

(19) **United States**

(12) **Patent Application Publication**  
Dally et al.

(10) **Pub. No.: US 2014/0266416 A1**  
(43) **Pub. Date: Sep. 18, 2014**

(54) **ON-PACKAGE MULTIPROCESSOR  
GROUND-REFERENCED SINGLE-ENDED  
INTERCONNECT**

**Publication Classification**

(51) **Int. Cl.**  
**H05K 1/11** (2006.01)  
(52) **U.S. Cl.**  
CPC ..... **H05K 1/11** (2013.01)  
USPC ..... **327/564**

(71) Applicant: **NVIDIA Corporation**, Santa Clara, CA (US)

(72) Inventors: **William J. Dally**, Los Altos Hills, CA (US); **Brucek Kurdo Khailany**, Austin, TX (US); **John W. Poulton**, Chapel Hill, NC (US); **Thomas Hastings Greer, III**, Chapel Hill, NC (US); **Carl Thomas Gray**, Apex, NC (US)

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

(21) Appl. No.: **13/946,980**

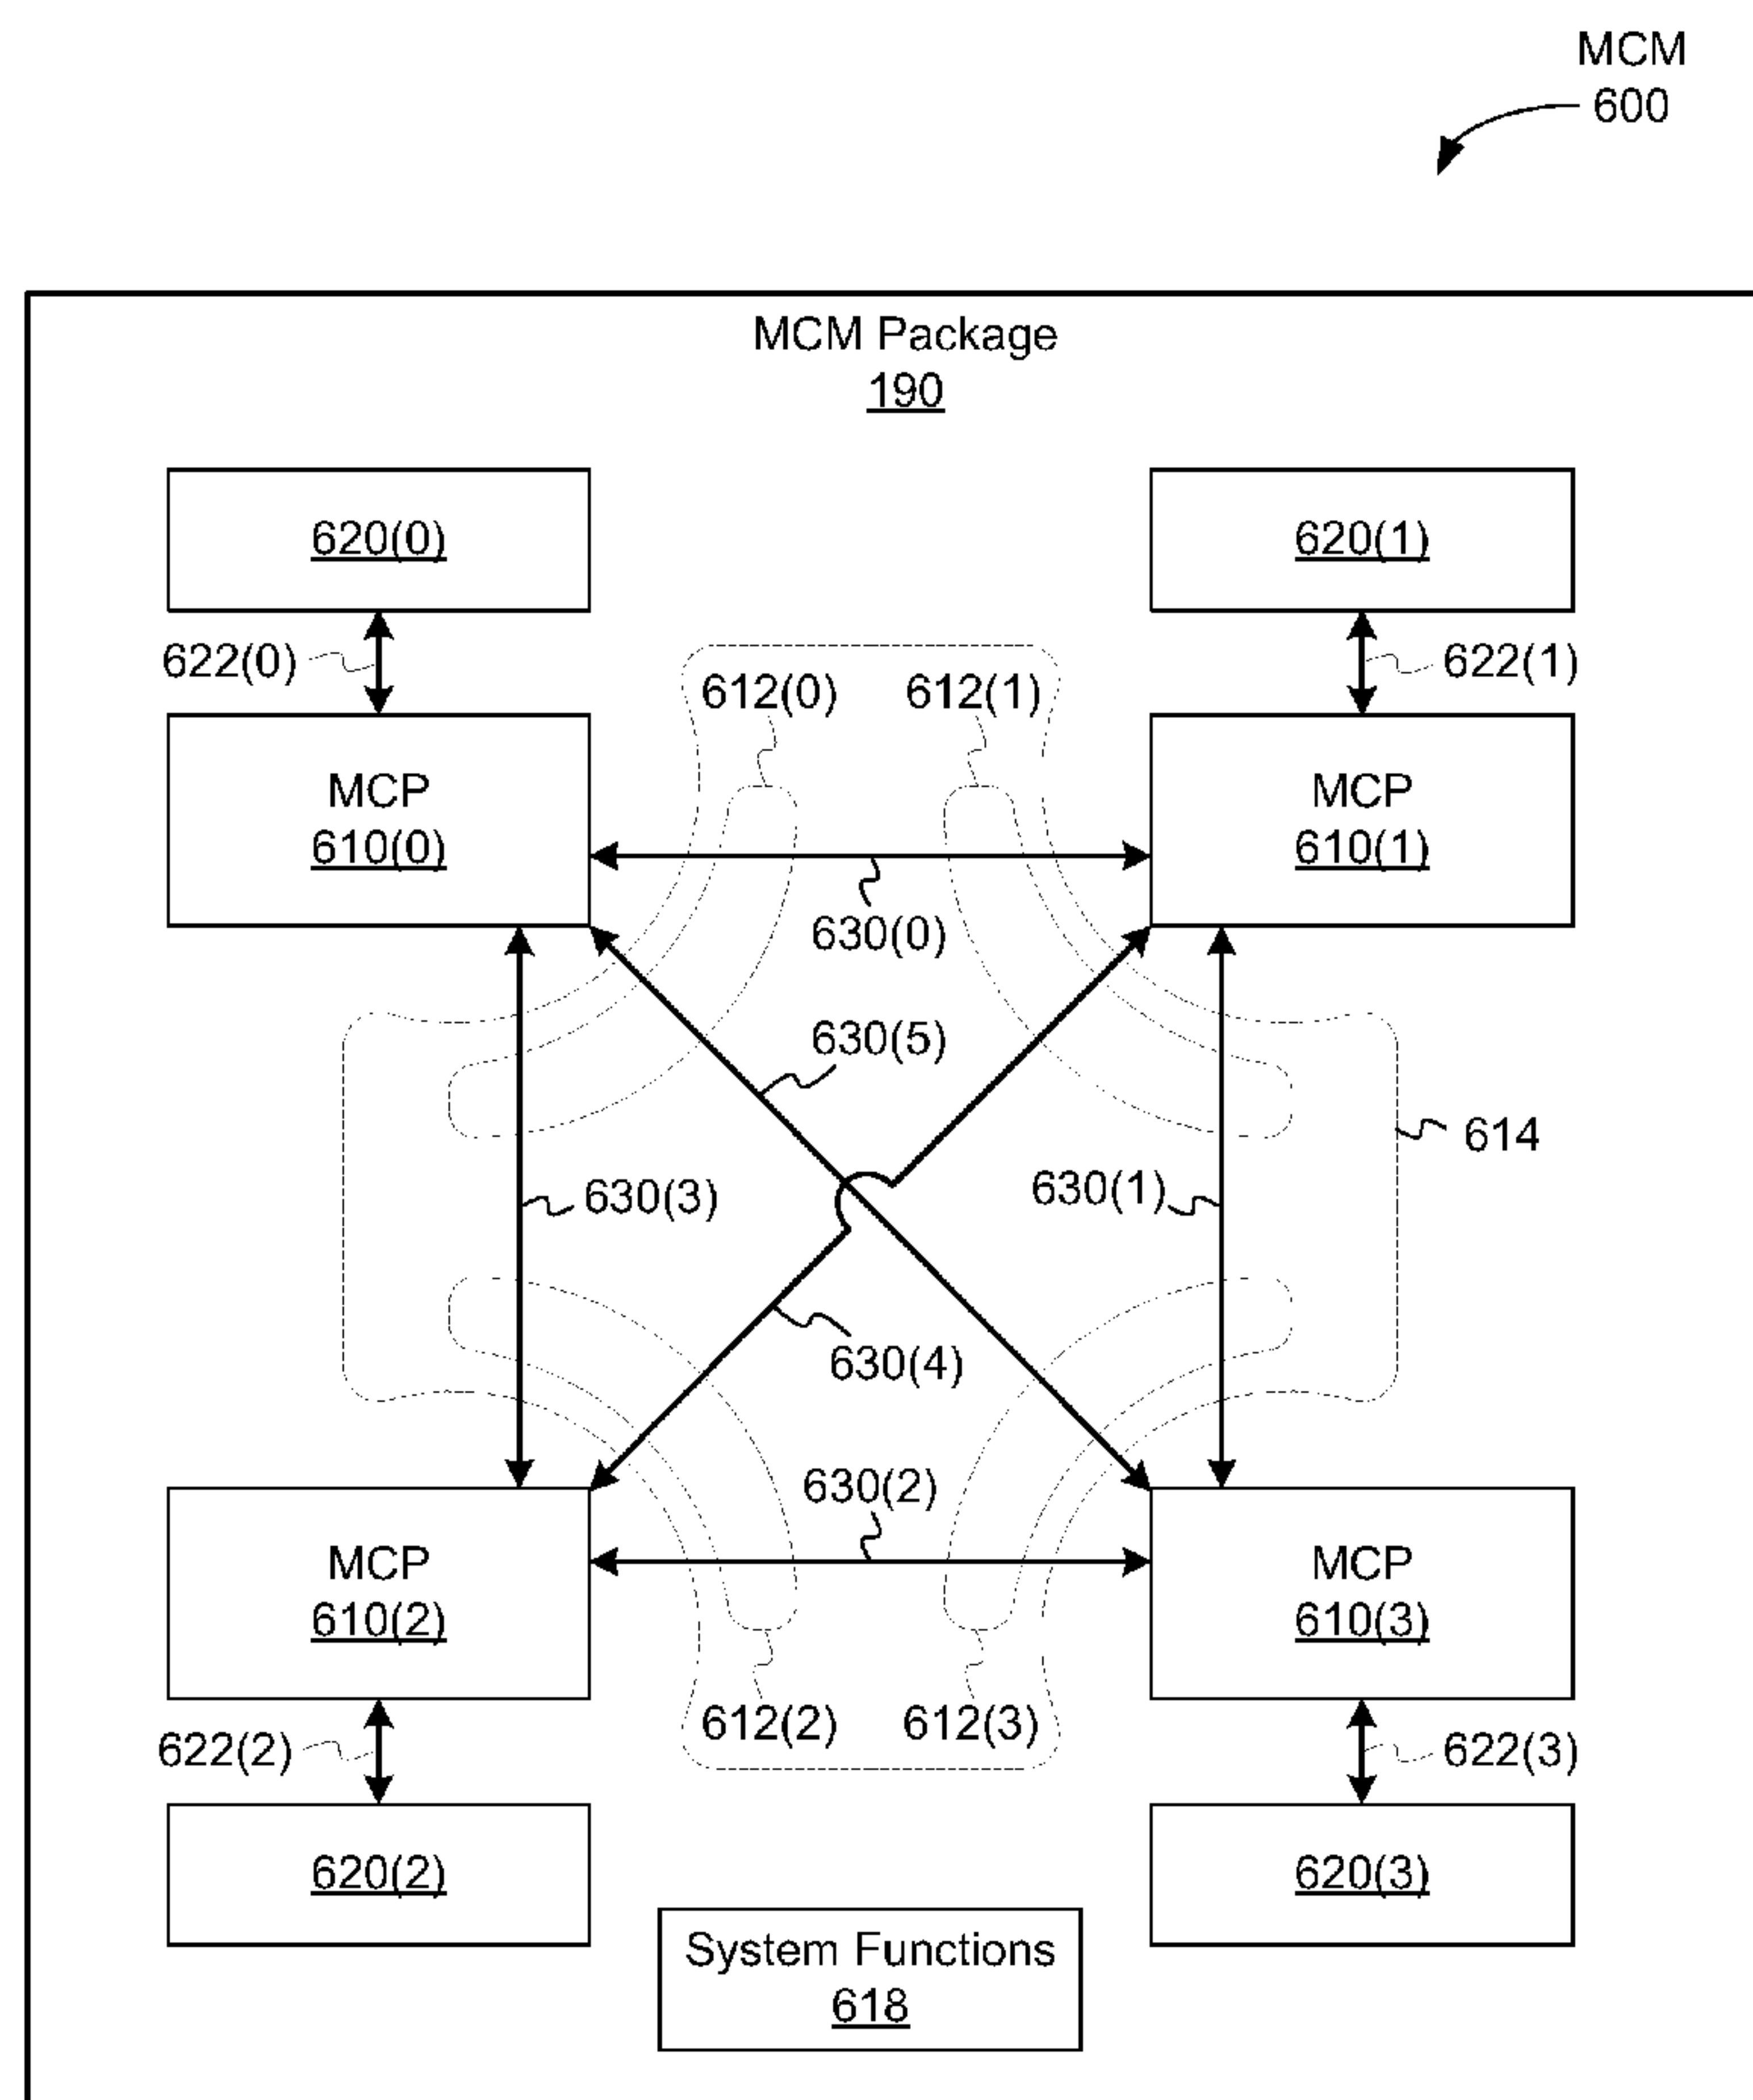
(22) Filed: **Jul. 19, 2013**

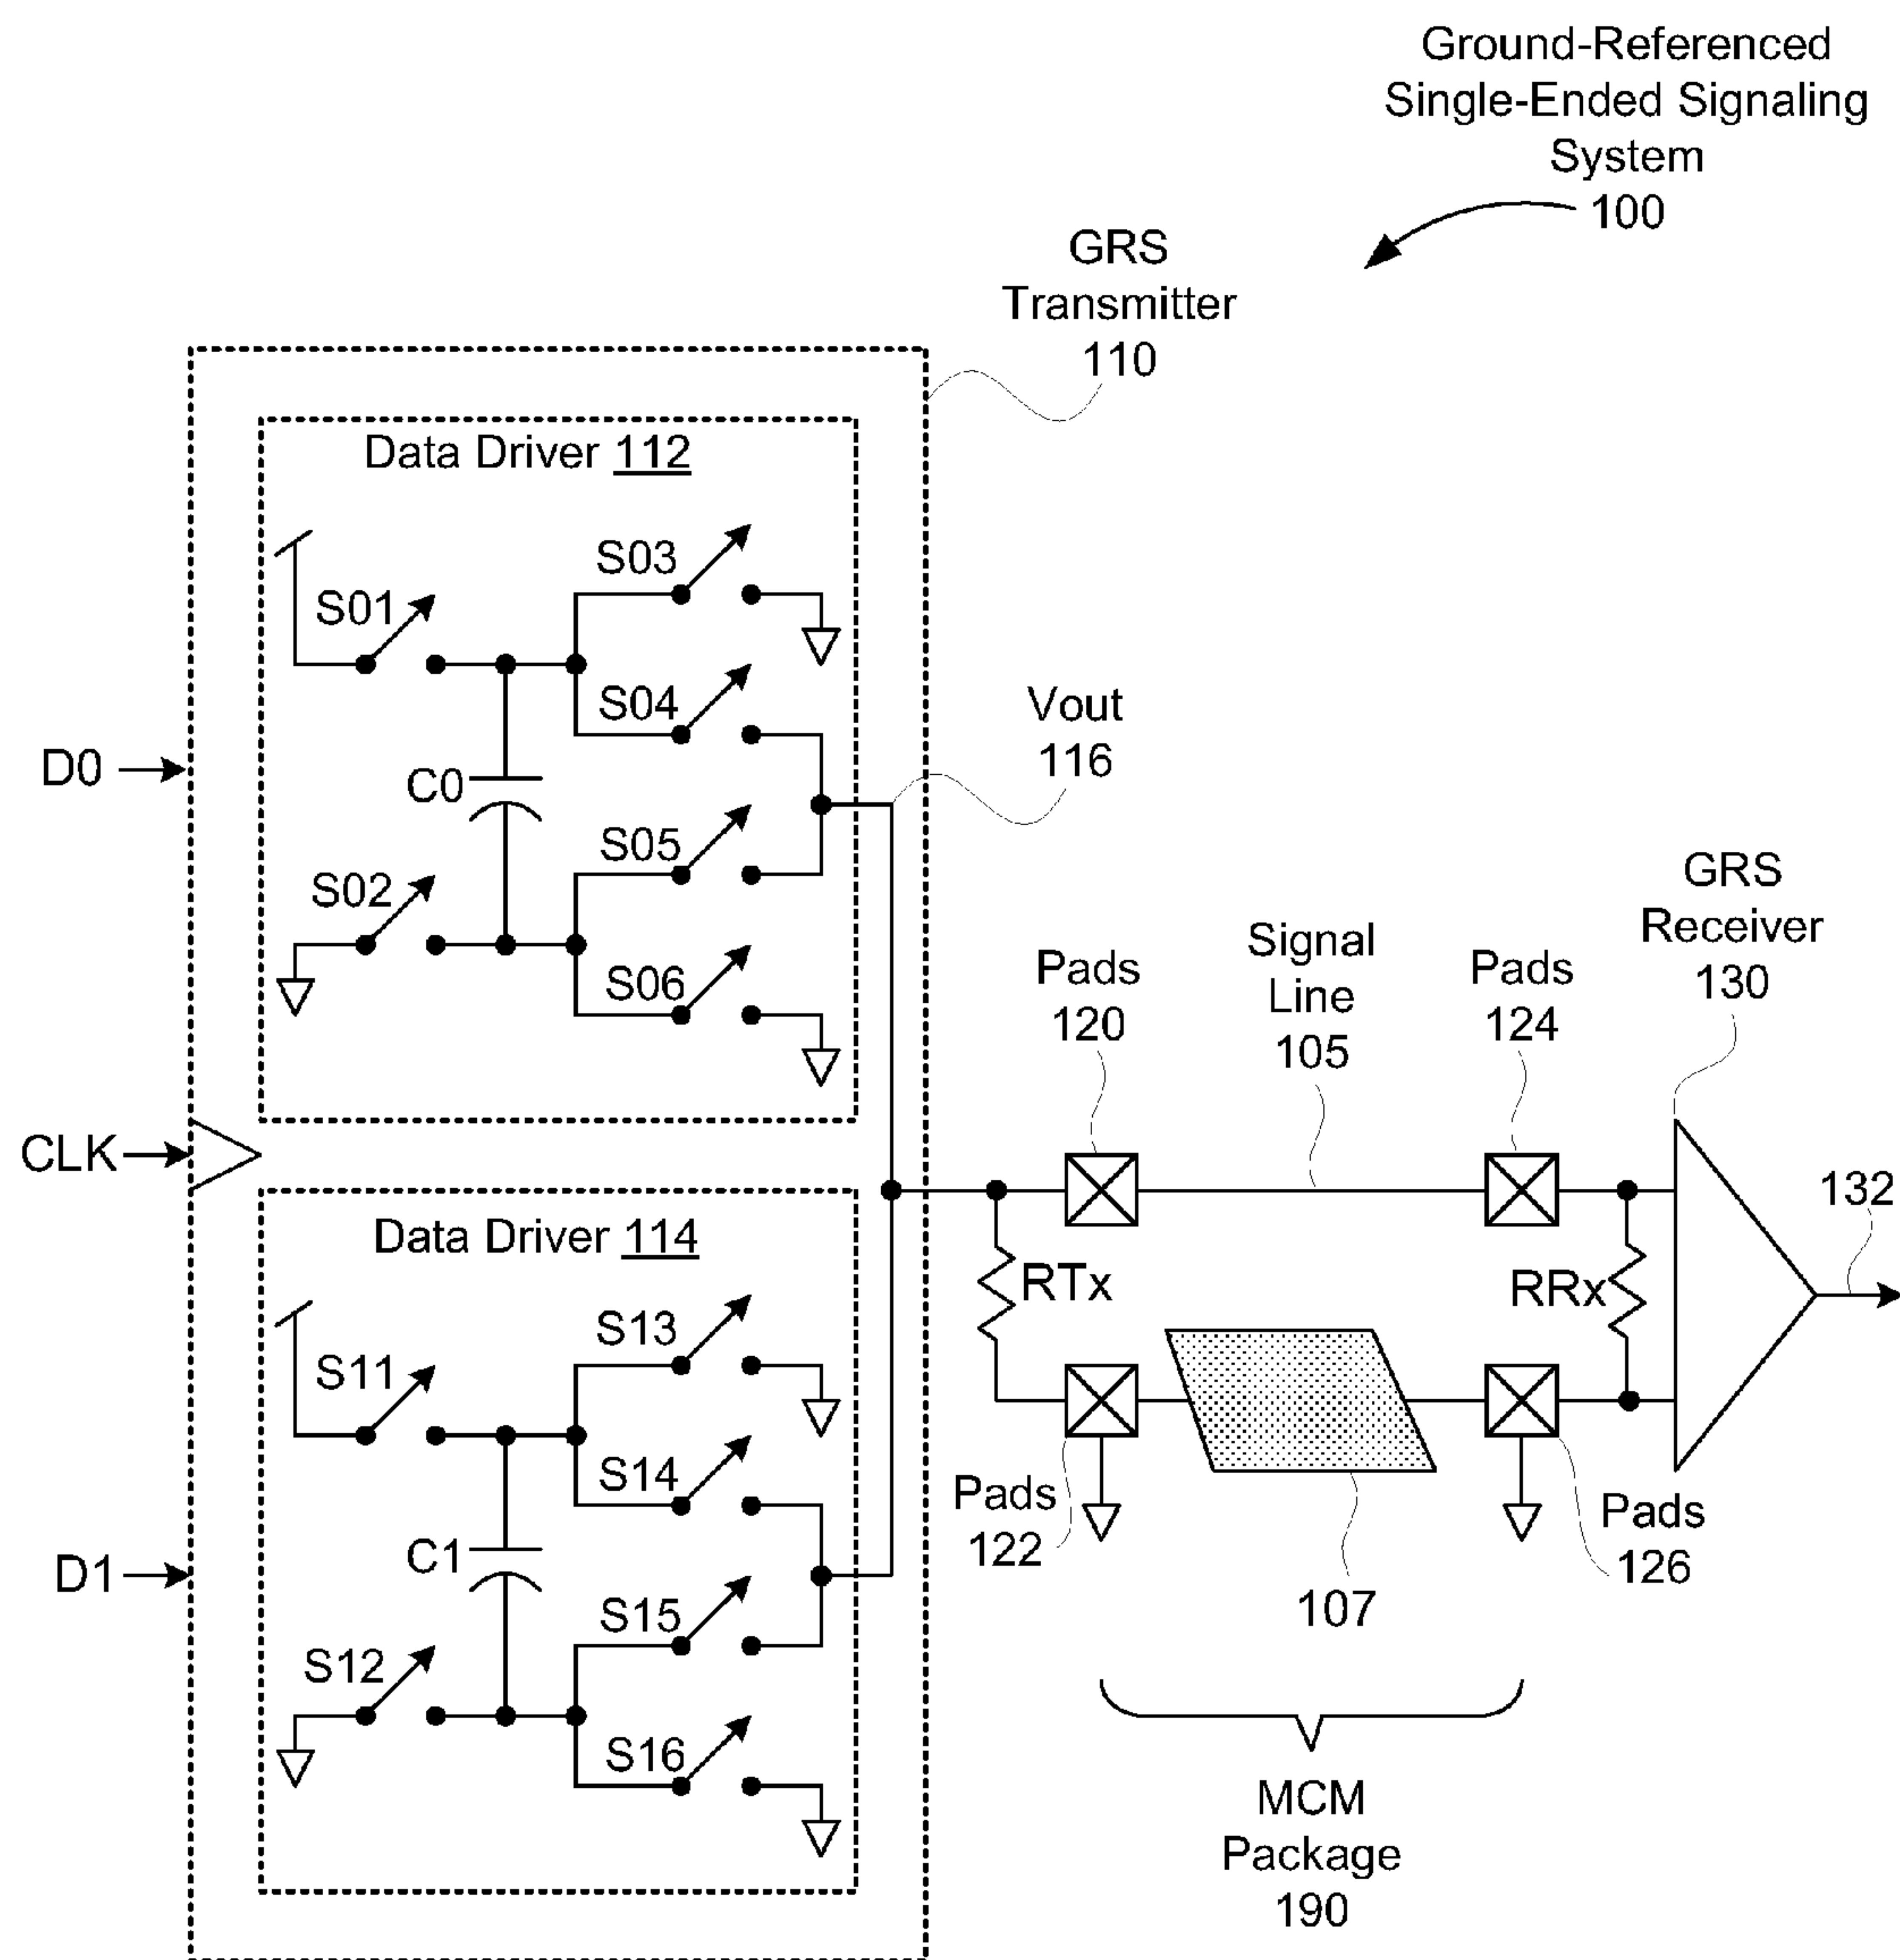
**Related U.S. Application Data**

(63) Continuation-in-part of application No. 13/844,570, filed on Mar. 15, 2013.

(57) **ABSTRACT**

A system of interconnected chips comprising a multi-chip module (MCM) includes a first processor chip, a second processor chip, and an MCM package configured to include the first processor chip, the second processor chip, and an interconnect circuit. The first processor chip is configured to include a first ground-referenced single-ended signaling (GRS) interface circuit. A first set of electrical traces fabricated within the MCM package and configured to couple the first GRS interface circuit to the interconnect circuit. The second processor chip is configured to include a second GRS interface circuit. A second set of electrical traces fabricated within the MCM package and configured to couple the second GRS interface circuit to the interconnect circuit.





