of the presence detect logic level. If another memory agent is coupled to the bit lane during a presence detect event, its transmitter on that bit lane will force the bit lane to the presence detect logic level.

If the inner agent detects the presence detect logic level on two of the three bit lanes, it knows that the outer agent is present and it may leave all or a portion of its outer port enabled. (In this example, the outer port includes the link interface for the outbound link 54B and the link interface for the inbound link 56A.) If the inner agent fails to detect the presence detect logic level on at least two of the three bit lanes, it may decide that an outer agent is not present and it may disable all or a portion of its outer port. The inner agent may be capable of reporting the presence or absence of an outer agent to another agent, for example to a memory controller in response to a status check command.

An agent utilizing redundant presence detect may also be capable of signaling a presence detect event to another

16

agent. For example, if a reset event is communicated to the buffer of FIG. 7 through a reset command on the outbound path, this command may be relayed to an outer agent, if any, by redrive circuit 60. This may place both agents in a presence detect mode.

Redundant presence detect according to the inventive principles of this disclosure is not limited to the specific embodiments discussed above. For example, only two bit lanes may be used for presence detect instead of three as in the example above, in which case the inner agent would only need to detect the presence detect logic level on a single bit lane to conclude that an outer agent was present. Likewise, redundant presence detect may be applied to systems and components utilizing various other types of memory architectures, e.g., an architecture that utilizes a ring-type arrangement of links such as RamLink.

Another inventive principle according to this disclosure relates to hot insertion and/or removal of components from a memory channel—that is, adding and/or removing components while the memory channel is operating. FIG. 24 illustrates an embodiment of a memory agent 134 according to the inventive principles of this disclosure. The embodiment of FIG. 24 may be a memory module, memory buffer, memory controller, etc. The agent includes a first port 136 and a second port 138. If the agent is assumed, for purposes of illustration only, to be a memory module such as one of modules 52 in the embodiment of FIG. 6, the first port may be designated as an inner port since it may be arranged to communicate with other agents on the memory channel that are located closer to the memory controller. Likewise, the second port may be designated as an outer port since it may be arranged to communicate with agents on the memory channel that are located further away from the memory controller. These designations are for purposes of illustration only, and the inventive principles are not limited to these details of the memory agent nor to the particulars of the memory channel shown in FIG. 6. These principles may also be applicable to other memory channel architectures such as the RamLink architecture shown in FIG. 1.

Each port of a memory agent according to the inventive principles of this disclosure has one or more link interfaces. In the embodiment of FIG. 24, each port has both a receive link interface and a transmit link interface. The inner port 136 has a receive link interface 140 which may be one or more receivers that are part of a redrive circuit 60, and a transmit link interface 142 which may be one or more transmitters that are part of another redrive circuit 62. The outer port has receive and transmit link interfaces 144 and 146, respectively, which are also part of redrive circuits 62 and 60, respectively. Link interfaces 140 and 146 may be coupled to outbound links 54A and 54B, respectively, and link interfaces 142 and 144 may be coupled to inbound links 56B and 56A, respectively. Each of the link interfaces may have one or more bit lanes, and the bit lanes and interfaces may be referred to using any combination of this terminology. For example, the bit lanes in interface 142 may be referred to as the inbound transmit or inbound Tx bit lanes. The bit lanes in interface 144 may be referred to as the inbound receive or inbound Rx bit lanes.

The embodiment of FIG. 24 is exemplary only, and memory agents and ports may be embodied in different ways. For example, link interfaces are not necessarily part of a redrive circuit. This is illustrated in the embodiment of a memory controller shown in FIG. 17 wherein a port may include the link interfaces 126 and 128 which are not part of redrive circuits. The link interfaces may include only one or

any number of bit lanes, and a port may only have a receive link interface or a transmit interface.

A memory agent according to the inventive principles of this disclosure may be capable of detecting the presence of another memory agent on one of its ports, and it may be capable of taking various actions depending on the presence or absence of another memory agent. For example, the memory agent of FIG. 24 may be capable of disabling all or a portion of its outer port if another memory agent is not present at the port. It may be capable of reporting the presence or absence of an outer agent to another agent, for example to a memory controller through its inner port. The memory agent of FIG. 24 may be capable of performing a presence detect operation which may include signaling a presence detect event to a potential outer agent on the outer port. It may also be capable of performing a fast reset operation.

