

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 May 2007 (10.05.2007)

PCT

(10) International Publication Number
WO 2007/052229 A2

(51) International Patent Classification: **Not classified**

[US/US]; c/o Philips Intellectual Property & Standards, Cross Oak Lane, Redhill Surrey RH1 5HA (GB).

(21) International Application Number:
PCT/IB2006/054056

(74) Agents: **WHITE, Andrew, G.** et al.; c/o Philips Intellectual Property & Standards, Cross Oak Lane, Redhill Surrey RH1 5HA (GB).

(22) International Filing Date:
2 November 2006 (02.11.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/264,821 4 November 2005 (04.11.2005) US

(71) Applicant (for all designated States except US): **KONINKLIJKE PHILIPS ELECTRONICS N.V.** [NL/NL]; Groenewoudseweg 1, NL-5621 BA Eindhoven (NL).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(71) Applicant (for AE only): **U.S. PHILIPS CORPORATION** [US/US]; 1251 Avenue Of The Americas, New York, New York NY-10020 (US).


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

(72) Inventor; and

(75) Inventor/Applicant (for US only): **PONTIUS, Tim**

[Continued on next page]

(54) Title: ALIGNMENT AND DESKEW FOR MULTIPLE LANES OF SERIAL INTERCONNECT



(57) Abstract: Methods and apparatus are provided for data communication between a transmitter and receiver over a plurality of serial links, which cause the transmitter to send serialized groups of bits down each lane, in such a way that the first bit of each group (each lane) is guaranteed to arrive in the correct order at the receiving end. Various embodiments of the present invention include declaring a budget for the maximum skew between lanes. In such embodiments, subsequent to determining the skew budget between lanes, the data to be transmitted is divided into groups of N bits, where N is any convenient number larger than M times S, with M being the number of lanes and S being the budgeted skew, in bit times.

WO 2007/052229 A2



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished
upon receipt of that report*

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a
patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the
earlier application (Rule 4.17(iii))*

*For two-letter codes and other abbreviations, refer to the "Guid-
ance Notes on Codes and Abbreviations" appearing at the begin-
ning of each regular issue of the PCT Gazette.*

DESCRIPTION

ALIGNMENT AND DESKEW FOR MULTIPLE LANES OF SERIAL INTERCONNECT

The present invention relates generally to high speed data communication, and more particularly relates to reducing the complexity of circuitry used to receive, deserialize, and align, multiple streams of high speed serial data.

Advances in semiconductor manufacturing technology have resulted in very highly integrated circuits, which contain millions of transistors, along with various other components and interconnections. The integrated circuits that result from the aforementioned advances provide significant electronic functionality at relatively low costs. Such advances have been incorporated into both digital and analog integrated circuits, and these advances have resulted in a wide variety of consumer electronic products, including products in which high speed communication, or transfer, of data between various integrated circuits or modules is required.

Conventionally, requirements for increased data transfer rates were satisfied by means such as increasing the number of parallel data communication paths. That is, wider busses were conventionally used to meet the requirements for increased data transfer rates. However, there are several disadvantages with such an approach. For example, wide parallel busses consume more board space and therefore lead to increased costs. Additionally, wide parallel busses typically require higher pin counts on the chips, or integrated circuits, that interface with those wide parallel busses which in turn consumes more chip area and requires larger packages, both of which result in higher costs.

Another approach to dealing with increased data transfer rates is to provide fewer signal paths, wherein those signal paths are used to transfer data at higher speeds. Typically such arrangements provide for parallel-to-serial conversion on the transmitter side, and serial-to-parallel conversion on the receiver side.

Many contemporary architectures provide high speed serial pathways for data communications between integrated circuits and/or modules, in place of wide

parallel busses. Such architectures provide advantages including, but not necessarily limited to, reduced pin counts, and reduced board space requirements.

In some instances, a plurality of high speed serial data paths are provided in order to increase the effective data transfer rate. Unfortunately, conventional systems that incorporate multiple high speed serial data paths, are subject to skewing problems and must provide additional circuitry, such as buffers and control logic to properly deskew and realign the data.