**Fig. 1A**

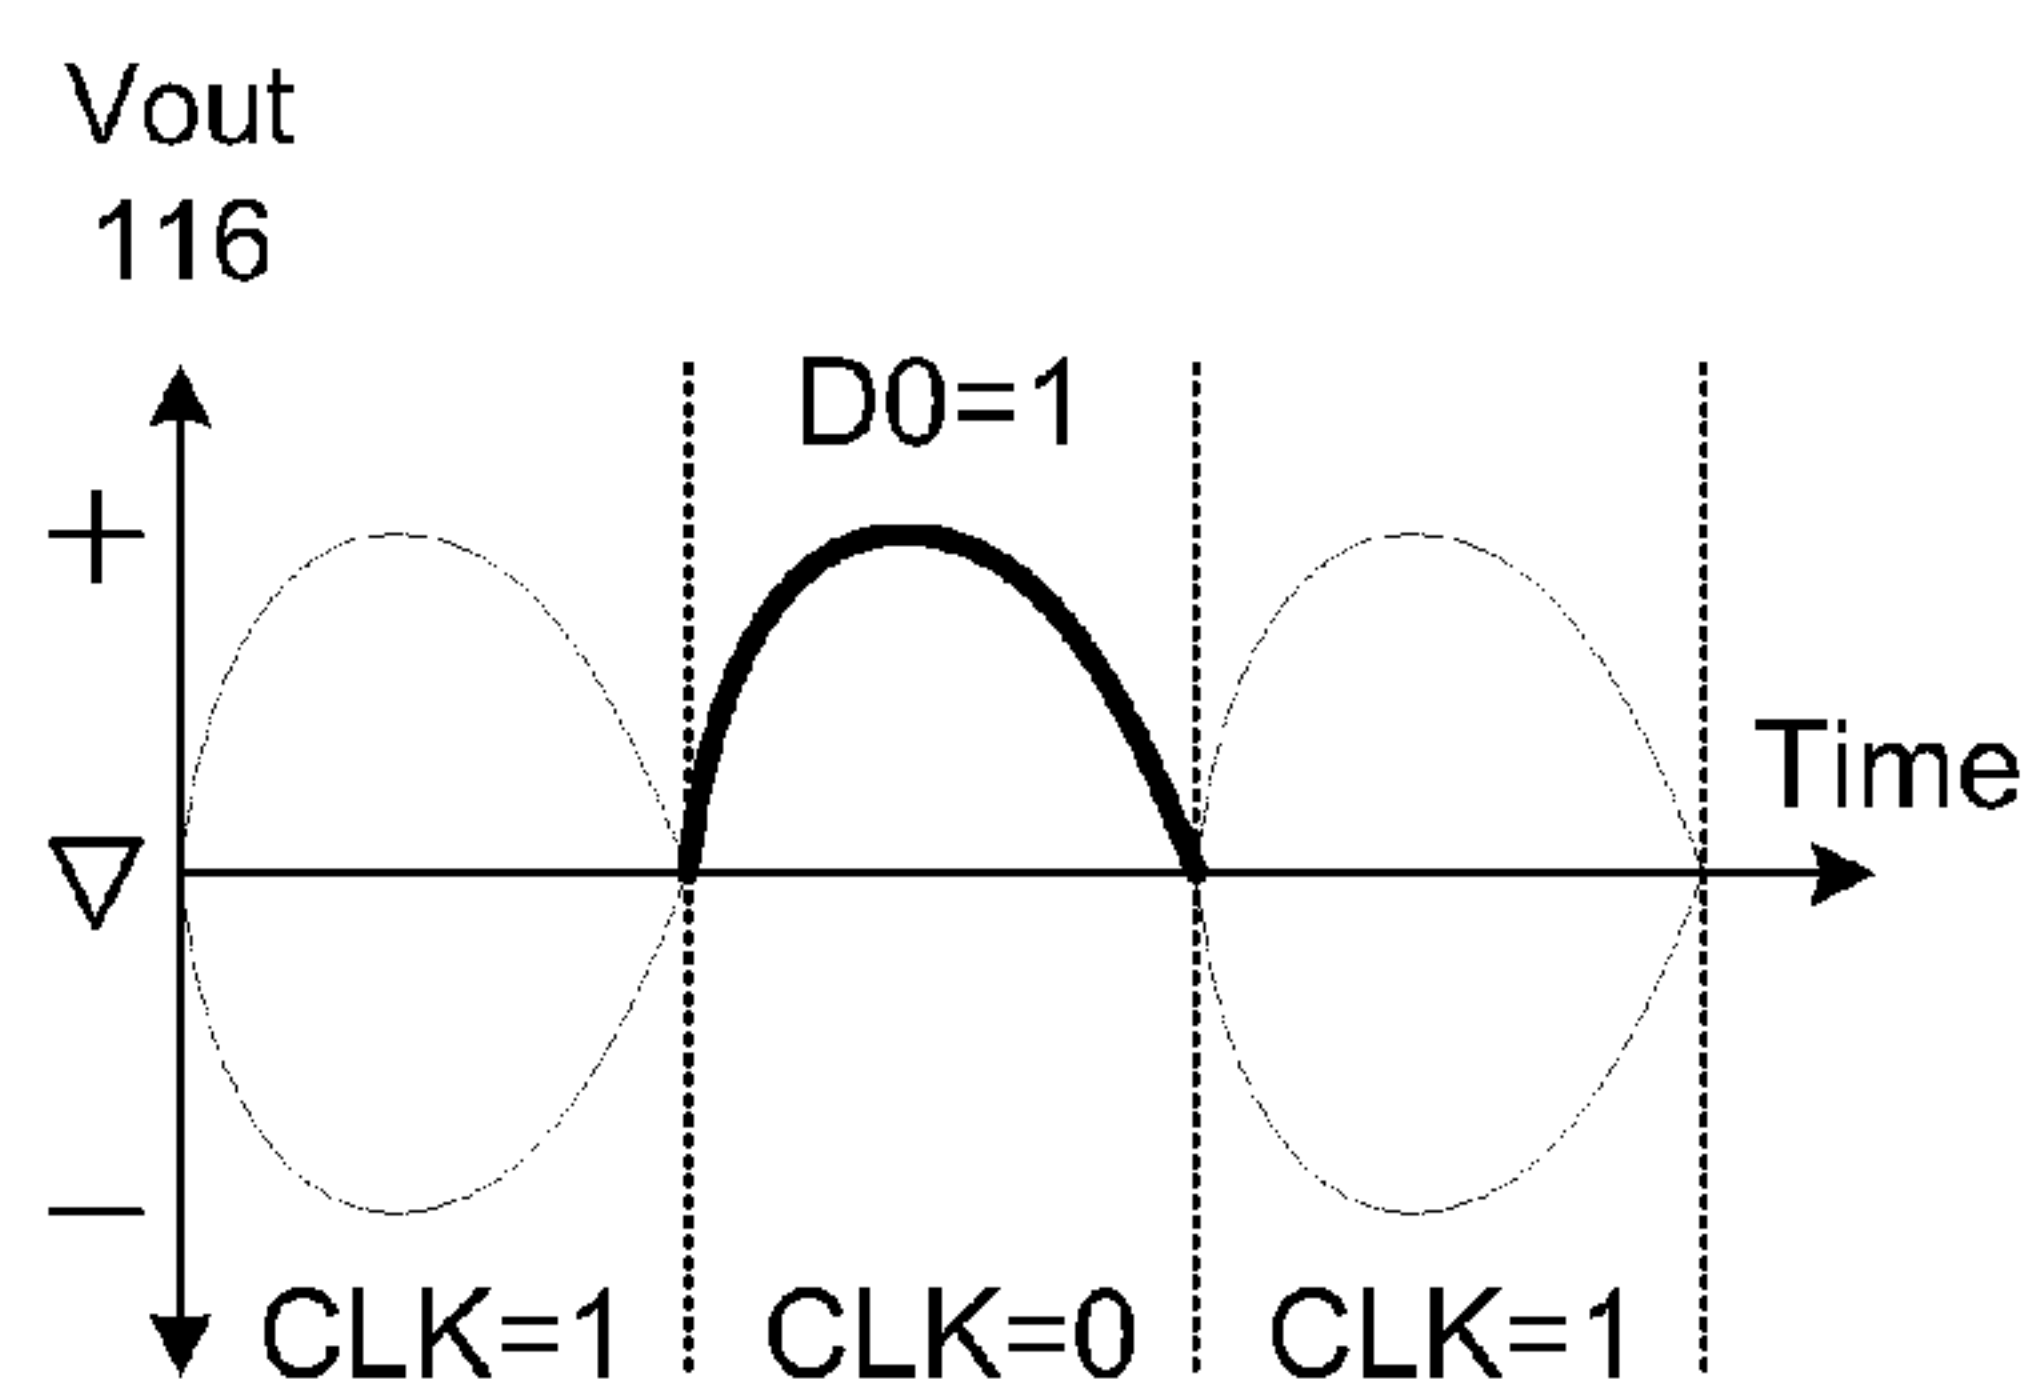