Another inventive principle which may facilitate hot add/removal in accordance with this disclosure will be described in the context of an example embodiment of a memory system. The example embodiment will be described with reference to the memory agent of FIG. 24 in the context of a memory system such as the embodiment of FIG. 6. In this example embodiment, it will be assumed that the memory agent of FIG. 24 is used to embody one or more of the buffers in FIG. 6, which in turn are part of modules having memory devices. All of these details, however, are for purposes of explanation only, and the inventive principles are not limited to these details.

In the example system, the memory agents may be capable of executing fast reset operations, full reset operations, and/or various polling or presence detect operations. In the example system, a minimum number of clock transitions may be necessary to keep the derived clocks on each bit lane locked to the data stream. Thus, the memory controller (or host) may initiate a reset operation by sending a continuous stream of ones or zeros on one or more of the bit lanes in the outbound path for a predetermined period of time. Since the data is redriven by each buffer on the path, all of the buffers receive the reset command, or event. In the example system, the three least significant bit (LSB) lanes may be used to signal a reset operation. The receiving agent may detect the reset event by sensing the stream of zeros or ones on any two of the three LSBs. This may assure that the presence of one failed bit lane does not interfere with a reset operation, but the inventive principles, which do not even require more than one bit lane, are not limited to such an implementation.

In the example system, the host may send a continuous stream of zeros hold all of the agents on the channel (in this example modules having buffers) in a first reset state indefinitely, for example while the host is held in reset by external conditions. The host may then send a stream of ones for a first amount of time, e.g., two frame periods, and then back to zeros to signal the other agents to execute a fast reset operation. Alternatively, the host may send a stream of ones for a second amount of time, e.g., more than two frame periods, to signal the other buffers to execute a full reset operation. A full reset may include various internal calibration operations such as impedance matching on the links, current source calibration in any receiver or drive circuitry, receiver offset cancellation, and the like. After the calibration operations are performed, the host may then signal the buffers to transition to the fast reset operation.

A fast reset operation may bypass certain operations performed during a full reset such as calibration operations. A fast reset operation may begin with a presence detect

operation. During a presence detect operation, each buffer on the channel may place a current on the three LSB inbound Rx bit lanes to force the bits to zero if they are not connected to an outer agent. Also during a presence detect operation, each buffer may drive the three LSB inbound Tx bit lanes to one. Each buffer may then check its three LSB inbound Rx bit lanes, and if it detects ones on two of the three lanes, it may leave its outer port enabled and update a status register accordingly. If the buffer does not detect two ones, it may assume that there is no outer agent, disable all or a portion of its outer port, configure itself to perform the functions of the outermost agent on the channel, and/or update a status register accordingly. A host may follow a similar presence detect operation to determine if any agents are on the channel. The buffers may relay the status information to the host in status frames in response to status requests from the host.

After a presence detect operation, the buffers in the example system may transition through various other operations during a fast reset such as a clock training state to train the local clocks on the buffers to lock onto the data stream, a frame training state to align frames that are sent over the channel, bit lane tests to check the operation of all bit lanes and place the buffers in fail-over mode if they have such a capability, etc. Also, once the host knows how many other agents are connected to the channel, it may adjust the frame size, timing, etc. to accommodate all of the agents.

In the example system, the memory agents may also or alternatively be capable of performing various polling operations to detect the presence of newly added agents on the channel. For example, each buffer may be capable of performing a polling operation on its outer port if it is the outermost buffer to determine if a new agent has been added to the channel. FIG. 25 illustrates an embodiment of such a polling operation according to the inventive principles of this disclosure.

At 148, the agent may disable all or a portion of its outer port. If the agent is a buffer or module, it may wait for a poll command from the host to transition to a hot reset operation at 150. If the agent is a host, it may disable all or a portion of its outer port and wait for a wake up command from a system environment. Upon receiving the wake up command, it may turn enable all or a portion of its outer port and transition to a reset state.

At 150, the agent may enable its outer port and drive zeros onto the three LSB outbound Tx bit lanes to send a reset to a potential new agent on its outer port. The agent may then transition to a hot calibration operation at 152.