What is needed are methods and apparatus for providing high speed serial data communication between integrated circuits and/or modules, in which the requirements for buffers and control logic used in deskewing and/or realignment operations are reduced or eliminated.

In accordance with the present invention methods and apparatus are provided for data communication between a transmitter and receiver over a plurality of serial links, or lanes, which cause the transmitter to send serialized groups of bits down each of the plurality of serial links, in such a way that the first bit of each group is guaranteed to arrive in the correct order at the receiving end.

Various embodiments of the present invention include declaring a budget for the maximum skew between lanes. In such embodiments, subsequent to determining the skew budget between lanes, the data to be transmitted is divided into groups of N bits, where N is any convenient number larger than M times S , with M being the number of lanes and S being the budgeted skew, in bit times.

Further features and advantages will now be described, by way of example only, with reference to the accompanying drawings in which :

Fig. 1 is a schematic block diagram of a conventional system including multiple serial links between a data transmitter and a data receiver, and further showing a byte reorder buffer used to support deskewing and realignment operations.

Fig. 2 is a schematic block diagram of a system in accordance with the present invention in which predetermined delays are introduced at the transmitter side of the serial links such that the data realignment operations on the receiver side are obviated.

Fig. 3 is a flow diagram illustrating a method in accordance with the present invention.

Generally, the present invention relates to providing methods and apparatus for reducing or eliminating the overhead of deskewing and realignment found in conventional data receivers. More particularly, by introducing predetermined amounts of delay into the transmissions on various ones of a plurality of serial links (also known as lanes), the data arriving at the receiver is known to be in order and therefore byte order realignment is not required.

Reference herein to "one embodiment", "an embodiment", or similar formulations, means that a particular feature, structure, operation, or characteristic described in connection with the embodiment, is included in at least one embodiment of the present invention. Thus, the appearances of such phrases or formulations herein are not necessarily all referring to the same embodiment. Furthermore, various particular features, structures, operations, or characteristics may be combined in any suitable manner in one or more embodiments.

The terms integrated circuit (IC), semiconductor device, monolithic device, microelectronic device, and chip are often used interchangeably in the field of electronics generally. The present invention is applicable to all the above as they are generally understood in the field.

When a serial link is used for data communication, the data, which is typically parallel, is converted to a serial bit stream and then transmitted over the serial link. At the receiving end, the serial bit stream is converted back to parallel data so that it may be used in parallel. It will be appreciated that the bandwidth of the interconnect is limited by the bit rate of the serial bit stream.

One method used to increase the bandwidth provided by the aforementioned serial link, is to provide multiple serial lanes. For the transmitter, this simply requires sending two, three, four, or possibly more bits at a time. For example, if there are two lanes, then two bits can be sent at once. For each subsequent bit period, two more bits are sent. If the bits are numbered 0,1,2,3,4,5,6,7..., then bits 0 and 1 are sent together in the first bit time, bits 2 and 3 are sent together in the second bit time, and so on. The number of serial bits sent at one time is commonly called the

number of lanes. Consider a four-lane road, which can carry twice as many cars (serial bits) as a two-lane road.

If the interconnect is synchronous (i.e., timing of the data is related to some synchronizing event, such as a clock edge), then the bits can be expected to arrive at the receiver at substantially the same time, that is at the same synchronizing event. However, a fully synchronous serial link is significantly restricted in frequency, which in turn restricts bandwidth. To operate at higher frequencies, the circuitry used to drive the serial lines is likely to implement embedded synchronization techniques, where the serial data is encoded in such a way that it is possible for the receiver to recover a "clock" from the serial data. In this scenario, that is where a clock is recovered from the serial data stream, it is likely that skew develops between the lanes. Using the example above, bit 1 might arrive before bit 0. Alternatively, bit 0 may arrive together with bit 3; then bit 2 arrives together with bit 5, and so on. At very high frequencies, there could actually be multiple bit times of skew between the lanes. To recover the original bit stream, the receiver must be able to re-align the bits. This process may be referred to as de-skewing the bit streams.