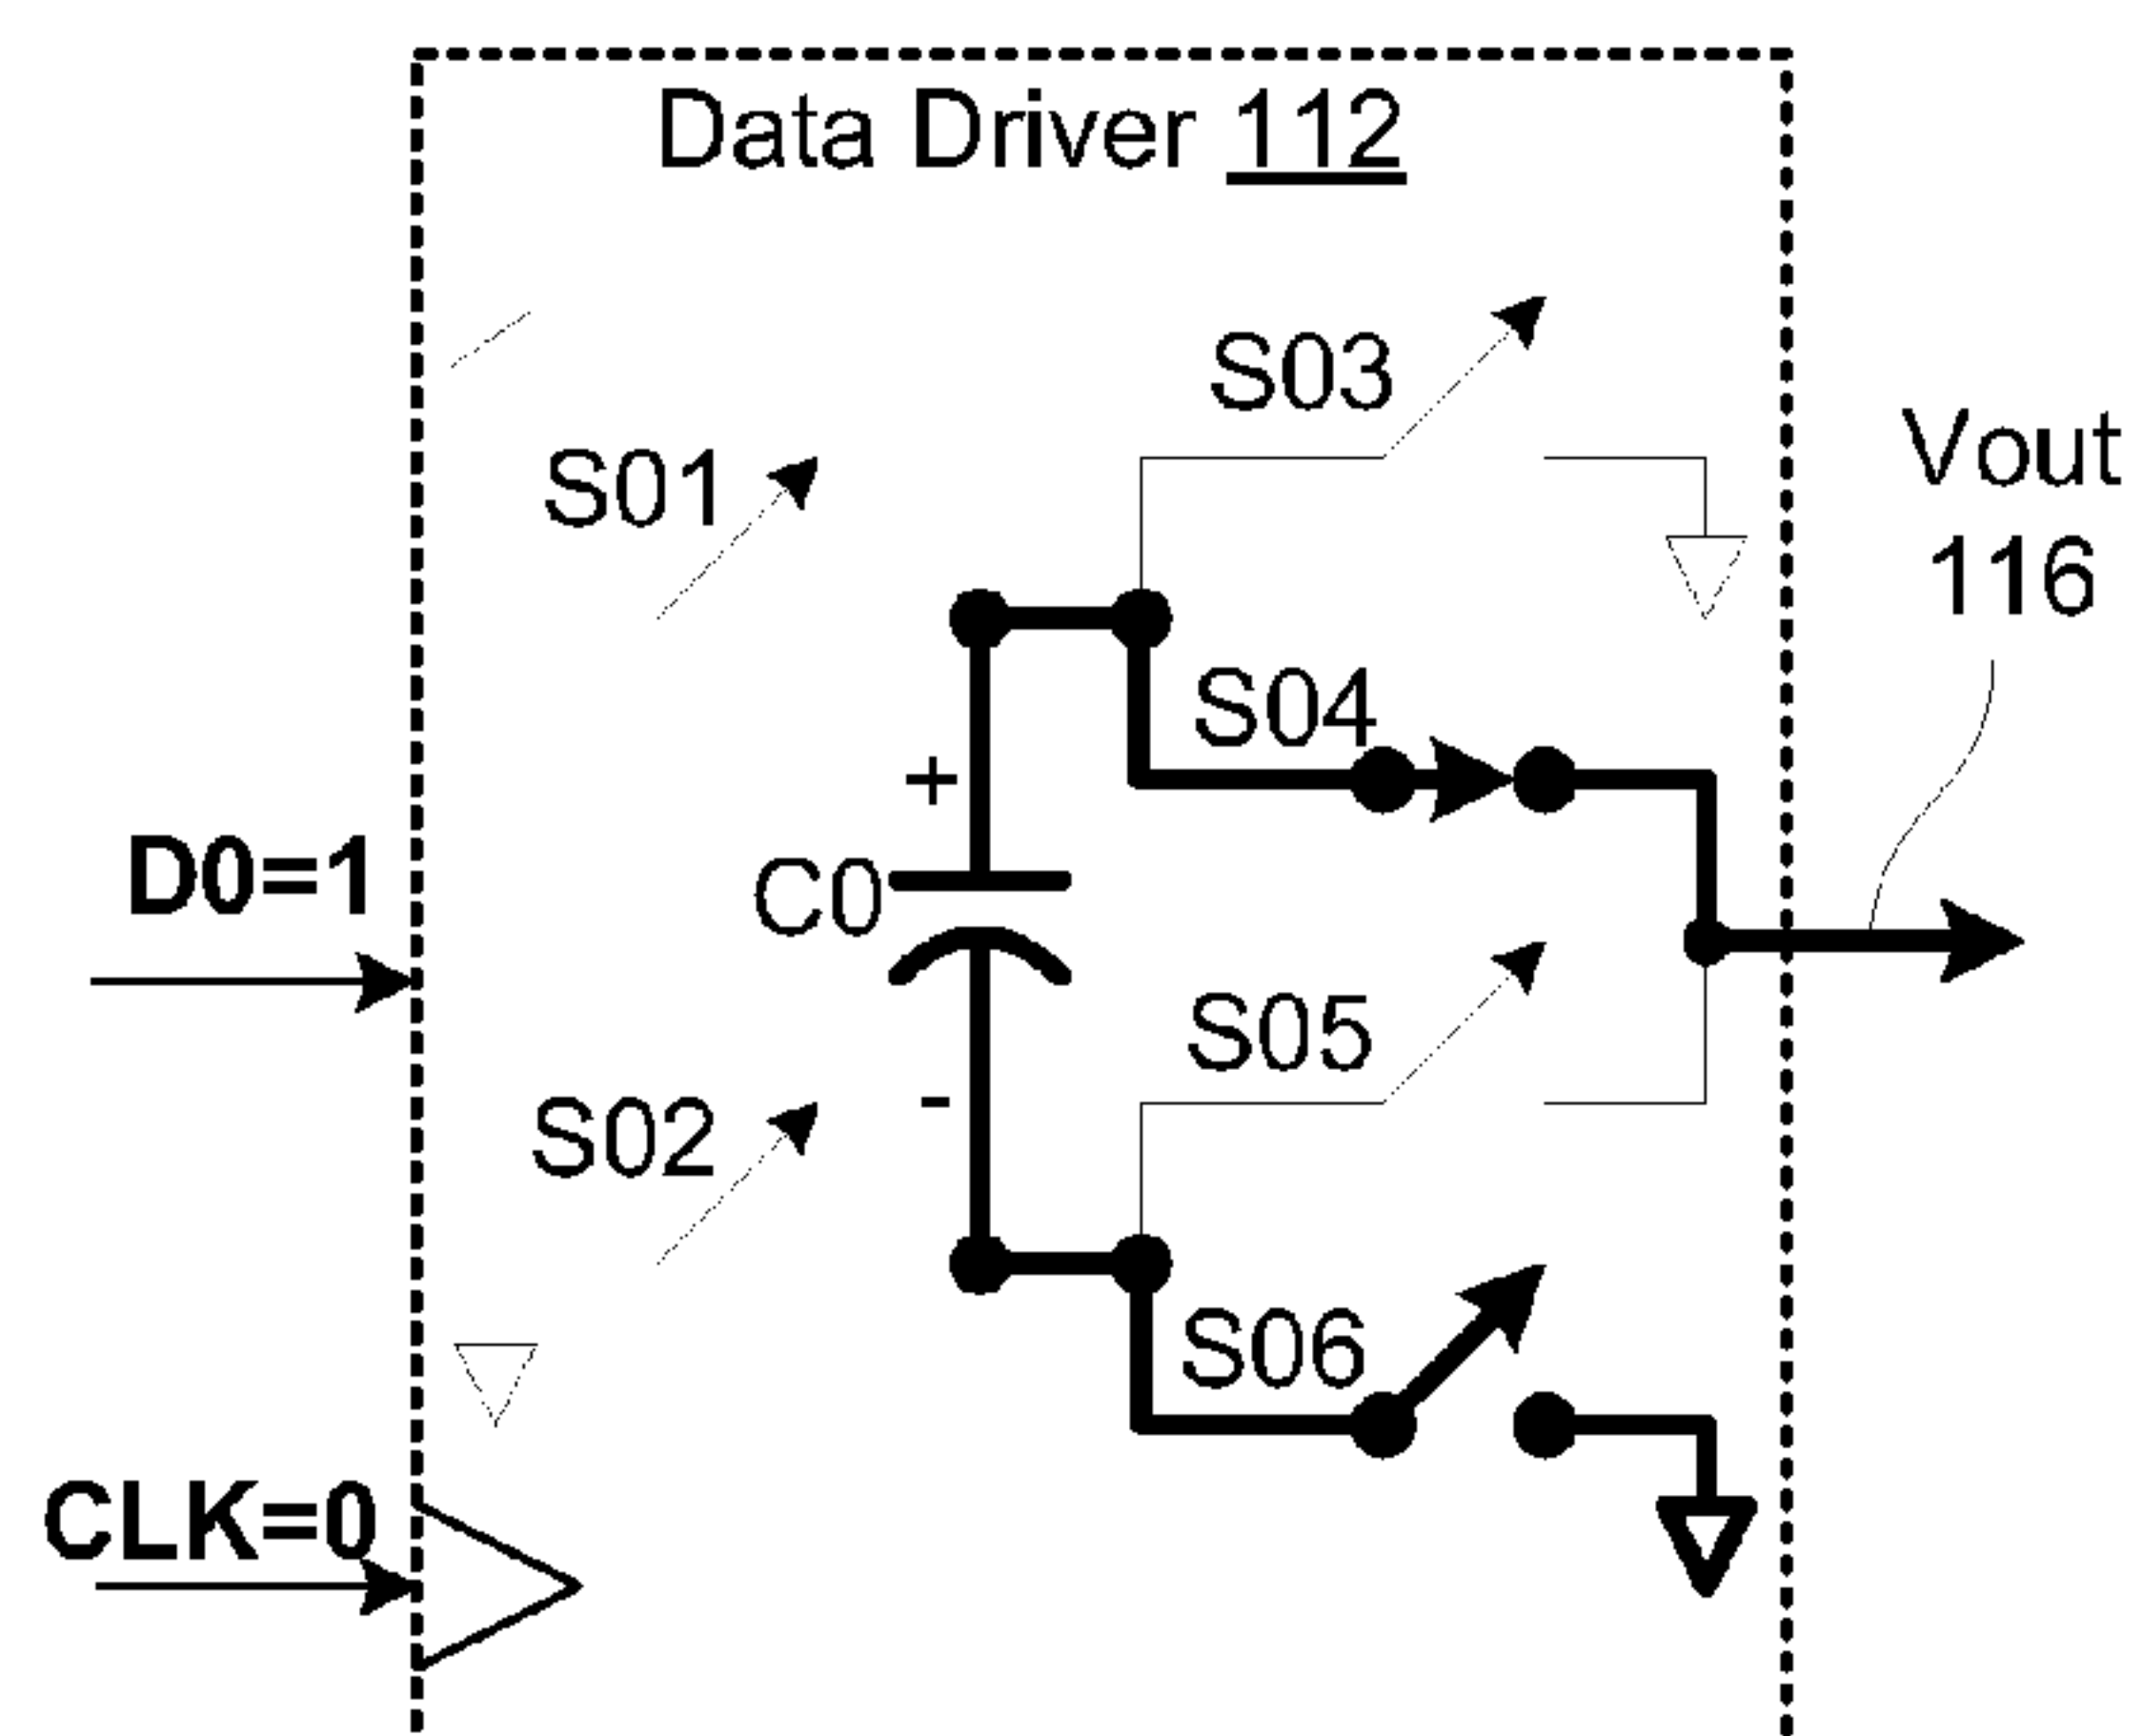
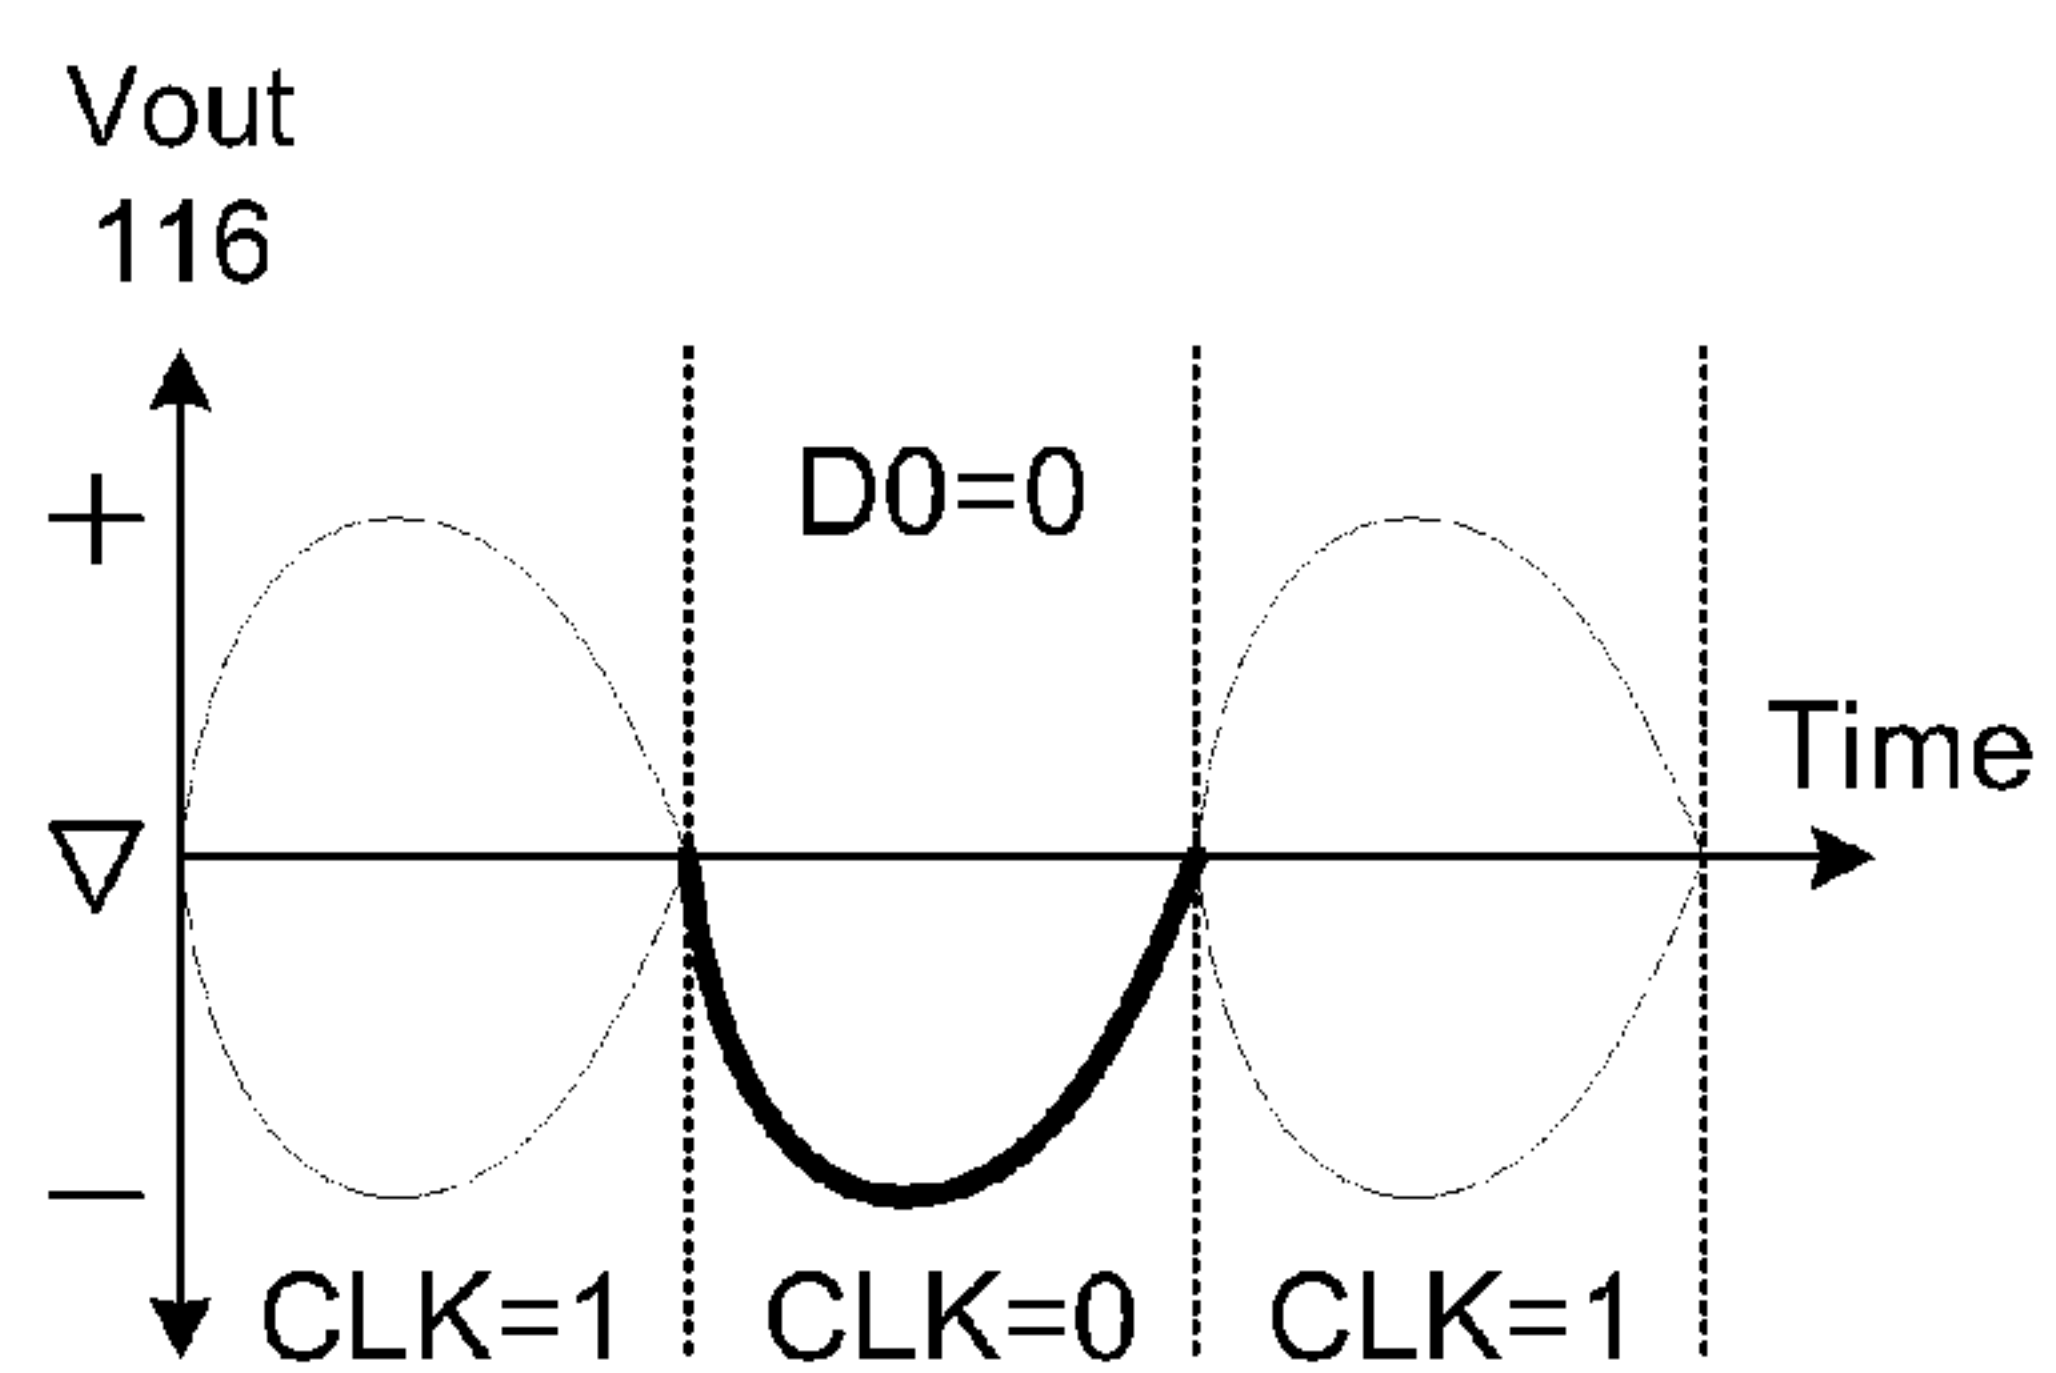