At 152, the agent may drive ones onto the three LSB outbound Tx bit lanes to force a potential new agent through a full reset including calibration operations, since a newly detected agent would presumably need to be calibrated. The agent may then transition to a hot detect operation at 154.

At 154, the agent may drive zeros onto the three LSB outbound Tx bit lanes and place a bias current on the three LSB inbound Rx bit lanes to force the bits to zero if they are not connected to an outer agent. The agent may then check the three LSB inbound Rx bit lanes, and if it detects at least two ones, it may decide at 155 that an outer agent is present and transition to a hot agent present operation at 156. Otherwise, the agent may decide at 155 that an outer agent is not present and transition back to the sleep operation at 148.

At 156, the agent may update a status register to indicate that it has detected an outer agent and then relay this information to the host, for example, in response to a status

request, or take some other action to relay the information to the host or other agent. The agent may also wait to receive a channel reset.

The host may become aware of the newly detected agent, either through periodic status requests, or other techniques and initiate a fast reset to re-initialize the entire channel with the new agent on the channel and accommodated in the channel timing.

The following are some additional embodiments of hot add/removal sequences according to the inventive principles of this disclosure. These additional embodiments are also described with reference to the embodiment of the memory system shown in FIG. 6 in the context of a larger system, such as a server having a user interface and system firmware, that employs the memory system of FIG. 6. The inventive principles illustrated by these additional embodiments, however, are not limited to the specific details described herein.

A hot add sequence according to the inventive principles of this disclosure may begin when a user appends a new agent onto the memory channel, for example on the outer port of the outermost agent. The user may inform the system firmware that an agent has been appended. The firmware may then cause power to be applied to the appended agent and inform the host through a wake up command that an agent has been appended. The host may then send a poll command to the previous outermost agent, which then may cycle through a polling operation such as the one described above with reference to FIG. 25. After the polling operation, the previous outermost agent may report the presence of a new outer agent. The host may then detect the presence of the new agent and issue a fast reset command to bring the new agent into operation and retime the entire channel. After the new agent is operational, the host may interrupt the system firmware to report that the new agent is operational. Alternatively, the host may wait for the system firmware to query the host to determine if the new agent is operational. The system firmware may then configure the host to accommodate any new hardware presented by the new agents such as new memory devices that may be present if the agent was a memory module or buffer.

A hot removal sequence according to the inventive principles of this disclosure may begin when a user informs the system that a specific agent on a memory channel is to be removed. The system may remove a corresponding host address range from a system map. If the system uses mirroring, the system may remap the host address ranges to agent mirrors. The system may then copy or move data from the host address range to other locations if not already mirrored. The system may then poll until all outstanding transactions are completed. The system may then cause the host to send a command to the agent just inside of the agent to be removed that causes this agent to assume it is the outermost agent on the channel, thereby causing it to disable its outer port and assume the functions of the outermost agent during subsequent fast resets. (A full reset would override this command.) The system may then initiate a fast reset to shut down the selected agent and any channel interfaces for components attached to the selected agent. The system may then disconnect power to the selected agent and notify the user that the agent may be removed.

A hot replace sequence according to the inventive principles of this disclosure may begin when the hot remove sequence described above is completed. The user may add a new agent in place of the one removed and then inform the system firmware that the new agent has been added. The running system may then prepare the host for the newly replaced component and supply power to the new compo-

nent. System firmware may then cause the host to send a command to the previous outermost agent to let it know that it should no longer assume that it is the outermost agent. This may cause the previous outermost agent to enable its outer port in response to the next reset, and wait for a poll command. Firmware may then instruct the host to send a poll command to the previous outermost agent which may then perform a polling operation such as the one described above with reference to FIG. 25, thereby initializing the new agent. The previous outermost agent may then report the presence of a new outer agent. The host may then detect the presence of the new agent and issue a fast reset command to bring the new agent into operation and retime the entire channel. After the new agent is operational, the host may interrupt the system firmware to report that the new agent is operational. Alternatively, the host may wait for the system firmware to query the host to determine if the new agent is operational.