A common approach to de-skewing the lanes is to use synchronization codes (sync codes) that are sent down each lane respectively. If the sync codes are sent down each lane at the same time, then the receiver can detect the delay between receiving the sync codes and use that information to know how to de-skew the received bit streams. One disadvantage with this approach is that the receiver does all of the difficult work. The transmitter does little to make deskewing simple for the receiver. This can be an issue if the receiver is implemented in an expensive manufacturing process, or one that is not well suited for digital logic (e.g., a DRAM process). It may also be the case that pushing all of the work to one end results in a more complicated design than that which could be achieved if the required work is split between the transmitter and the receiver.

Various embodiments of the present invention cause the initiator, that is, the transmitter, to send serialized groups of bits down each lane, in such a way that the first bit of each group (each lane) is guaranteed to arrive in the correct order at the receiving end.

Various embodiments of the present invention include declaring, or determining, a budget for the maximum skew between lanes. In such embodiments, subsequent to determining the skew budget between lanes, the data to be transmitted is divided into groups of N bits, where N is any convenient number larger than M times S , with M being the number of lanes and S being the budgeted skew, in bit times.

For example, with 2 lanes, and a skew budget of 5 bit times, N must be greater than 10. In this example 16 (two bytes) would be a good choice. With 4 lanes and 5 bit times, N must be greater than $4*5=20$, so perhaps 24 bits (three bytes) would be a convenient choice.

The transmitter then captures the first N bits of the data to be transmitted, and begins transmitting those bits on lane 0. The next N bits will be transmitted on lane 1, but not until after N/M bit times have passed. This continues on each lane until bit groups have been started on all lanes. At approximately N/M bit times later (could be one more or one fewer, because N/M is generally rounded to an integer), the transmitter may send another group of N bits over lane 0.

Consider an illustrative example of the present invention with two lanes and a skew of 5 bits times (i.e., $M*S = 10$). The number of bits might be chosen as $N=16$ (two bytes) for convenience. The bytes zero and one are captured and transmitted down lane 0. Subsequently, beginning at a point in time that is eight bit periods later ($16/2=8$), bytes two and three begin transmitting over lane 1. Eight bit times later (i.e., a total of 16 bit times from the beginning), bytes four and five begin transmitting over lane 0. In this example, at the receive end, the bits on lane 1 may arrive 8 bit times after the bits on lane 0. However, because of the 5-bit time skew budget, lane 1 bits may arrive as few as 3 bit times after lane 0 or as many as 13 bit times after lane 0. In any case, the first group of N bits sent over lane 1 always arrives after the first group of N bits on lane 0 and before the second group of N bits on lane 0. This makes it very easy for the receiver to deskew the groups because the order of arrival is guaranteed.

Consider a more complex illustrative example with 2 lanes and a skew budget of 9 bits. Here assume that $N=24$ is selected (must be greater than $2*9=18$). The first group of 24 bits starts on lane zero at time 0. The second group of 24 bits starts on lane 1 twelve bit times later ($24/2=12$). At the receiving end, lane 1 data

may arrive as early as three bits times later (12-9) or as late as 21 bit times later (12+9), but it will still always be between the first and second groups transmitted over lane 0.