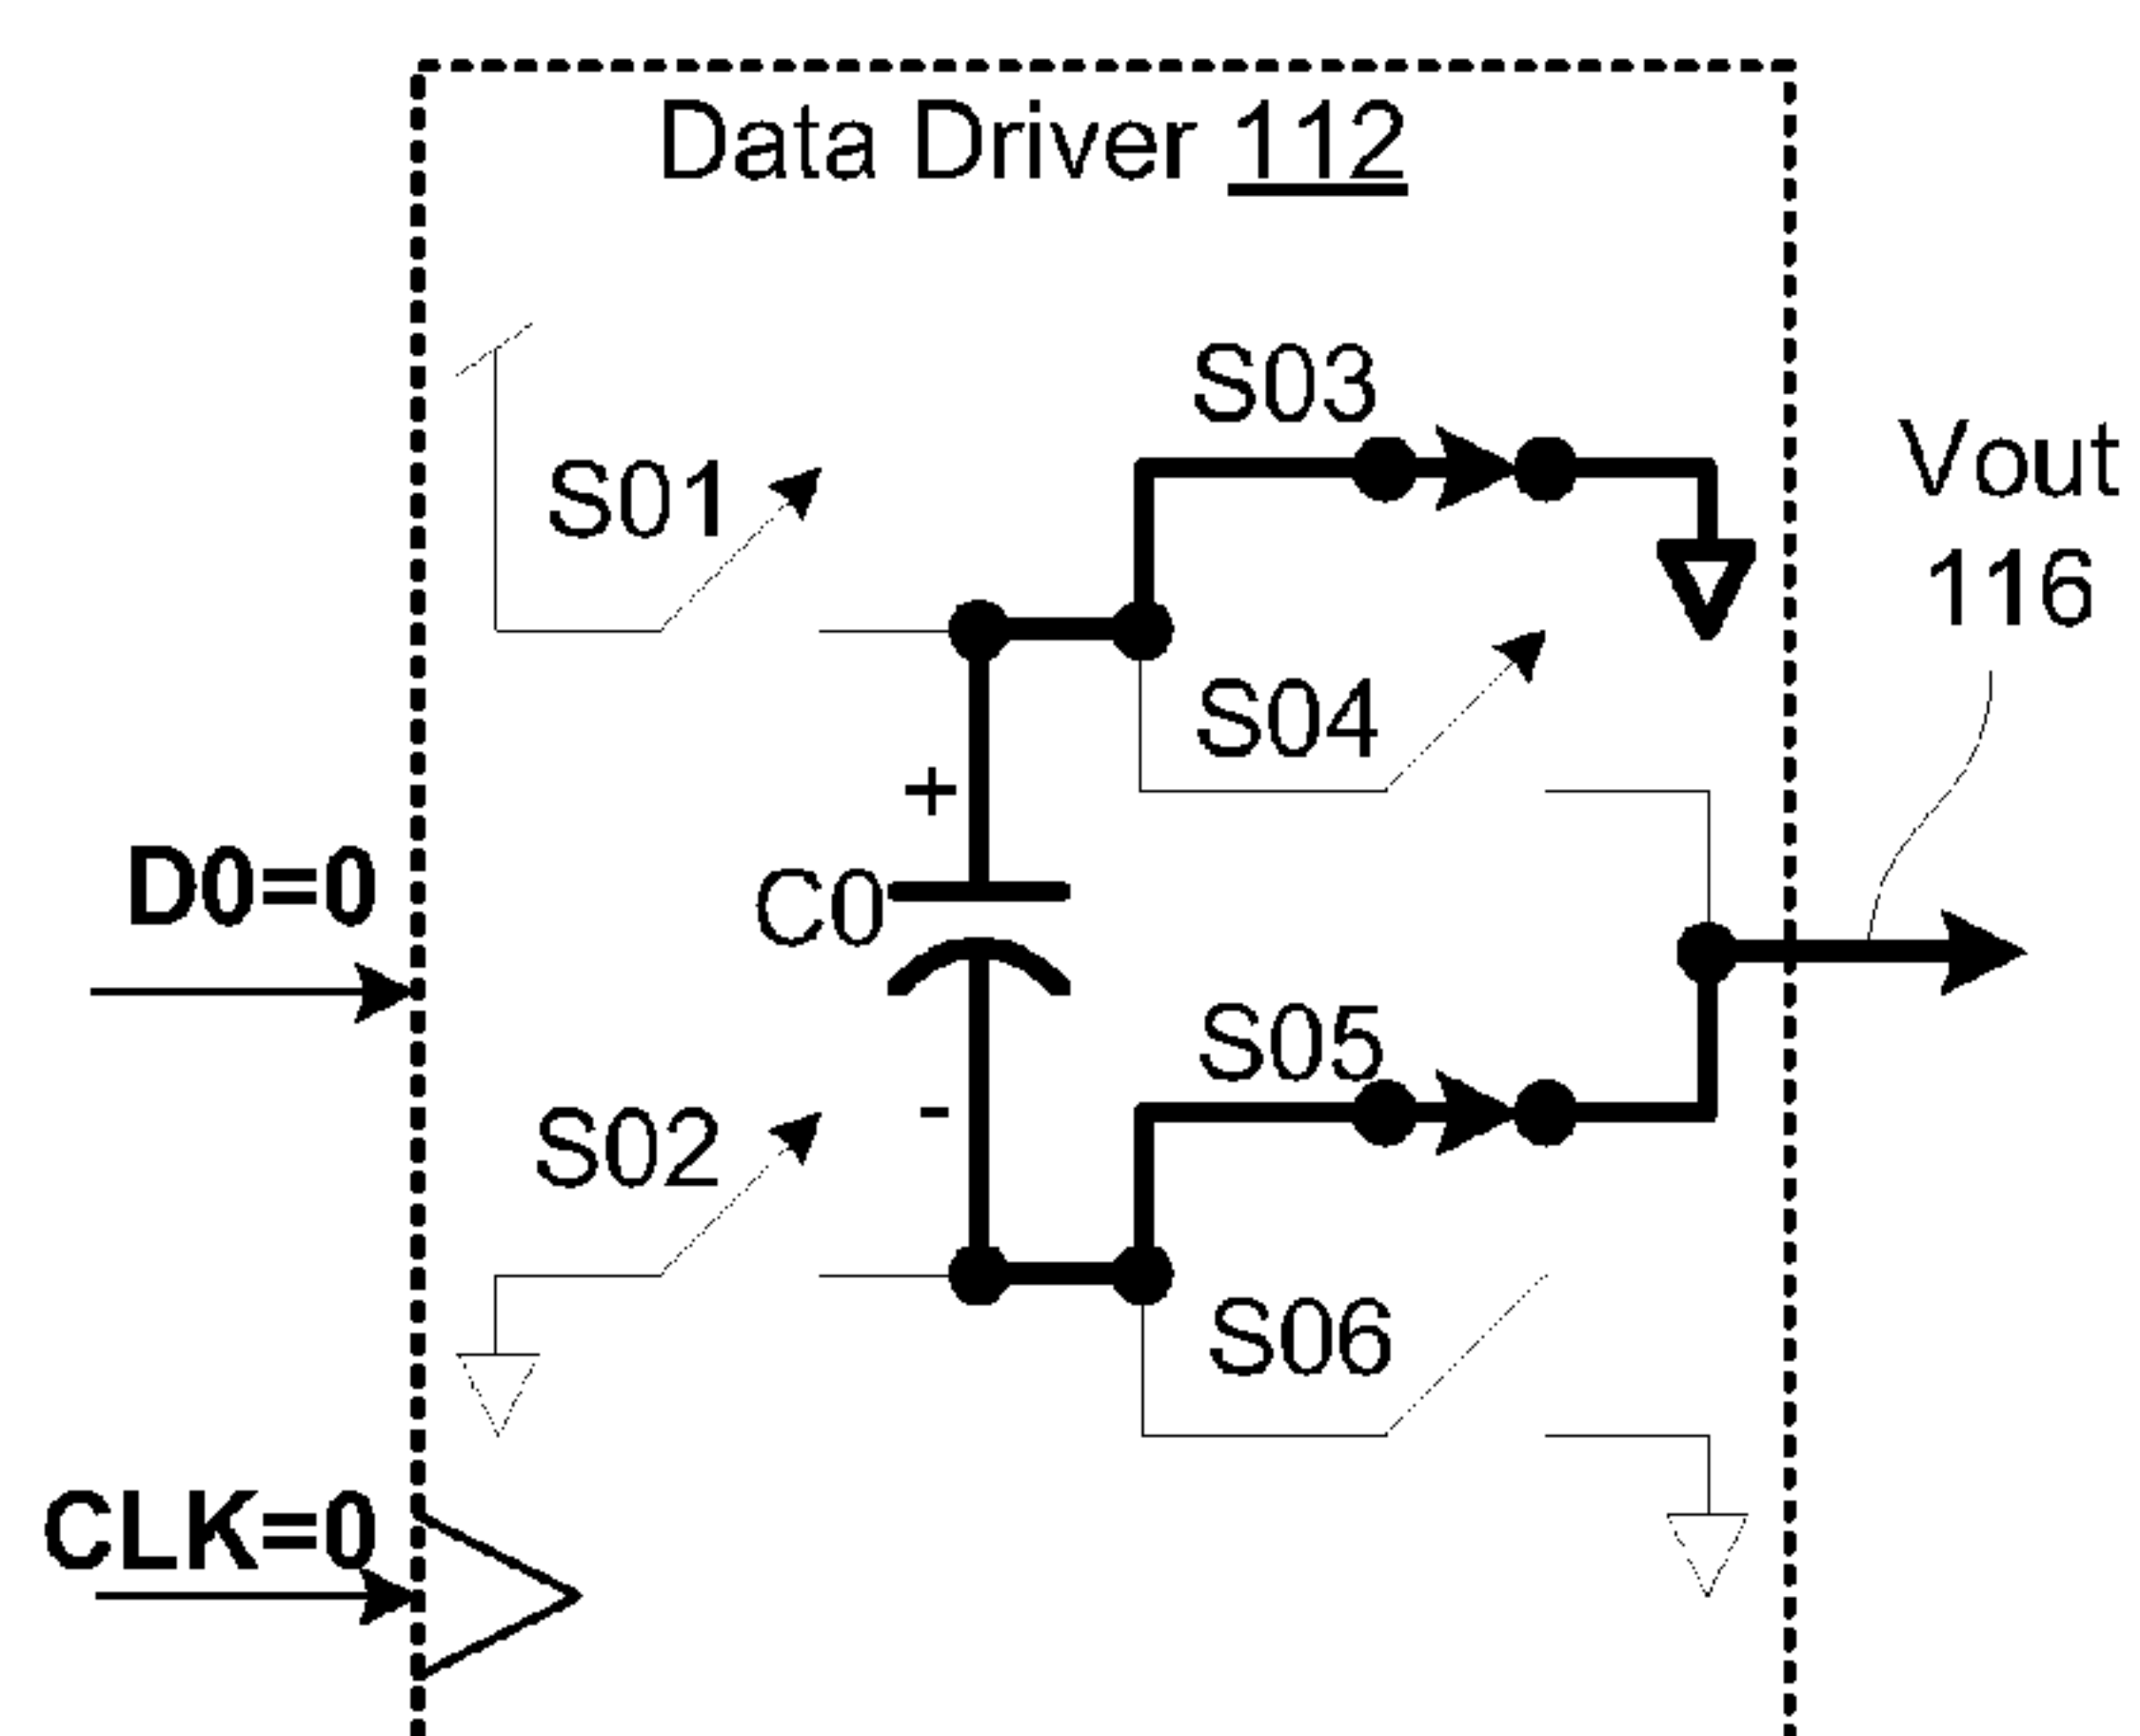
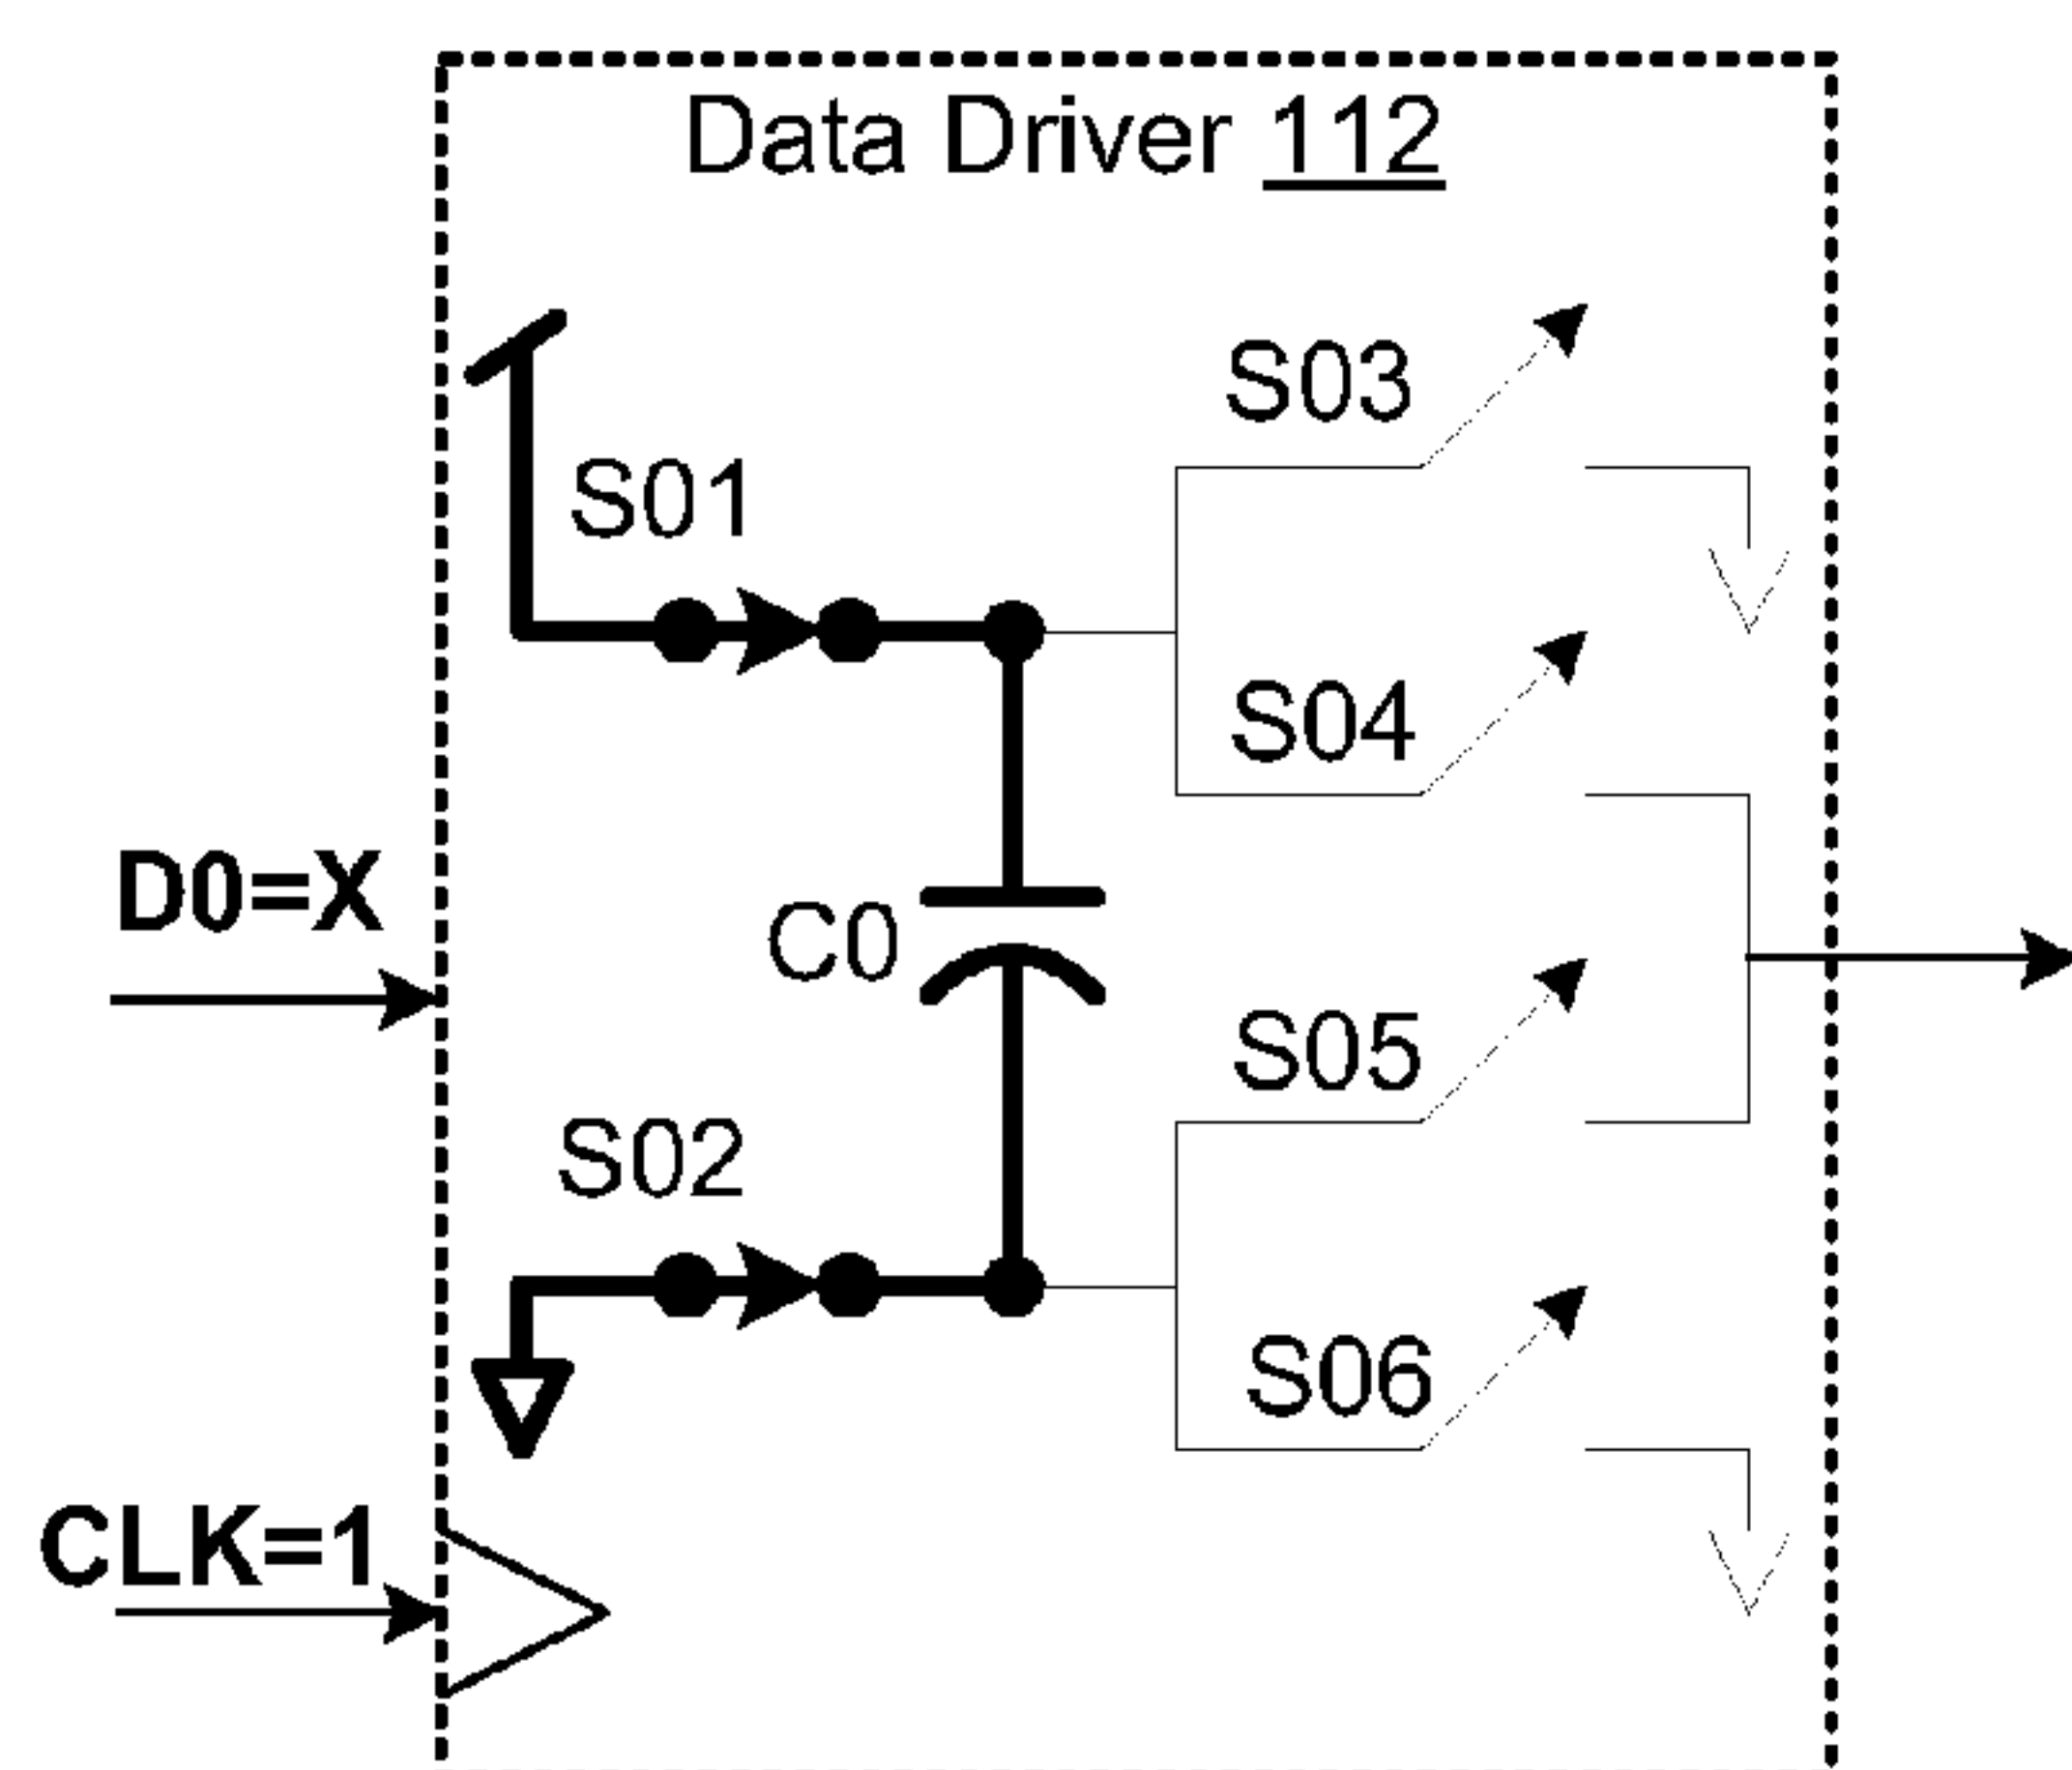