Another inventive principle of this disclosure relates to accumulating data between a data path and a memory device. FIG. 26 illustrates an embodiment of a memory module utilizing data accumulation according to the inventive principles of this disclosure. The module 174 of FIG. 26 includes one or more memory devices 180 and a redrive circuit 176 to receive one or more signals on point-to-point link 178A, and then redrive the signals on point-to-point link 178B. A data accumulator 182 is disposed between the redrive circuit and a memory device. An optional second redrive circuit 184 is arranged to receive one or more signals on point-to-point link 186A and redrive the signals on point-to-point link 186B. In this embodiment, the point-to-point links are shown as unidirectional links, but the inventive principles are not limited to unidirectional links.

The data accumulator 182 may be a first-in, first-out (FIFO) data structure or any other type of suitable queue or buffer. The use of a data accumulator may allow for bandwidth mismatches. For example, a memory device having a high-bandwidth burst mode may be used for the memory device 180. The bandwidth of the data path formed from the unidirectional links may be less than the burst mode of the memory device in order to reduce pin count, power consumption, and manufacturing and operating costs. The memory device, however, may need to receive data at full bandwidth for proper operation in burst mode. By utilizing a data accumulator, write data from the data path may be accumulated at a rate less than the burst rate of the memory device, and then delivered to the memory device at its full burst rate.

The module of FIG. 26 is not limited to any particular arrangement of unidirectional links or any particular arrangement for transferring data to and/or from the redrive circuits. Data is shown flowing from the redrive circuit 176 to the one or more memory devices 180 through data accumulator 182, but the direction could be reversed, and additional data accumulators may be added between the redrive circuit and memory devices. Likewise, if the optional second redrive circuit 184 is included, data may flow may be either to or from a memory device. One or more data accumulators may be included between the second redrive circuit and the memory devices to accumulate write data to a memory device and/or read data from a memory device.

FIG. 27 illustrates another embodiment of a memory module and an embodiment of a memory buffer utilizing data accumulation according to the inventive principles of this disclosure. The module 174 of FIG. 27 includes a memory buffer 188 having two redrive circuits 176 and 184,

21

and a memory interface **190** arranged to transfer data to and from one or more memory devices **180**. As with the module of FIG. **26**, only one of the redrive circuits is needed in the buffer of FIG. **27**. The memory interface **190** includes a data accumulator **182** which may be arranged to accumulate data as it flows between the redrive circuit **176** and memory devices **180**. Data flow may be in either direction, and an additional accumulator may be included to accommodate data accumulation in both directions. Likewise, if the second redrive circuit **184** is included, one or more additional data accumulators may be included to interface with the second redrive circuit.

FIG. **28** illustrates another example embodiment of a memory buffer utilizing data accumulation according to the inventive principles of this disclosure. The buffer **64** of FIG. **28** is similar to that of FIG. **7**, but the memory interface **66** includes a FIFO-type write data accumulator **192** and a FIFO-type read data accumulator **194**. Some possible implementation details will now be described with reference to the embodiment of FIG. **28**, but the inventive principles of this disclosure are not limited to these details.

Write data from the first redrive circuit **60** is accumulated in the write FIFO at whatever data rate the outbound path happens to be operating at. Once enough write data is accumulated, it may be written to one or more memory devices at full burst rate through memory bus **68**. The read FIFO **194** may perform data capture from the memory device at full burst rate, and levelize the data prior to transferring the read data to the second redrive circuit **62** through multiplexer **74**.

The write FIFO may be constructed so that it can accumulate multiple bursts of data prior to bursting the data to a memory device. This allows the read-write-read memory bus turn around penalty to be amortized over a number of write operations. The write FIFO may also be constructed so that additional data may be loaded into the FIFO while data is being delivered to the memory device. This allows the depth of the FIFO to be smaller than the number of transfers in a burst. As a further refinement, a data pre-accumulator may be located ahead of the write FIFO and set up to speculatively capture write data from the data path without regard as to whether the data is intended for this particular memory buffer **64**. Once the target memory buffer is identified, the data in the pre-accumulator may be transferred to the write FIFO if this memory buffer was the intended recipient, otherwise it may simply be discarded.

As with the other memory modules and/or buffers disclosed in this disclosure, the embodiments of memory modules and buffers described with reference to FIGS. **26–28** are not limited to any particular mechanical arrangements, and may be used to interface devices other than memory devices to a channel.