Referring to Fig. 3, an illustrative process **300**, in accordance with the present invention is described. In a first operation, a determination is made **302** as to the skew between lanes. This determination can be made in any suitable manner, for example the transmitter may reference a table that is dynamically updated. Such a table may be disposed within the transmitter integrated circuit or it may be external thereto. Alternatively, the skew between lanes may have been characterized for a particular design and the skews are pre-programmed into the transmitter device at the factory. A determination is made **304** as to the number of bits, N, that will be taken in groups for transmission. As noted above, N is greater than M times S, where M is the number of lanes, and S is the skew in bit times. N bits are then captured **306** from a memory. Such a memory may be referred to as transmit buffer, or a packet FIFO. Any suitable storage elements may be used to contain the data to be transmitted. A transmission **308** of the captured data is then begun, in which the data is sent on a designated one of the plurality of lanes. Any suitable means may be used for indicating the appropriate lane upon which to transmit. In one embodiment a pointer is used to indicate the lane, and this pointer is then incremented **310**, or otherwise updated, so that it point to the next lane upon which data will be transmitted. The next N bits of transmit data are then captured **312**. A determination **314** is made regarding whether N/M bit times have elapsed since the last transmission began. If the determination of **314** is negative, the process **300** waits for the determination to become affirmative. When the determination of **314** is affirmative, a determination **316** is made regarding whether more data is to be sent. If there is no more data to send, then process **300** ends at **318**. However, if there is more data to send, control passes to step **308**, and process **300** loops until the appropriate amount of data has been transmitted.

Embodiments of the present invention can be applied to serial data transmission systems where the allowed complexity of the receiver is limited. The simplification of delaying the transmission of bits over subsequent lanes shifts some

of the complexity from the receiver to the transmitter, making the receiver's task much easier without significantly increasing the complexity of the transmitter.

It is to be understood that the present invention is not limited to the embodiments described above, but encompasses any and all embodiments within the scope of the subjoined Claims and their equivalents.

CLAIMS

1. A method (300) of transmitting data between a transmitter having data to send, and a receiver, wherein a plurality of serial links are disposed between the transmitter and receiver, comprising:
 - a) determining (302) a skew budget between the plurality of serial links;
 - b) determining (304) a number of bits N , where N is an integer greater than $M \times S$, where M is the number of serial links, and S is the skew budget in bit times;
 - c) capturing (306) a first N bits of the data to be sent, and beginning the transmission (308) of the first N bits on a first one of the plurality of serial links; and
 - d) capturing (312) a second N bits of the data to be sent, at approximately N/M bit times subsequent to the beginning of the transmission on the first one of the plurality of serial links, and beginning the transmission of the second N bits on a second one of the plurality of serial links.
2. The method of Claim 1, wherein capturing the first N bits comprises accessing a memory disposed within the transmitter and serializing the first N bits.
3. The method of Claim 2, wherein capturing the second N bits comprises accessing the memory disposed within the transmitter and serializing the second N bits.
4. The method of Claim 1, further comprising the first N bits and the second N bits at the receiver.
5. The method of Claim 1, further comprising capturing and transmitting additional groups of N bits on sequentially designated lanes, the beginnings of the transmissions being separated by N/M bit times, until all the data to be sent has been transmitted.
6. A transmitter adapted for use with a plurality of serial links, comprising:
 - a packet memory operable to store packets for transmission, the packet memory producing a multi-bit output when accessed;

a means for determining a skew budget the plurality of serial links;
a means for determining a number of bits N , where N is greater than $M \times S$, where M is the number of serial links, and S is the skew budget in bit times;
a means for capturing a first N bits of the data to be sent, and beginning the transmission of the first N bits on a first one of the plurality of serial links; and
a means for capturing a second N bits of the data to be sent, at approximately N/M bit times subsequent to the beginning of the transmission on the first one of the plurality of serial links, and beginning the transmission of the second N bits on a second one of the plurality of serial links.

7. The transmitter of Claim 6, wherein the means for capturing the first N bits comprises a means for accessing the packet memory and serializing the first N bits.

8. The transmitter of Claim 7, wherein the means for capturing the second N bits comprises a means for accessing the packet memory and serializing the second N bits.

9. A method (300) of transmitting data between a transmitter and a receiver over a plurality of serial links, comprising:

- a) determining (302) a skew between the serial links;
- b) determining (304) a number, N , where N is greater than M times S , M is the number of serial links, and S the skew;
- c) capturing (306) N bits from a transmit buffer;
- d) begin transmitting (308) the captured N bits on a serial link designated by a serial link pointer;
- e) updating (310) the serial link pointer;
- f) capturing (312) a next N bits from the transmit buffer;
- g) determining (314) whether N/M bit times have elapsed since the beginning of the last transmission;
- h) determining (316) whether there is more data to be sent;
- i) repeating steps (d) through (h).