Fig. 1B

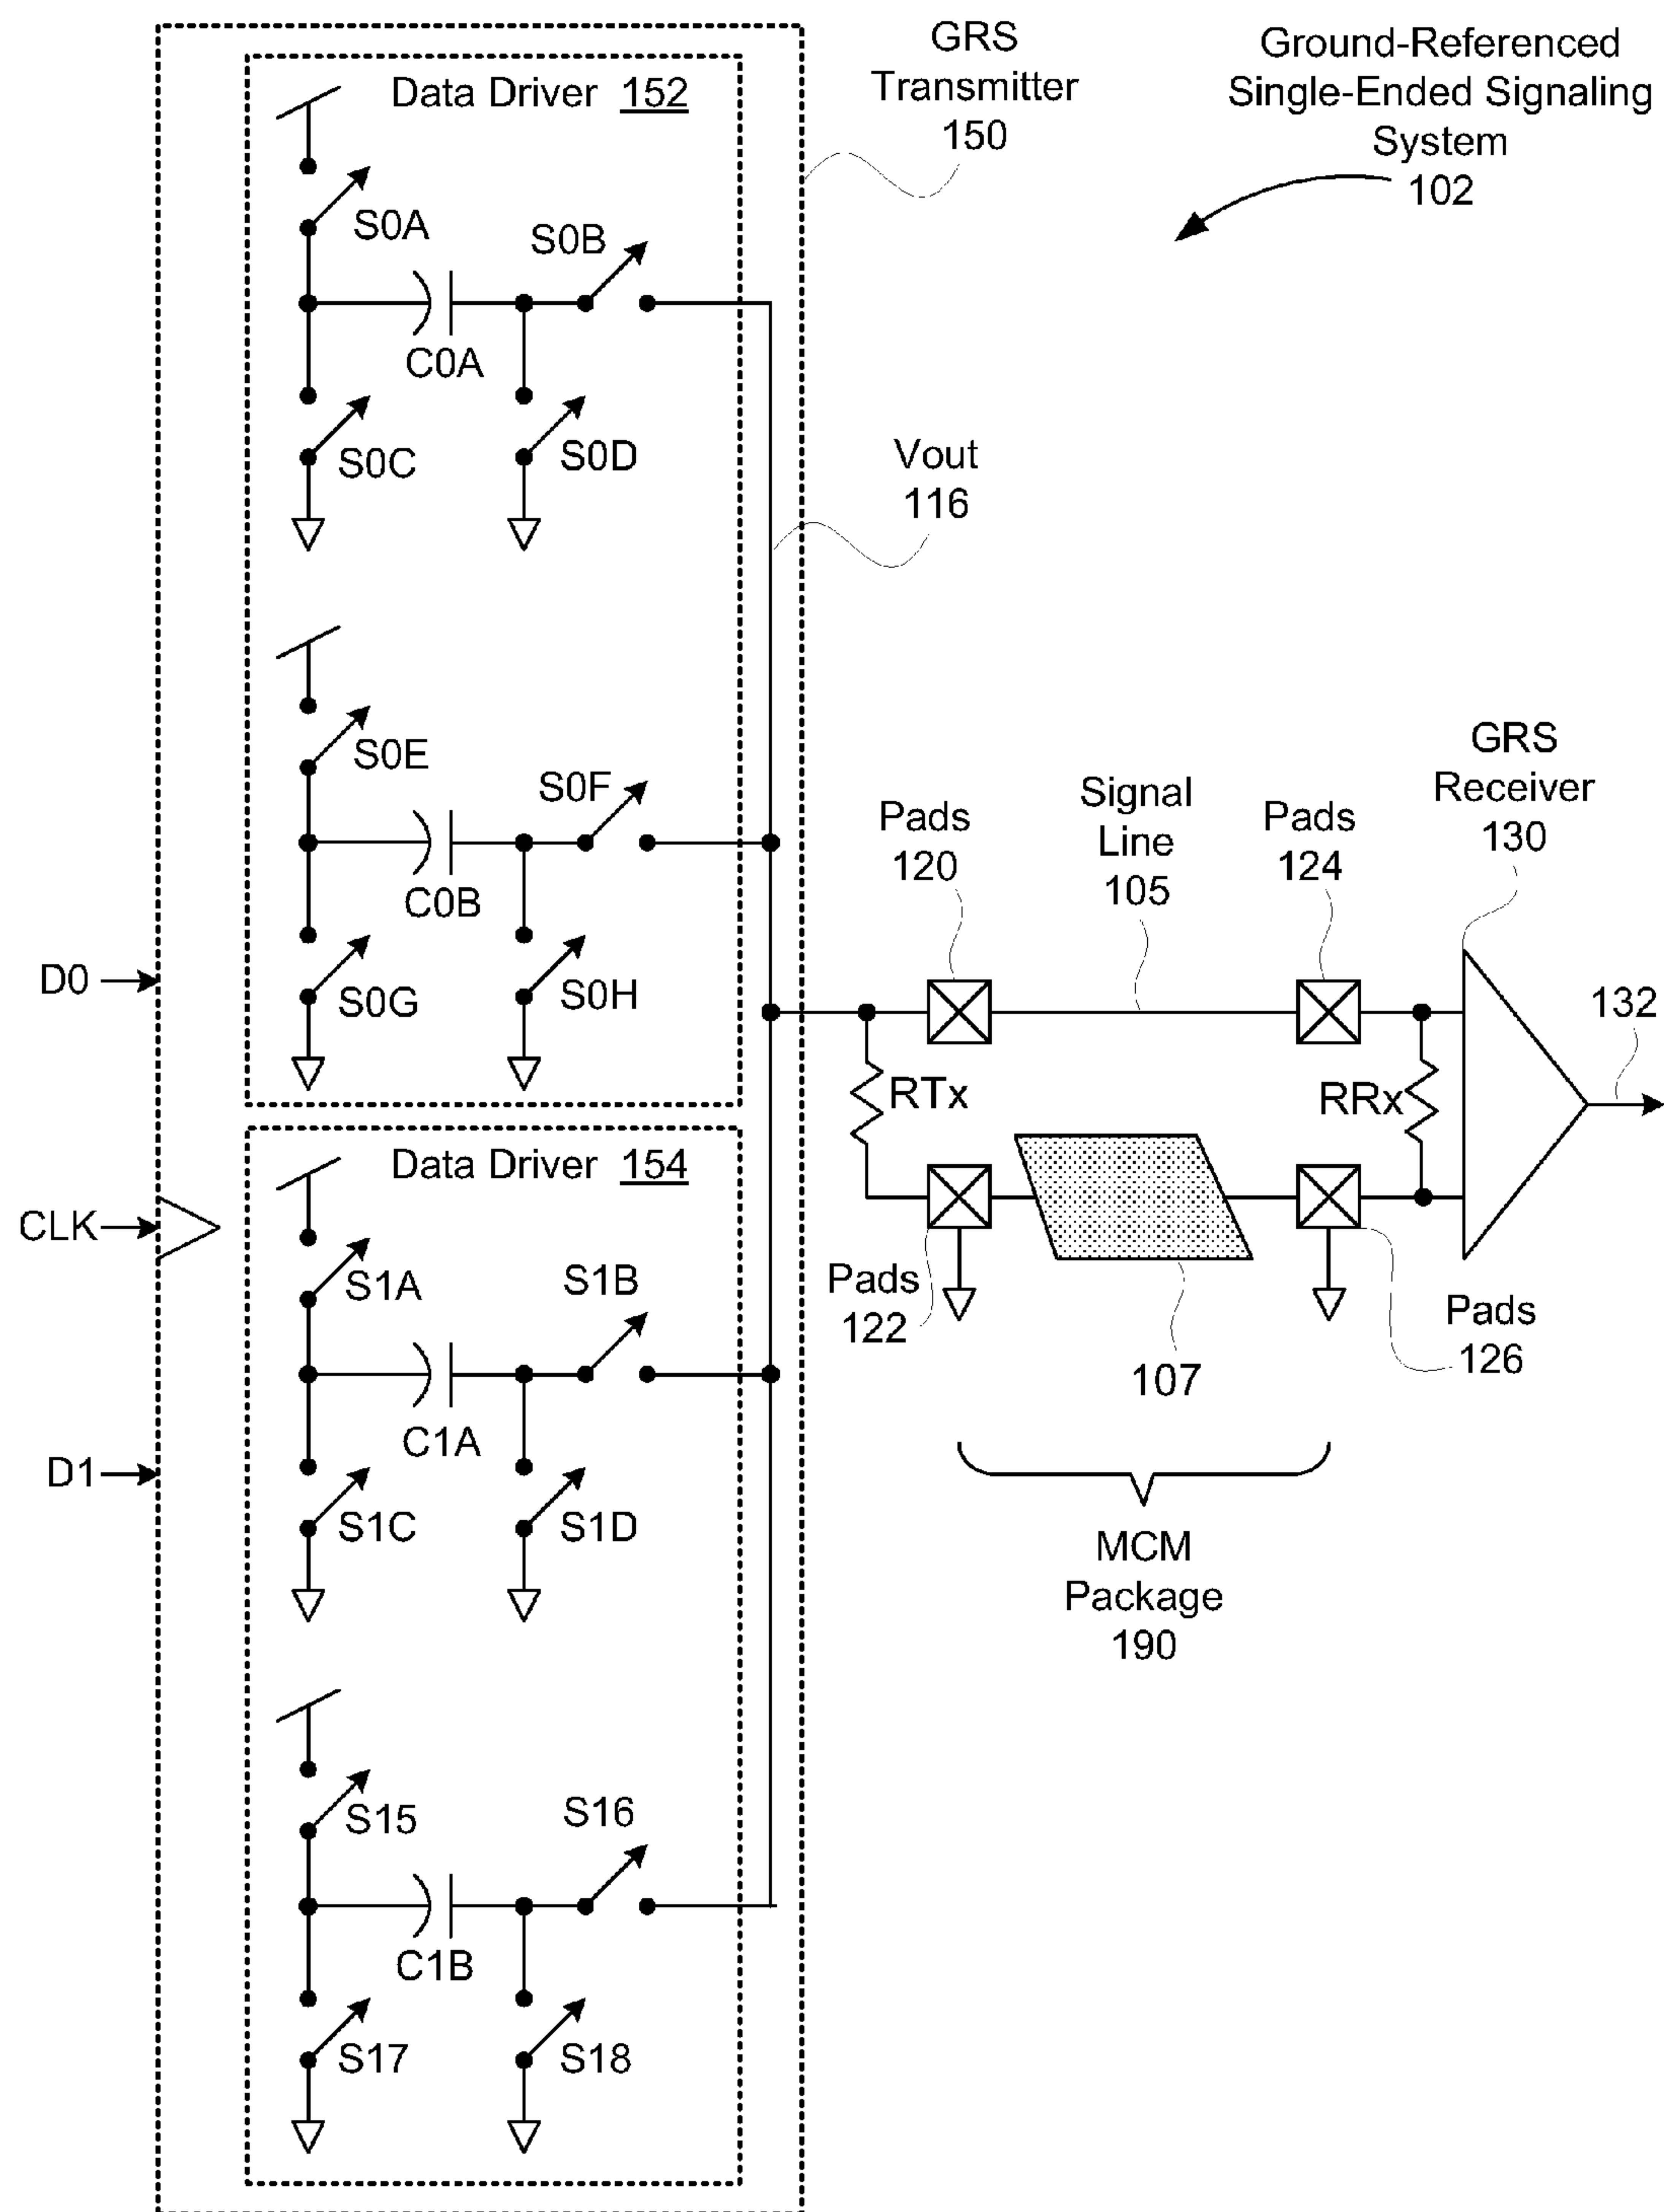
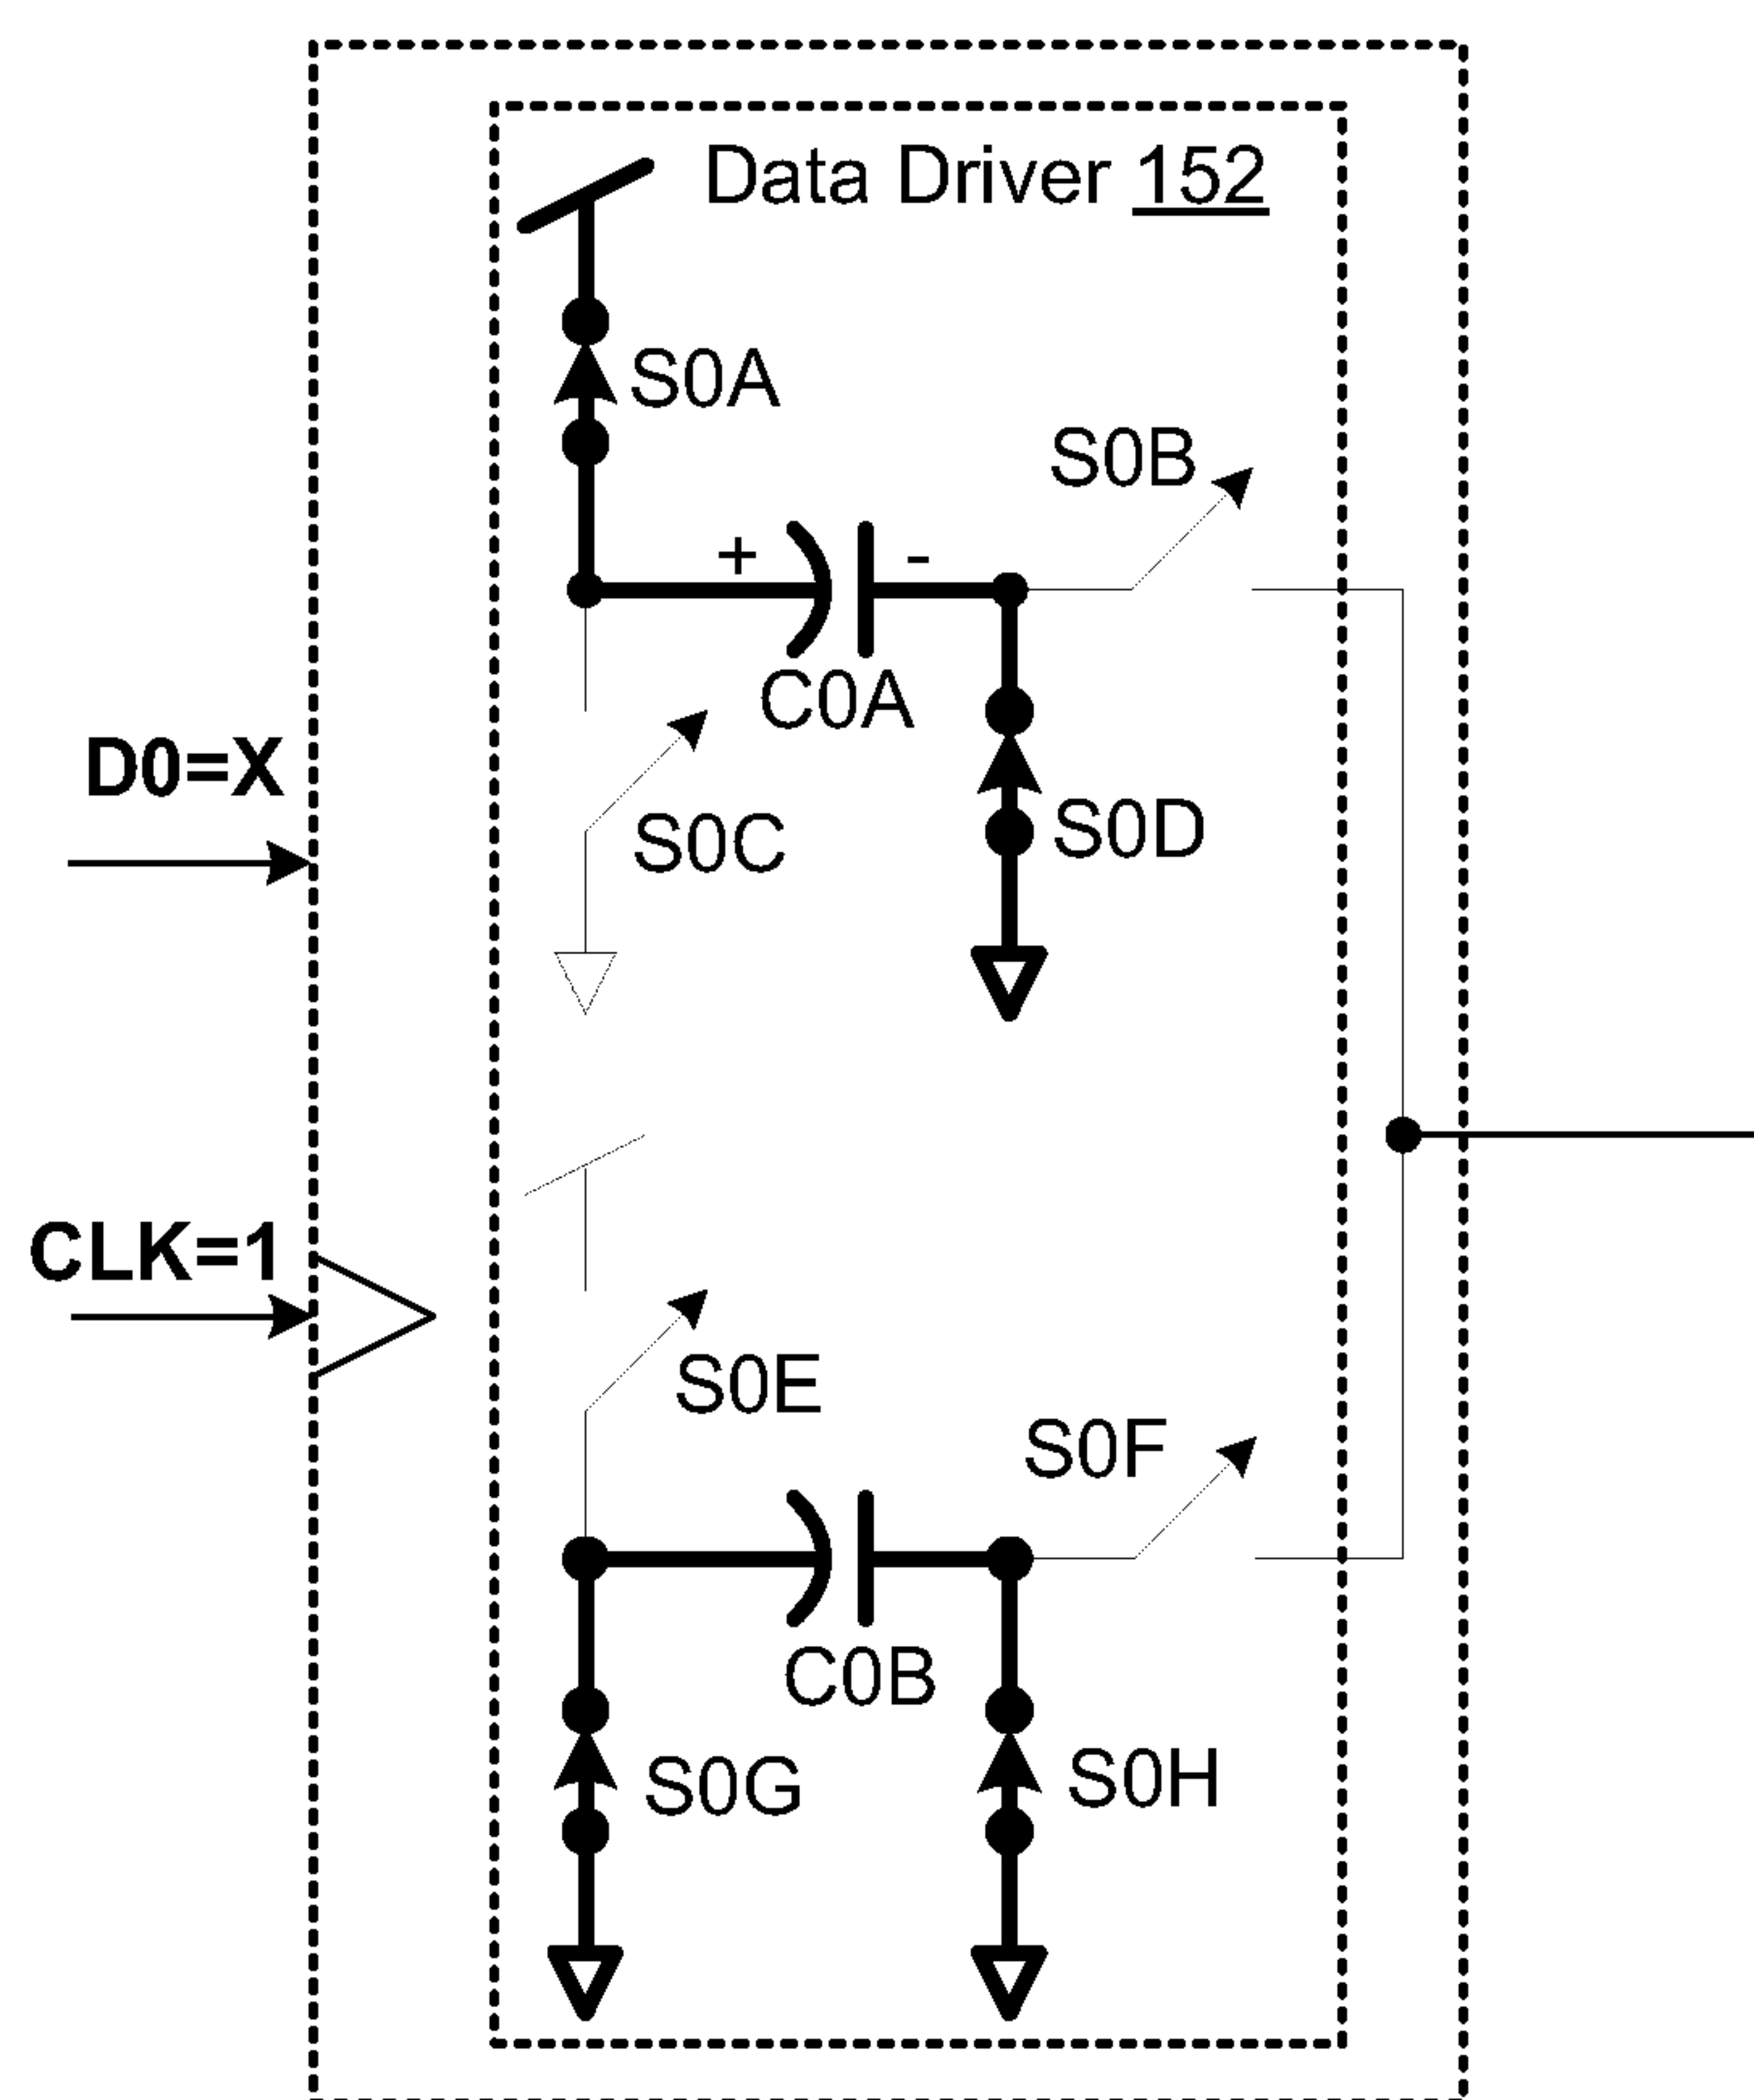