Another inventive principle of this disclosure relates to transmitting frames with early delivery of a check code for a portion of the frame. FIG. **29** illustrates an example embodiment of a frame according to the inventive principles of this disclosure. In this example, the frame is 10 bits wide and 12 transfers long, but any other width and length may be used. The bits in the row identified as transfer “**0**” are transferred first, followed by the bits in the row identified as transfer “**1**”, etc. The frame is completed when the bits in row “**11**” are transferred. The first four rows (shown in shading) are designated arbitrarily as a first portion of the frame, while the remainder of the frame is designated as a second portion of the frame.

In prior art frame transfer schemes, a check code for error checking the entire frame is typically placed at the end of the

22

frame. According to the inventive principles of this disclosure, a check code for the first portion of a frame may be transferred before completing, or preferably even beginning, the transfer of the second portion. This early delivery of a check code allows the memory agent receiving the frame to error check the first portion of the frame, and preferably begin utilizing any information contained therein, before the second portion of the frame is completed.

For example, if this technique is used with a memory agent having apparatus that buffers memory devices such as DRAM chips from a communication channel, a DRAM command may be placed in the first portion of the frame, while a data payload may be placed in the second portion. Early delivery of the check code for the first portion allows apparatus in the memory agent receiving the frame to error check the command in the first portion of the frame and forward it to a DRAM chip before the data payload in the second portion is received.

The check code for the first portion may be included as part of the first portion of the frame. It may be placed at the end of the first portion, distributed throughout the first portion, contained only partly in the first portion, or transferred in any other suitable manner such that it is received before the end of the second portion. The check code for the first portion may be combined with other check codes to create compound codes, or may be the result of compounding with other codes both within and outside of the frame.

As used herein, the first portion need not be literally first in the frame, but may also be any portion that is received earlier than a later portion. Likewise, the second portion may be the next portion after the first, but there may also be other portions between the first and second portions or after the second portion, and the first and second portions might even overlap, so long as the effect is that a check code for the first portion may be transferred early so that error checking of the first portion may begin before the frame is completely transferred.

A second check code for the second portion of the frame may be placed at the end of the second portion, distributed throughout the second portion, contained only partly in the second portion, or transferred in any other suitable manner. The second check code may cover only the second portion of the frame, may cover the entire frame, or may be compounded with other check codes in other ways.

Another inventive principle of this disclosure relates to organizing check codes across multiple frames. FIG. **30** illustrates an example embodiment of a scheme for delivering a check code across multiple frames according to the inventive principles of this disclosure. In this example, a frame is 10 bits wide and 12 transfers long, but any other width and length may be used. The bits in the row identified as transfer “**0**” are transferred first, followed by the bits in the row identified as transfer “**1**”, etc. The frame is completed when the bits in row “**11**” are transferred. The frames in FIG. **30** are transferred sequentially with frame N-1 being transferred first, followed by frame N.

A portion of a check code for frame N-1 is shown shaded in frame N-1 arbitrarily in the position of bit “**9**” in the rows identified as transfers “**4**” through “**11**” of frame N-1. Another portion of the check code for frame N-1 is shown shaded in frame N arbitrarily in bits “**9**” through “**7**” in rows “**0**” through “**1**” and in bits “**9**” through “**6**” in rows “**2**” through “**3**”.

The check code for frame N-1 (which is shown distributed over frames N-1 and N) may be intended for error checking all of frame N-1, just a portion of frame N-1, or all or a portion of some other frame. Multiple check codes (or

23

portions of check codes) for different portions of frame N-1 may also be contained in the same frame or combined with check codes (or portions of check codes) from other frames to create compound check codes.

The inventive principles of this disclosure relating to organizing check codes across multiple frames are independent of those inventive principles relating to early delivery of check codes. These different principles may be combined, however, thereby giving rise to additional inventive principles. FIG. 31 illustrates an embodiment of a frame transfer scheme incorporating early check code delivery and distributed check codes according to various inventive principles of this disclosure. Numerous implementation details (such as the use of a CRC code for a check code) will be described by way of example, but the inventive principles are not limited to these details.

In the example embodiment of FIG. 31, each frame is again 10 bits wide and 12 transfers long, and frame N-1 is transferred first, followed by frame N. The first four rows of Frame N-1 are once again arbitrarily designated as a first portion and will be referred to as the command portion. The last four rows are arbitrarily designated as a second portion and will be referred to as the data portion.