10. The method of Claim 9, wherein determining a skew between links comprises

referring to a table.

11. The method of Claim 9, wherein capturing N bits comprises reading from a memory and serializing the data.

12. The method of Claim 10, wherein the table is disposed external to an integrated circuit on which the transmitter is disposed.

13. The method of Claim 9, wherein updating the serial link pointer comprises incrementing a counter.

14. The method of Claim 9, wherein N is a multiple of 8.

1/2

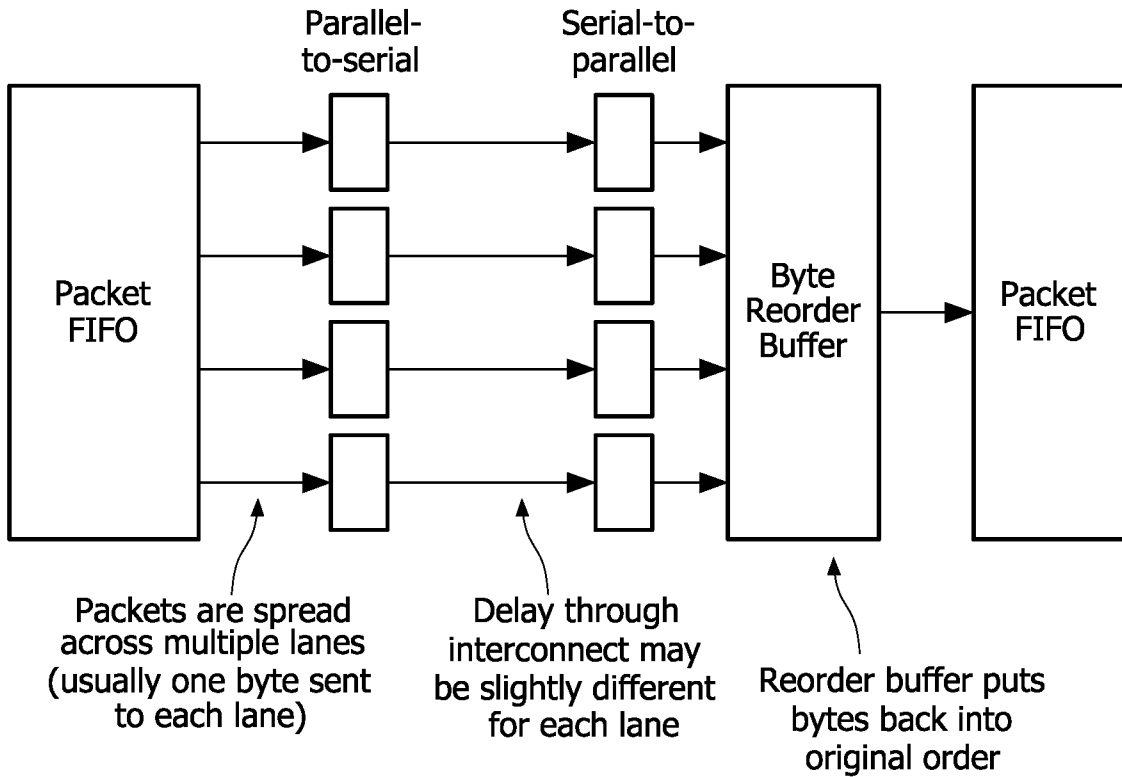


FIG. 1 PRIOR ART

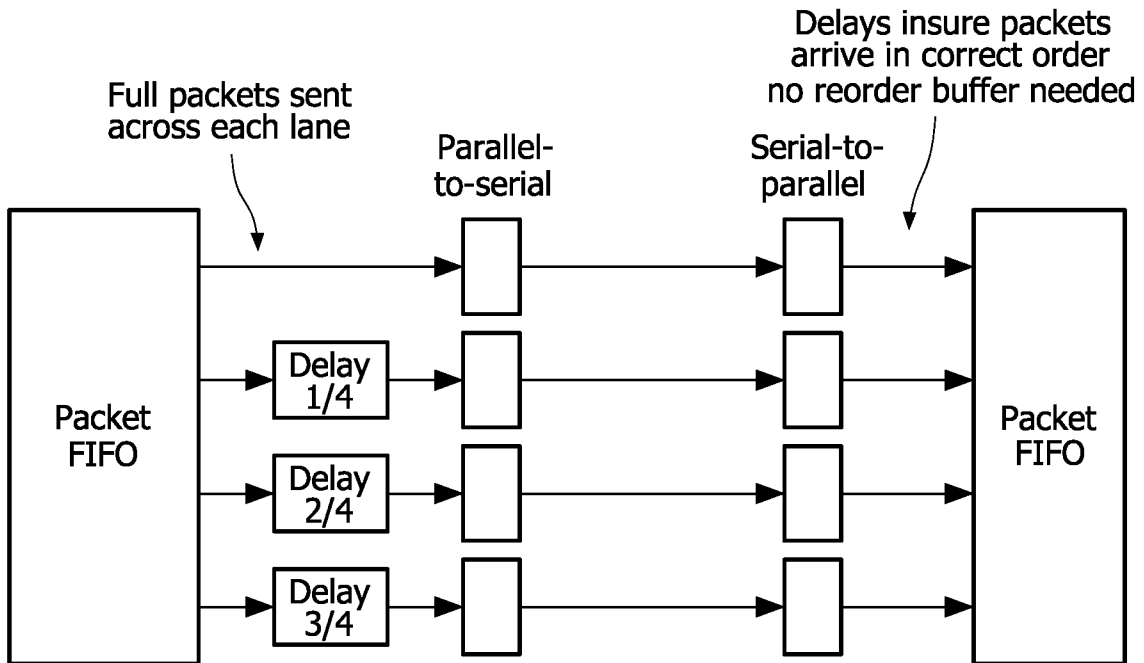


FIG. 2 INVENTION

2/2

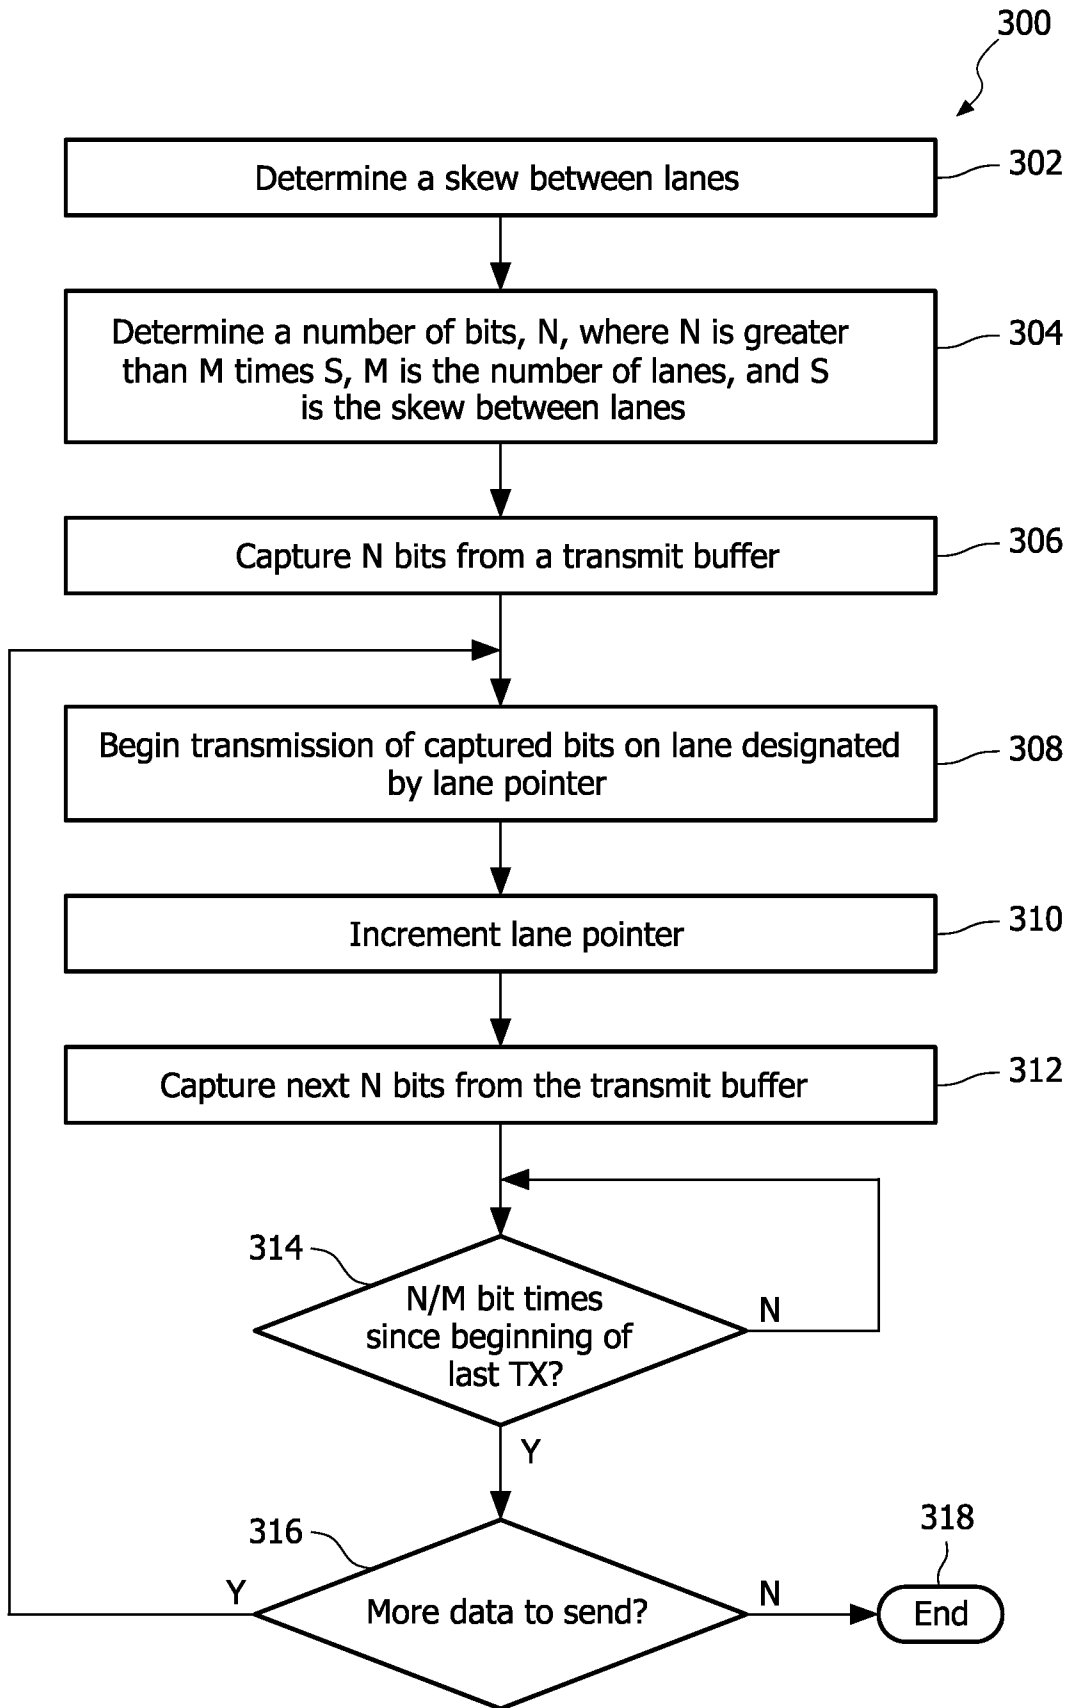


FIG. 3