Fig. 1C



*Fig. 1D*

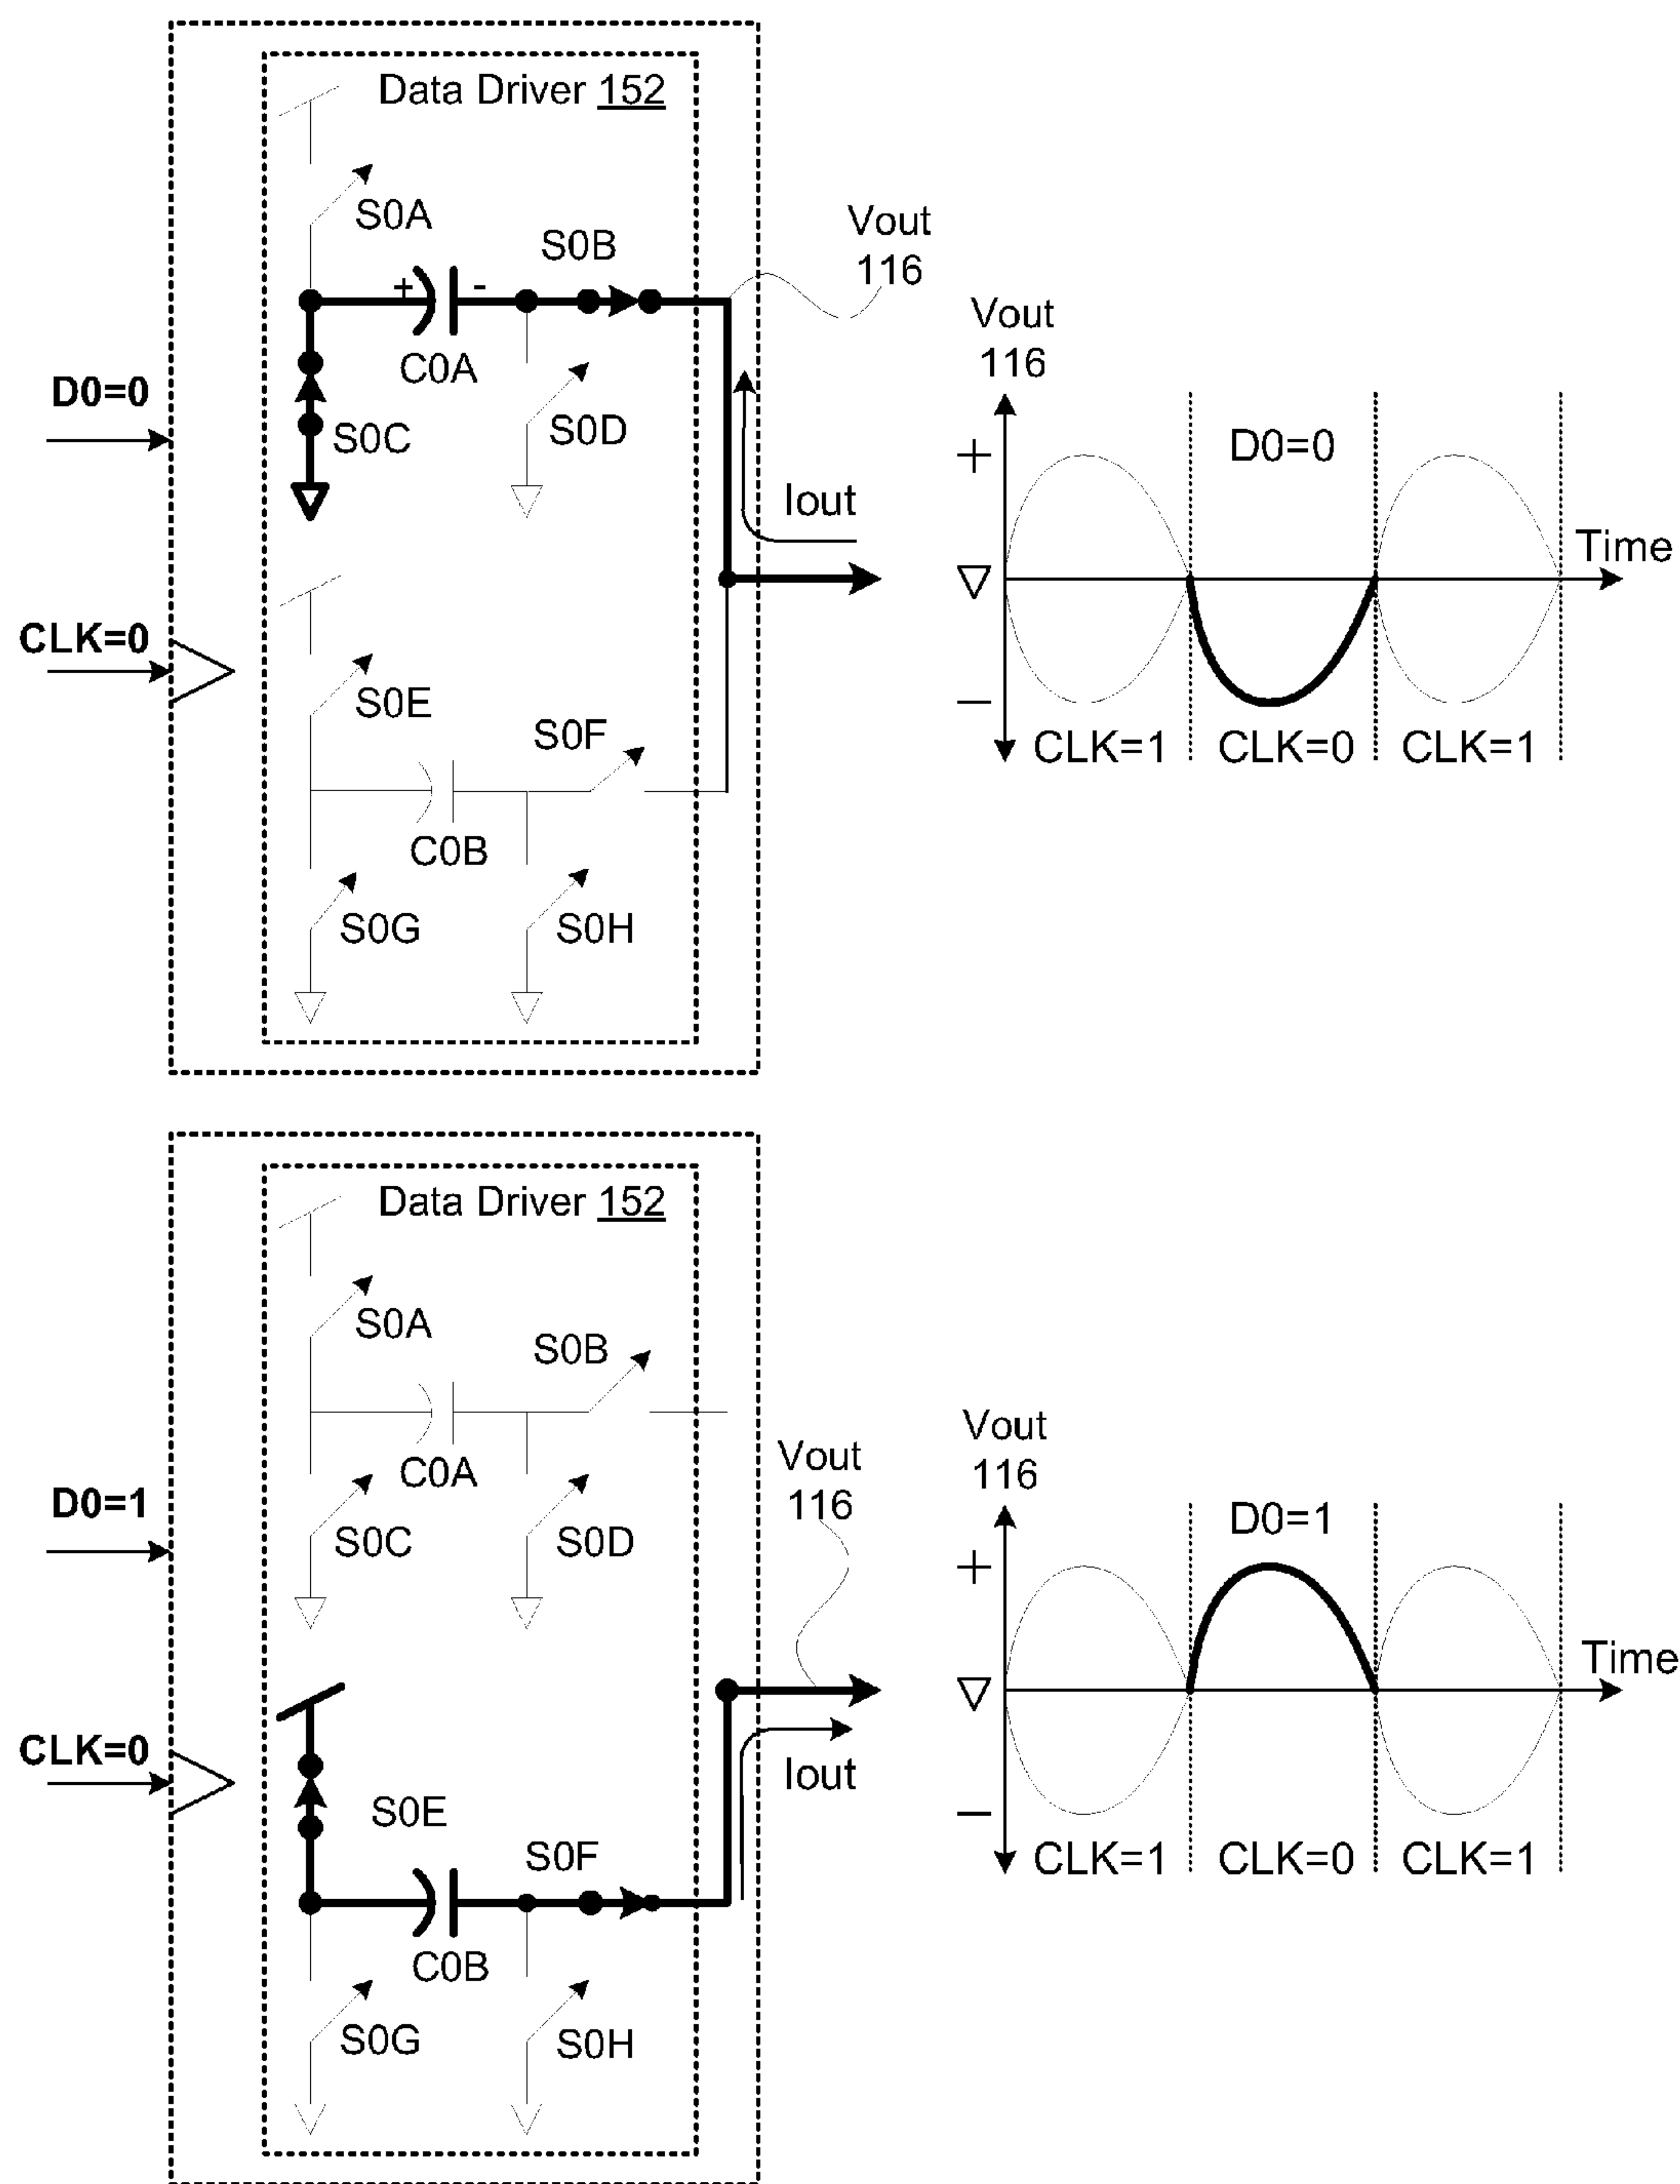