The command portion includes 24 bits of command information in the aC[23:0] field, two additional bits of command information or frame type encoding in the F[1:0] field, and a 14-bit compound CRC checksum in the aE[13:0] field. The aC[23:0] field and the F[1:0] field will be referred to collectively as the "A" command. The aE[13:0] field provides error detection coverage across the F[1:0], aC[23:0] and aE[13:0] fields.

The data portion includes 72 bits of data in the B[71:0] field which may be used for additional commands, command extensions, data transfer, etc., and a portion of a 22-bit compound CRC checksum identified as FE[21:0]. Eight of the 22 bits are located in the FE[21:14] field in frame N-1 (the current frame). The other 14 bits are located in the FE[13:0] field which is shown outside of the frame in FIG. 31 because this portion of the checksum will be latched and compounded with a CRC code in frame N (the next frame).

To generate a frame at the transmitting memory agent, a 22-bit CRC (referred to as CRC22[21:0] or the "data CRC") is generated from the 72-bit data B[71:0]. A 14-bit CRC (referred to as CRC14[13:0] or the "command CRC") is generated from the 26-bit "A" command F[1:0]aC[23:0]. Eight bits of the 22-bit data CRC are used directly as FE[21:14] and are located in the 10th bit lane (bit lane "9") of the current frame. The remaining 14 bits of the 22-bit data CRC become FE[13:0] and are combined with the 14-bit CRC generated from the 26-bit "A" command in the next frame using a bit-wise exclusive-or (XOR) function to create the compound checksum aE[13:0] which will be transmitted in the next frame. The compound 14-bit checksum aE[13:0] in the current frame is generated by an XOR operation of the "A" command 14-bit CRC from this frame, with the latched FE[13:0] generated from the 72-bit data of the previous frame.

To decode a frame at the receiving memory agent, a 14-bit command checksum CRC14[13:0] is generated from the 26-bit command, and a 22-bit data checksum CRC22[21:0] is generated from the 72-bit data in the current frame. CRC22[13:0] is latched as FE[13:0] for future compound CRC checks in the next frame. A test compound checksum TESTaE[13:0] is generated through a bitwise XOR of FE[13:0] from the previous frame with the new aE[13:0] from the current frame. If the generated test compound checksum TESTaE[13:0] matches the compound checksum

24

aE[13:0] transmitted with the current frame, there are no errors in the "A" command of the current frame.

To complete the detection of faults in the 72-bits of data from the previous frame, the 14-bit command checksum CRC14[13:0] generated from the current 26-bit command is XORed with the new aE[13:0] from the current frame, thereby generating a result which is compared to the latched FE[13:0] from the previous frame.

To start the fault detection of the 72-bits of data transferred in the current frame, FE[21:14] transmitted with the current frame is compared with the new CRC22[21:14] generated from the 72-bit data in the current frame. The completion of fault detection for the 72-bits of data transferred in the current frame is done when the next frame arrives.

A fault in aE[13:0] indicates that both the "A" command in the current frame could be faulted, and that the 72-bit data in the previous frame could be faulted. A comparison fault in the transmitted FE[21:14] partial checksum indicates that the 72-bit data in the previous frame could be faulted.

The CRC of the "A" command may be checked as soon as the first 4 transfers of the frame are received and the information in the "A" command may be used immediately without waiting for the remainder of the frame to arrive. This mechanism may provide strong check code protection of the 72 data bits of the previous frame while reducing latency in the delivery of the "A" command in the current frame.

The inventive principles are not limited to number or position of bits shown in the embodiment described with respect to FIG. 31. The first and second portions of each frame have been referred to as command and data portions, but any type of information may be transferred in each portion, and the data portion may contain additional commands, command extensions, etc. The tasks described do not necessarily need to be performed in order described. Thus, the embodiment of FIG. 31 may be modified to accommodate different system requirements or circumstances.

Another inventive principle of this disclosure relates to the use of variable mapping for testing lanes. FIG. 32 illustrates an embodiment of a memory agent, in this example assumed to be a memory module or buffer, according to the inventive principles of this disclosure. The memory agent 134 of FIG. 32 includes a receive link interface 140 which may have one or more receivers and a transmit link interface 142 which may have one or more transmitters. A loopback unit 196, which may be a multiplexer as shown here or other type of redirection device, is capable of selectively mapping receive bit lanes to transmit bit lanes so that the memory agent may retransmit training sequences received from a memory controller back to the controller as return sequences on the transmit bit lanes. By selectively remapping receive bit lanes to transmit bit lanes using different mappings, the controller may analyze the return sequences to identify not only whether there is a failed bit lane, but also whether the failure is on a receive or transmit data path as well as which receive or transmit bit position has failed.

The embodiment of FIG. 32 is shown with a receive and transmit link interfaces having unidirectional bit lanes, but the inventive principles are not limited to this particular configuration of interfaces or types of lanes.

FIG. 33 illustrates example embodiments of two possible bit lane mappings according to the inventive principles of this disclosure. The mappings shown in FIG. 33 assume, for purposes of illustration only, that the memory agent has 10 bit lanes in the receive link interface and 14 bit lanes in the transmit link interface. Using mapping A, the loopback unit

25

redirects the training sequences received on the lower five bits of the receive bit lanes to the transmit bit lanes such that each of the lower five bit lanes is redirected to multiple transmit bit lanes. Using mapping B, the training sequences received on the upper five bits of the receive bit lanes are retransmitted to the multiple transmit bit lanes.

A training sequence may contain a mapping indicator to instruct the memory agent which mapping to use. A training sequence may also contain various groups of bit transmissions that provide test parameters to the memory agent or that provide electrical stress patterns that test the signal integrity of each bit lane. Each of the bit lanes may receive the same training sequences, or different bit lanes may receive different sequences, for example, sequences having different electrical stress patterns.

The training sequences received by the memory agent may be retransmitted without modification so that they function as the return sequences, or the memory agent may modify the sequences or create entirely different sequences. For example, the memory agent may retransmit most of the training sequence as the return sequence while modifying only a small group within the sequence to provide identifying or status information to the memory host.

If memory agents having multiple ports and variable mapping capabilities according to the inventive principles of this disclosure are utilized, for example, in a multiple-agent configuration such as that shown in FIG. 3, the agents may be constructed so that only the outermost agent provides loopback operation, while the other agents operate in a pass-through mode during a testing operation.

Another inventive principle of this disclosure relates to aligning a core clock in a memory agent to a lane. FIG. 34 illustrates an embodiment of a memory agent, in this example assumed to be a memory module or buffer, according to the inventive principles of this disclosure. The memory agent of FIG. 34 includes I/O cells 74, which receive data signals from lanes, which in this example are bit lanes, and generate frame signals FRAME. A logic circuit 198 generates a bit align signal BA in response to frame signals FRAME received from the I/O cells. A core clock generator 200 generates a core clock CORE CLK which is used internally by the memory agent to synchronize one or more of its operations. The core clock generator generates the core clock in response to the bit align signal BA, and it may also synchronize the core clock with a reference clock REF CLK such as that shown in FIG. 8. A base clock generator 202 generates a base clock BASE CLK in response to the core clock. The base clock is utilized by a memory interface 66 to synchronize its internal operations, for example to send commands to a memory device that may be coupled to the memory interface.

The I/O cells may be arranged as part of a redrive circuit 60. A deskew circuit 70 may be coupled between the bit lanes and memory interface to transfer data from the bit lanes to the memory interface. The BA signal may be used by the deskew circuit to align data with one of the bit lanes.

In one example embodiment, the core clock may be aligned to the slowest bit lane. This may reduce or minimize latency if the arrangement of FIG. 34 is used, for example, in a system in which commands are passed from the bit lanes to a memory device that is coupled to the memory interface.

Some possible implementation details for another example embodiment of the arrangement of FIG. 34 are as follows. Each I/O cell may establish its own alignment and drives its FRAME signal to indicate when the bit or frame alignment has been established. The logic circuit 198 may OR the FRAME signals to generate the BA signal which is

26

used by the deskew circuit to align data from all of the bit lanes with the slowest bit lane. Data from bit lanes that are too fast are delayed in the deskew circuit. The core clock generator 200 may align the core clock to the slowest bit lane. This may possibly reduce or minimize the time from the last bit of an address/command arriving at the memory agent to the time it can be used by a memory device coupled to the memory interface. The third base clock generator 202 may generate a base clock (which may be used as a command clock) at some multiple frequency of the core clock. Additional logic circuitry may generate a data interface clock DATA CLK for the memory interface as a delayed version of the core clock. If the arrangement of FIG. 34 is used, for example, in a system that transfers frames as described above with respect to FIG. 31, one or more of the clock signals used by the memory interface may be delayed only long enough to allow the "A" command to propagate through the memory interface.