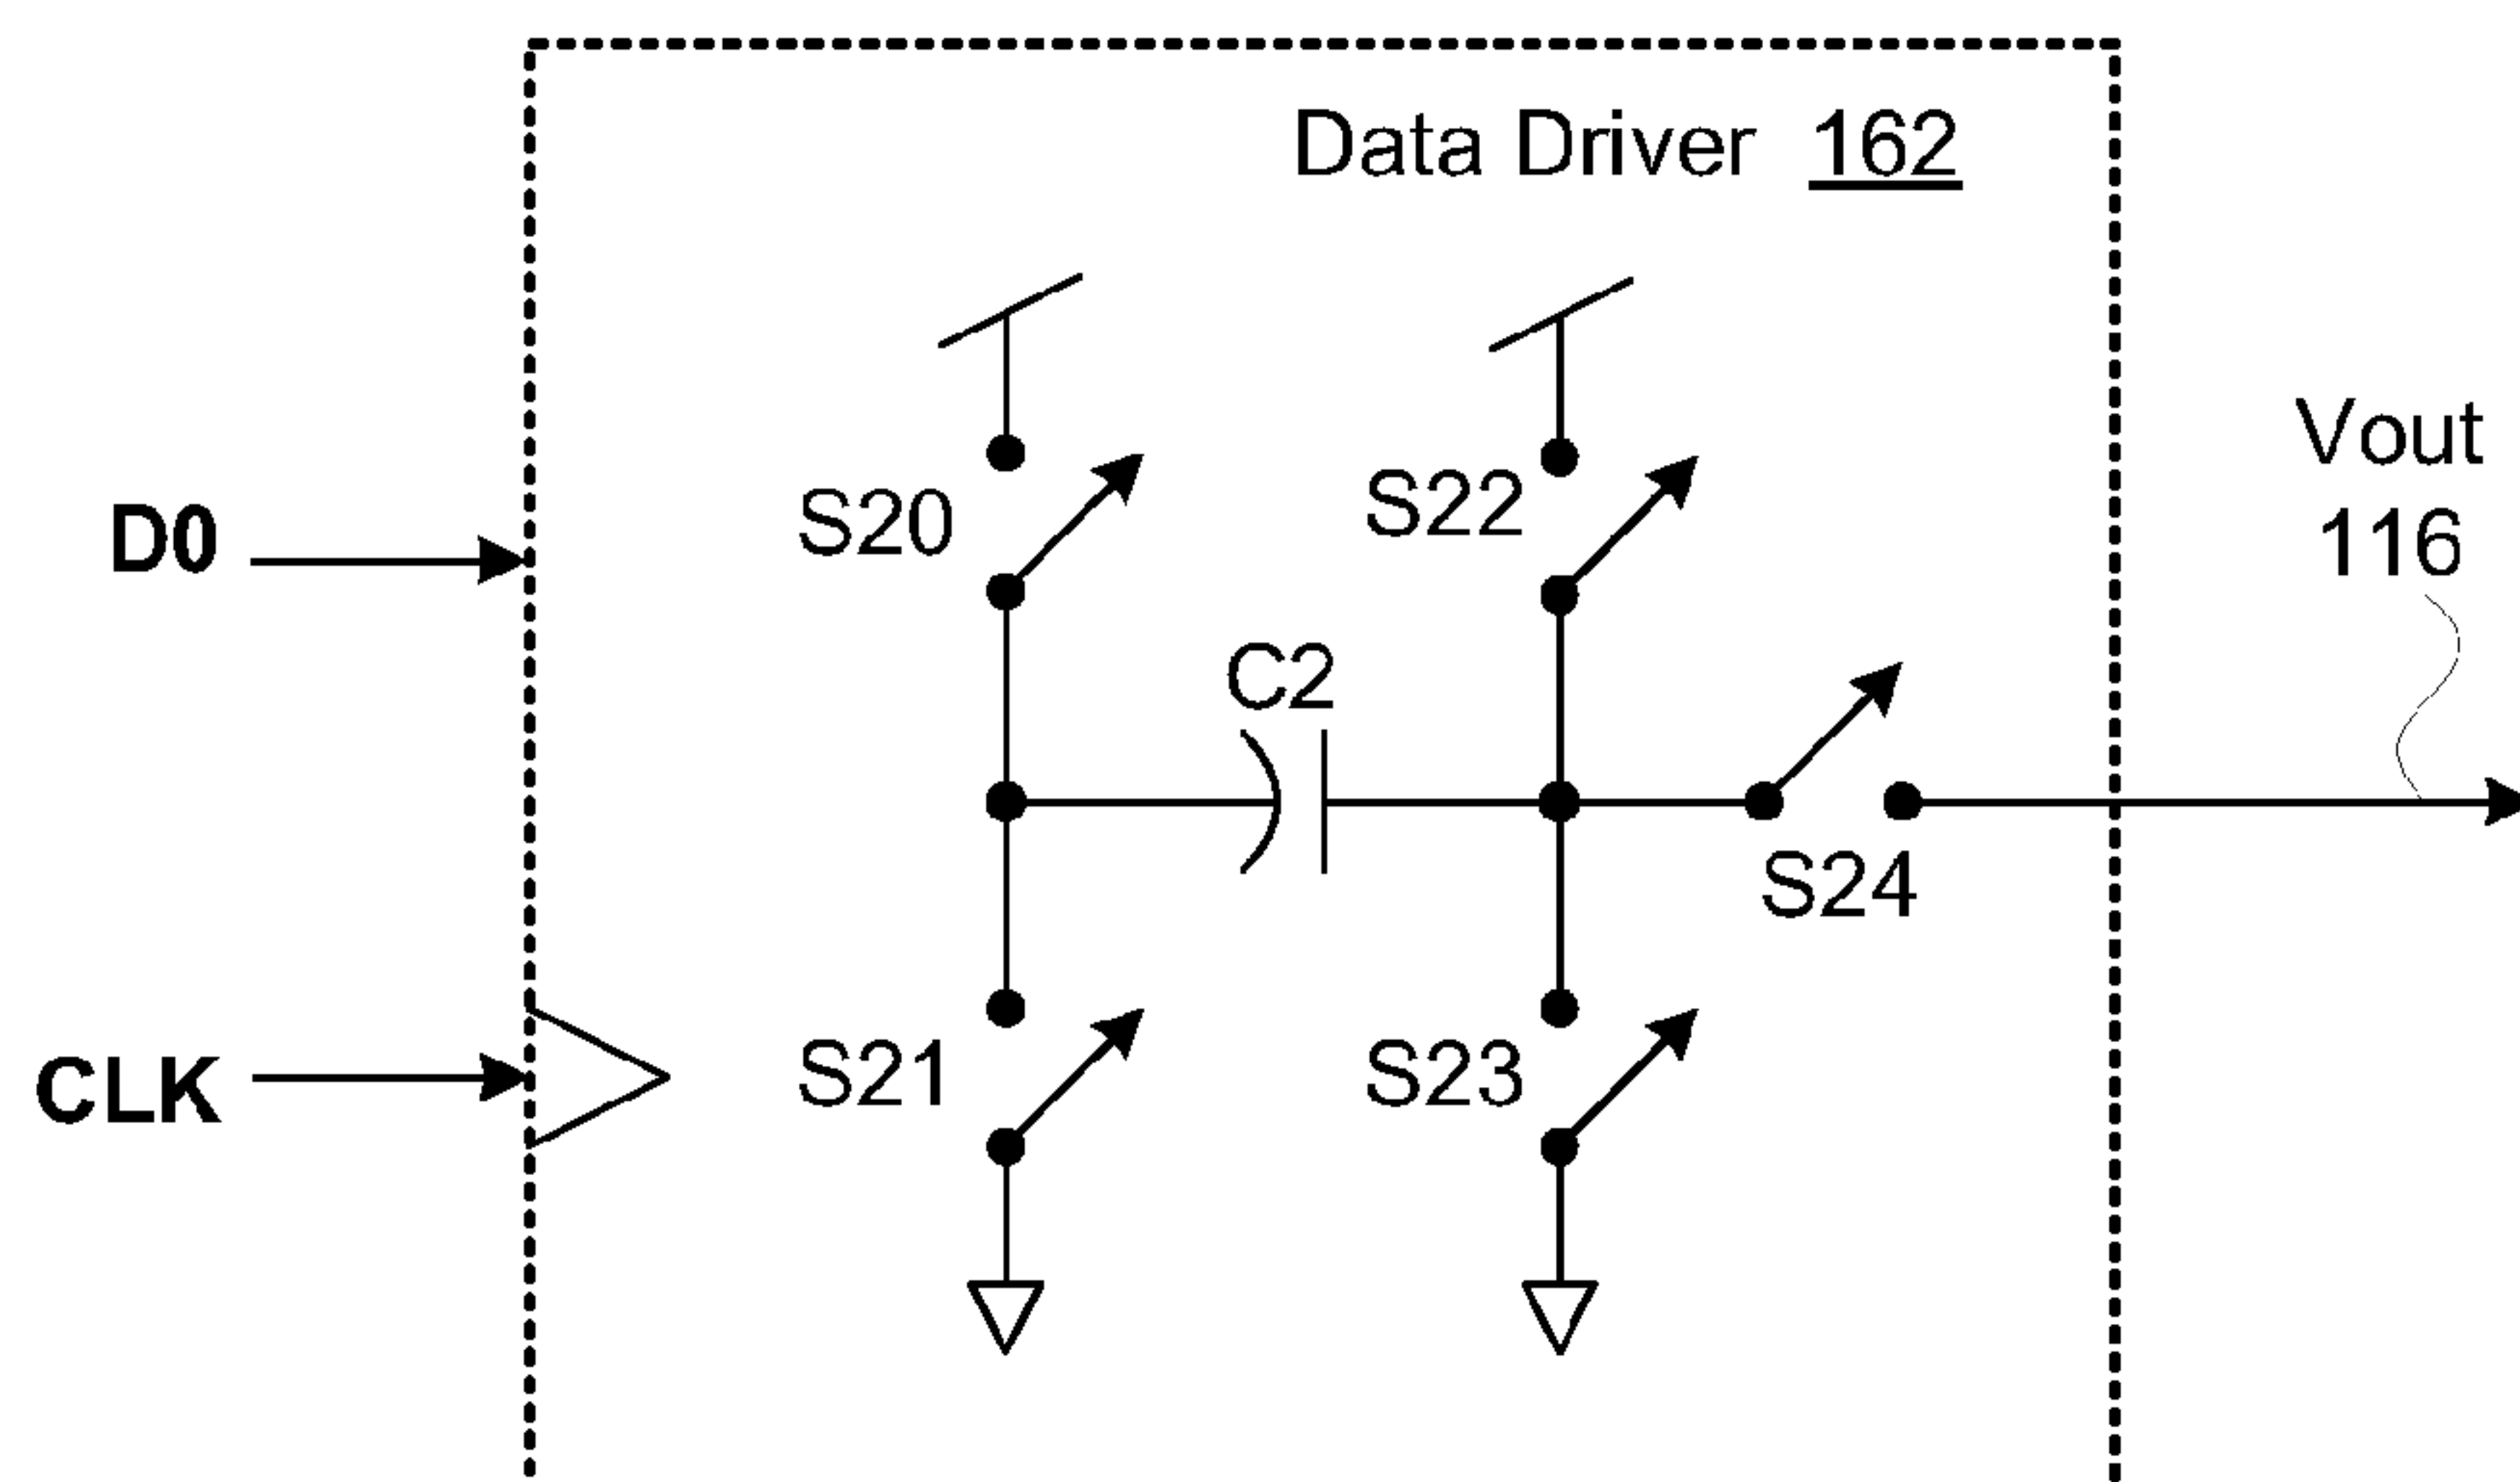
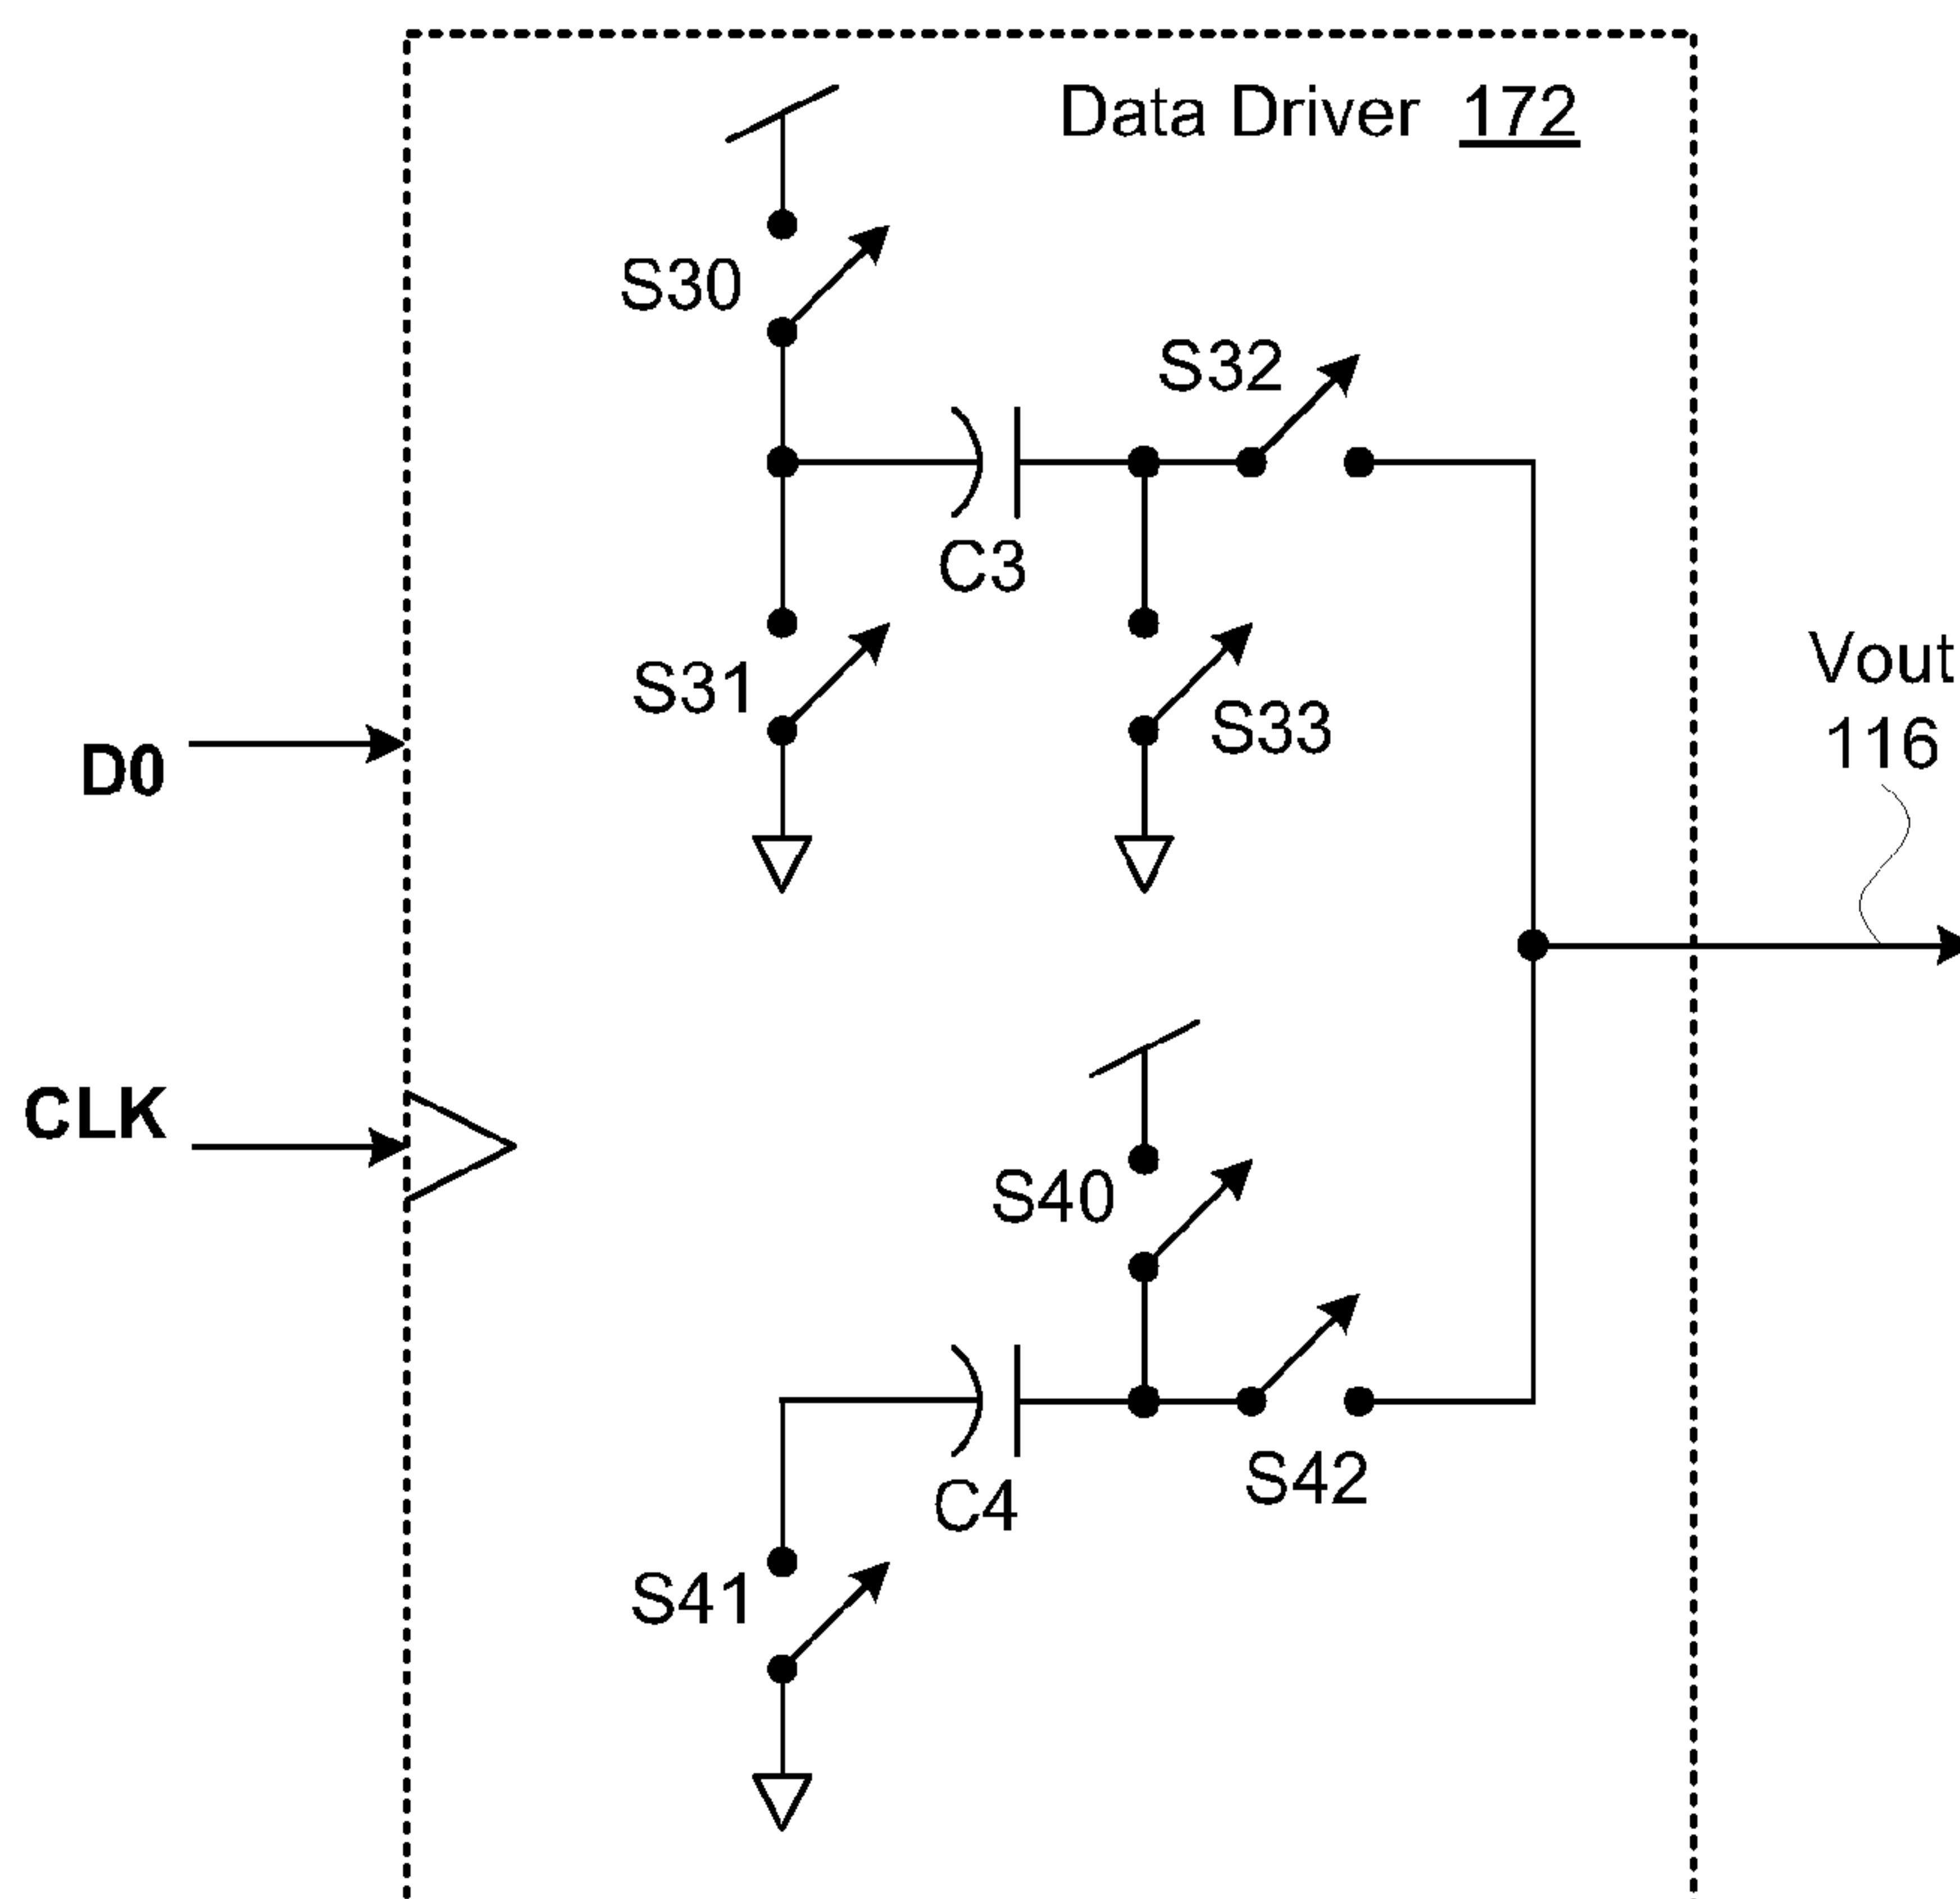