The inventive principles are not limited to the various embodiments described above with respect to FIG. 34. For example, in FIG. 34, the deskew circuit is shown separately from the I/O cells. The deskew circuit, however, may be integrated into the I/O cells, possibly between a redrive path and a deserializer, or possibly even directly in a redrive path. As a further example, the deskew circuit of FIG. 34 is shown utilizing the same bit align signal that may be used to align the core clock to the bit lanes, but the deskew circuit may also utilize or include separate logic circuitry to align data to one of the bit lanes.

The embodiments described herein may be modified in arrangement and detail without departing from the inventive principles. Accordingly, such changes and modifications are considered to fall within the scope of the following claims.

The invention claimed is:

1. A memory agent comprising:

a memory interface;

a memory channel having a plurality of data lanes;

a core clock generator coupled to the memory interface; and

a logic circuit coupled between the data lanes and the core clock generator to align the core clock to one of the data lanes.

2. A memory agent according to claim 1 wherein the logic circuit is capable of aligning the core clock with the slowest one of the data lanes.

3. A memory agent according to claim 1 further comprising a deskew circuit coupled to transfer data between the data lanes and the memory interface, and coupled to the logic circuit to align data with one of the data lanes.

4. A memory agent according to claim 1 wherein each data lane comprises an I/O cell to generate a frame signal.

5. A memory agent according to claim 4 wherein the logic circuit is coupled to the I/O cells to generate an align signal responsive to the frame signals.

6. A memory agent according to claim 5 wherein the logic circuit is capable of generating the core clock responsive to the align signal.

7. A memory agent according to claim 6 further comprising a base clock generator coupled to the core clock generator to generate a base clock responsive to the core clock.

8. A memory agent according to claim 1 wherein the memory agent comprises a memory buffer.

9. A memory agent according to claim 1 wherein the memory agent comprises a memory module.

10. A memory agent according to claim 1 wherein the memory agent comprises a memory controller.

27

11. A memory agent according to claim 1 wherein the data lanes comprise bit lanes.

12. A system comprising:

memory agent comprising:

a memory interface;

a memory channel having a plurality of data lanes,

a core clock generator coupled to the memory interface, and

a logic circuit coupled between the data lanes and the core clock generator to align the core clock to one of the data lanes; and

a memory controller coupled to the data lanes.

13. A system according to claim 12 wherein the logic circuit is capable of aligning the core clock with the slowest one of the data lanes.

14. A system according to claim 12 wherein the memory agent further comprises a deskew circuit coupled to transfer data between the data lanes and the memory interface, and coupled to the logic circuit to align data with one of the data lanes.

15. A system according to claim 12 wherein the memory agent is a first memory agent, and further comprising a second memory agent coupled to the first memory agent.

16. A system according to claim 15 wherein the second memory agent comprises:

a second core clock generator, and

a second logic circuit coupled between the data lanes and second core clock generator to align the second core clock to one of the data lanes.

28

17. A system according to claim 15 wherein the second memory agent further comprises a second deskew circuit coupled to the plurality of data lanes.

18. A method of operating a memory agent comprising:

receiving a plurality of data signals on a plurality of data lanes in a memory channel;

aligning a core clock with one of the data lanes; and

coupling the core clock to a memory interface.

19. A method according to claim 18 wherein aligning the core clock with one of the lanes comprises aligning the core clock with the slowest one of the data lanes.

20. A method according to claim 18 further comprising deskewing the plurality of signals received on the data lanes.

21. A method according to claim 18 further comprising generating a base clock responsive to the core clock.

22. A method according to claim 18 wherein aligning the core clock comprises generating a plurality of frame signals responsive to the plurality of signals received on the data lanes.

23. A method according to claim 22 wherein aligning the core clock further comprises generating an align signal responsive to the plurality of frame signals.

24. A method according to claim 23 wherein generating the align signal comprises ORing the plurality of frame signals.

25. A method according to claim 18 further comprising generating the core clock responsive to a reference clock.

* * * * *