Fig. 1E

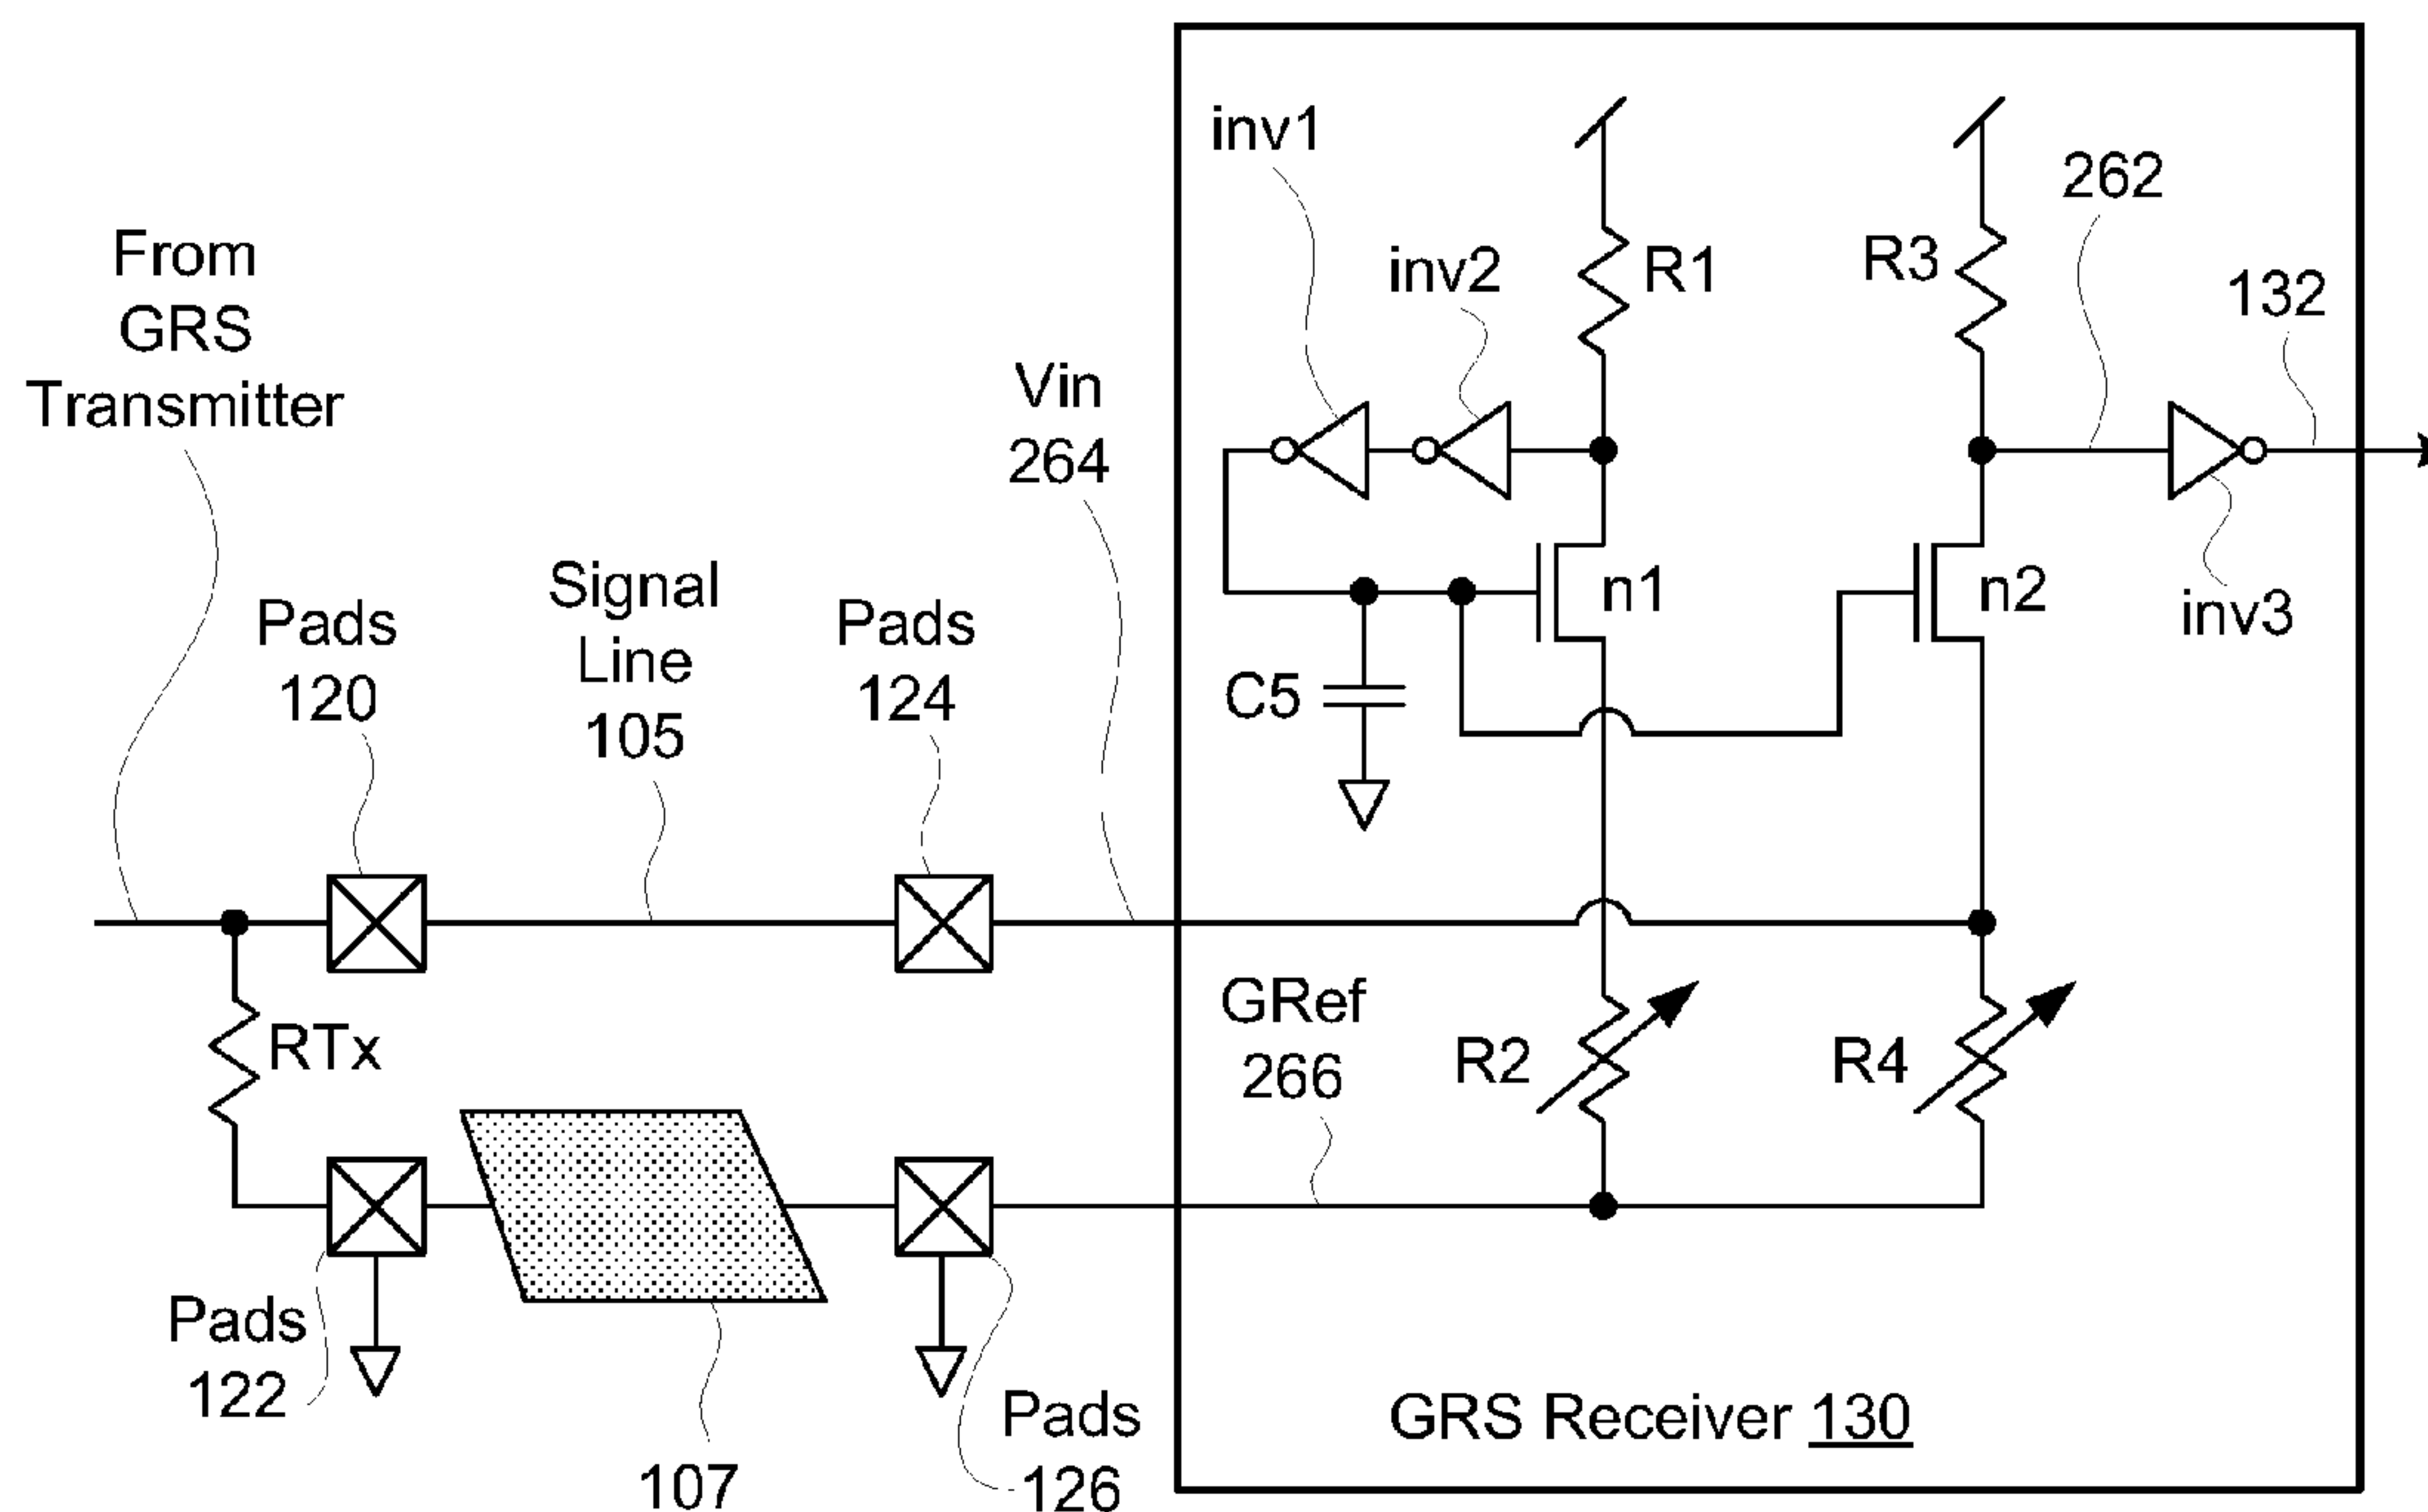




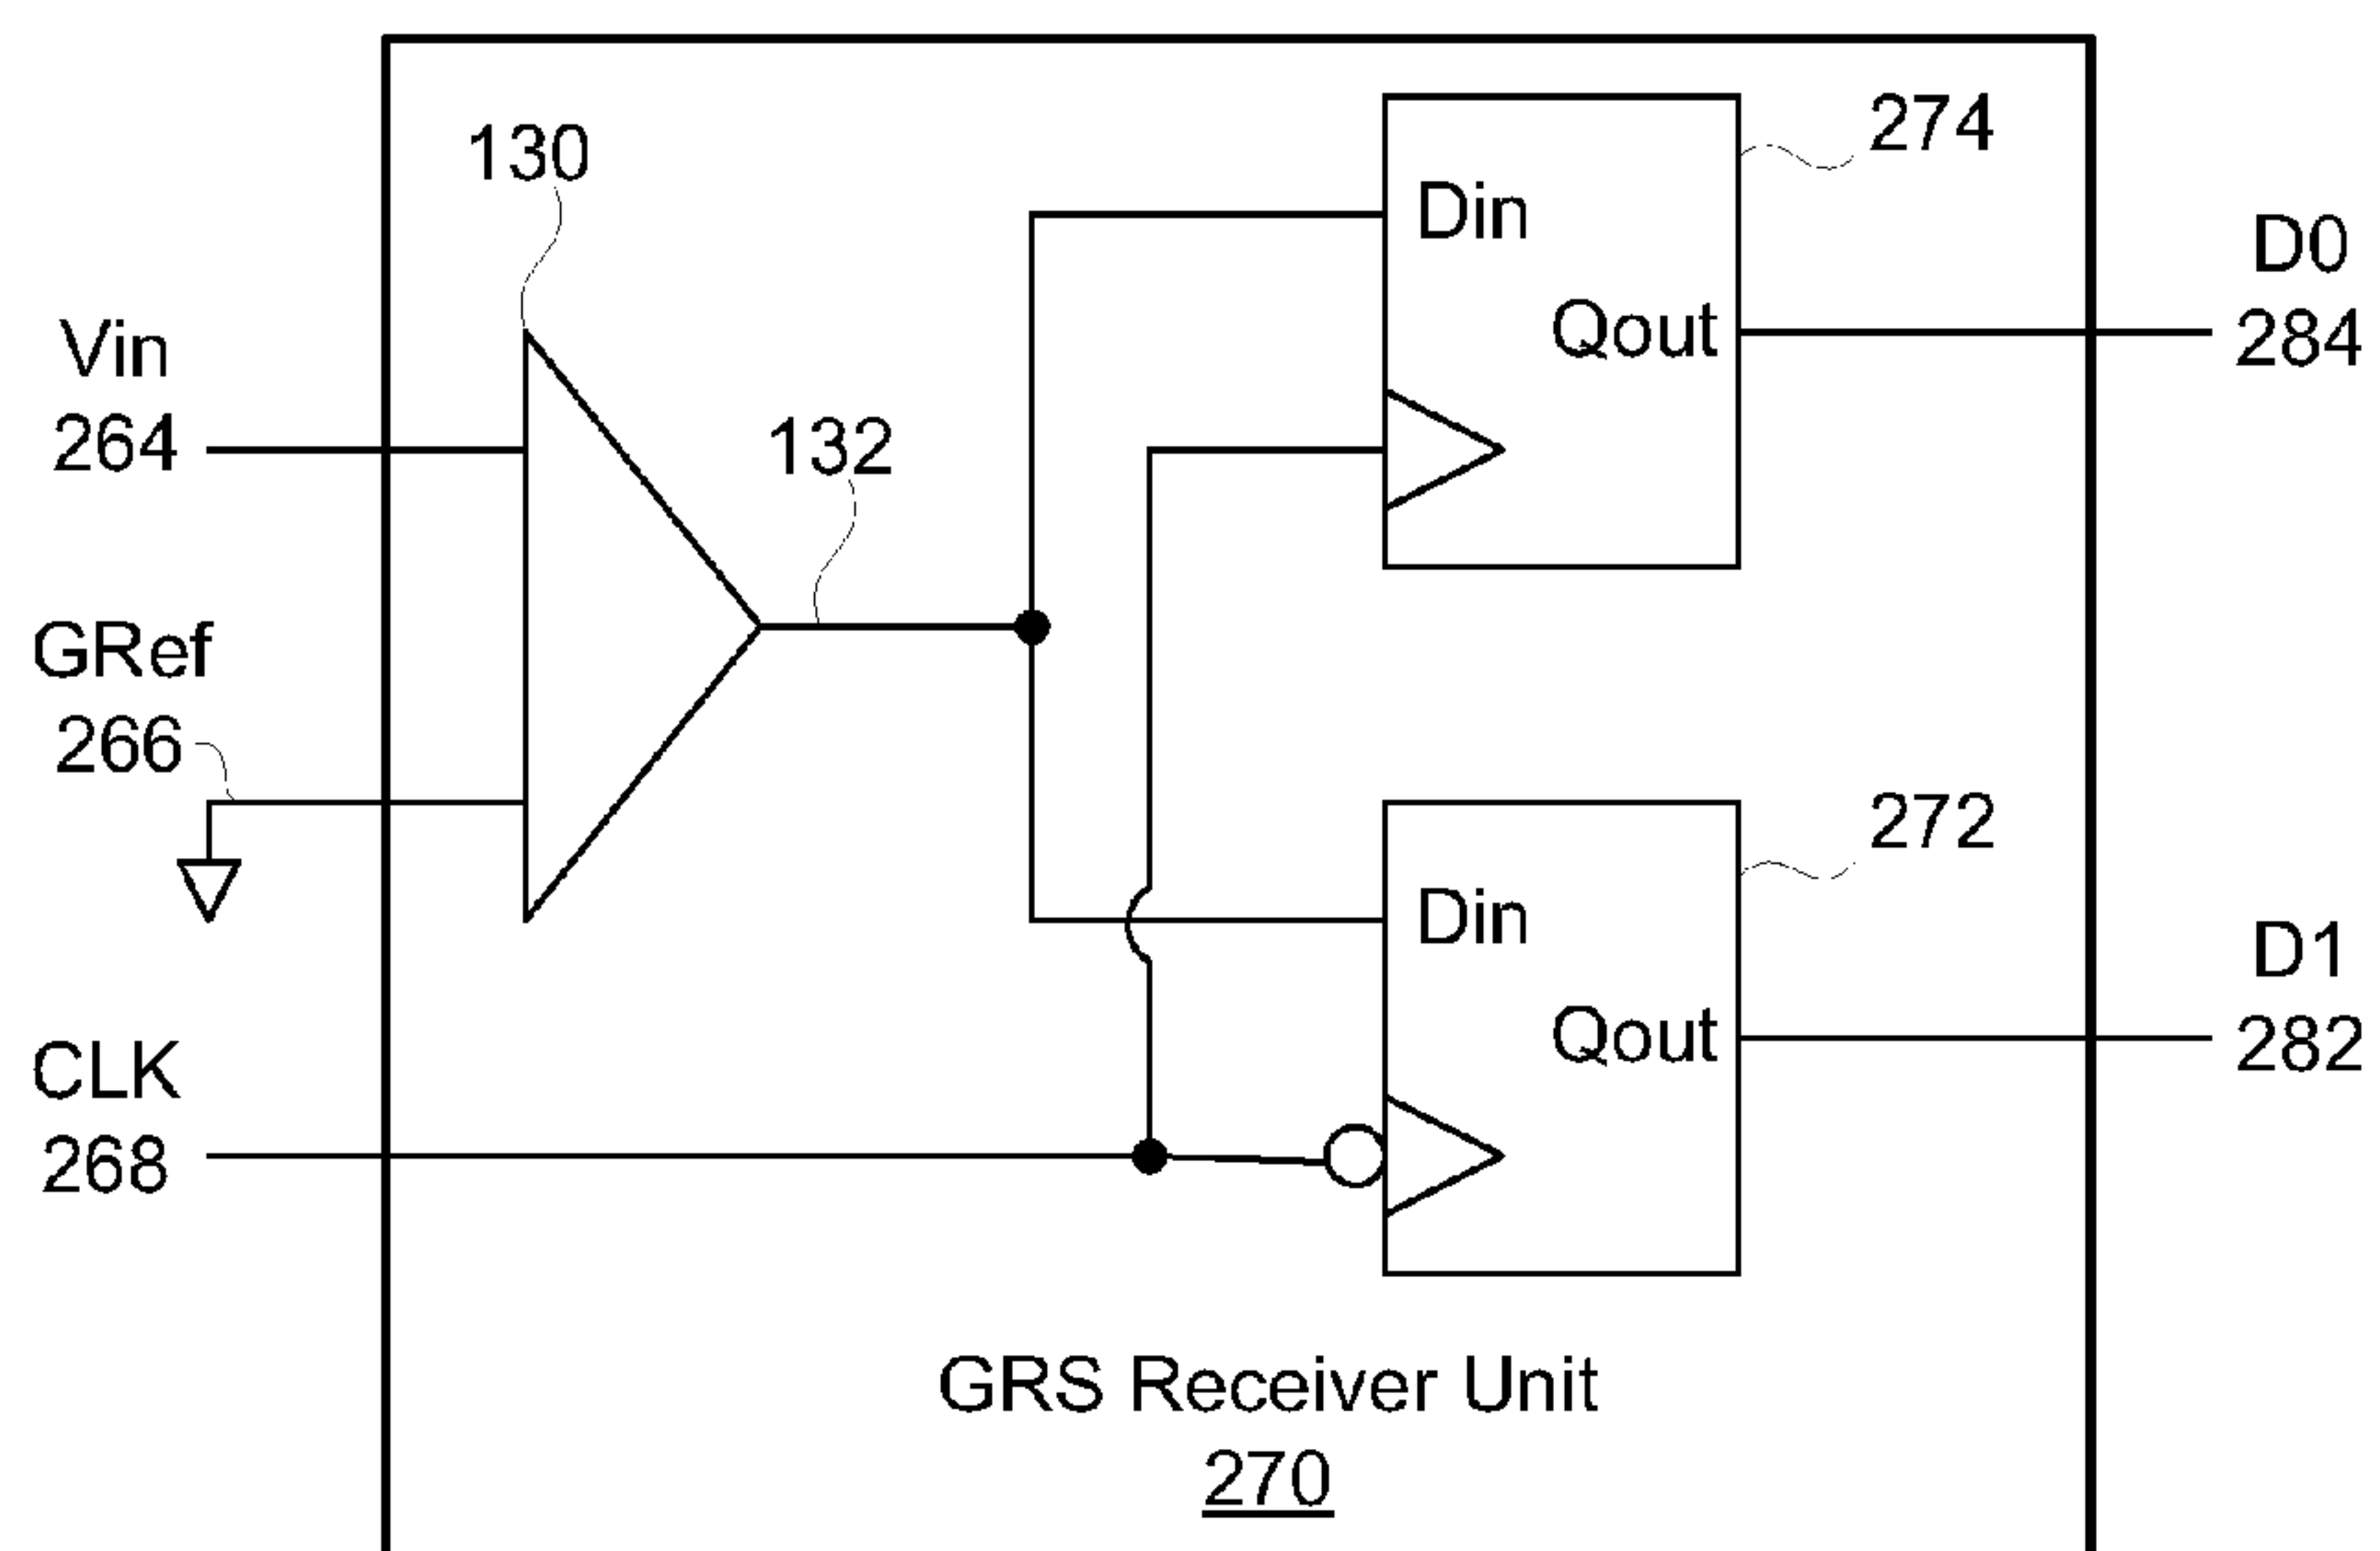
*Fig. 1F*



*Fig. 1G*



**Fig. 2A**



**Fig. 2B**



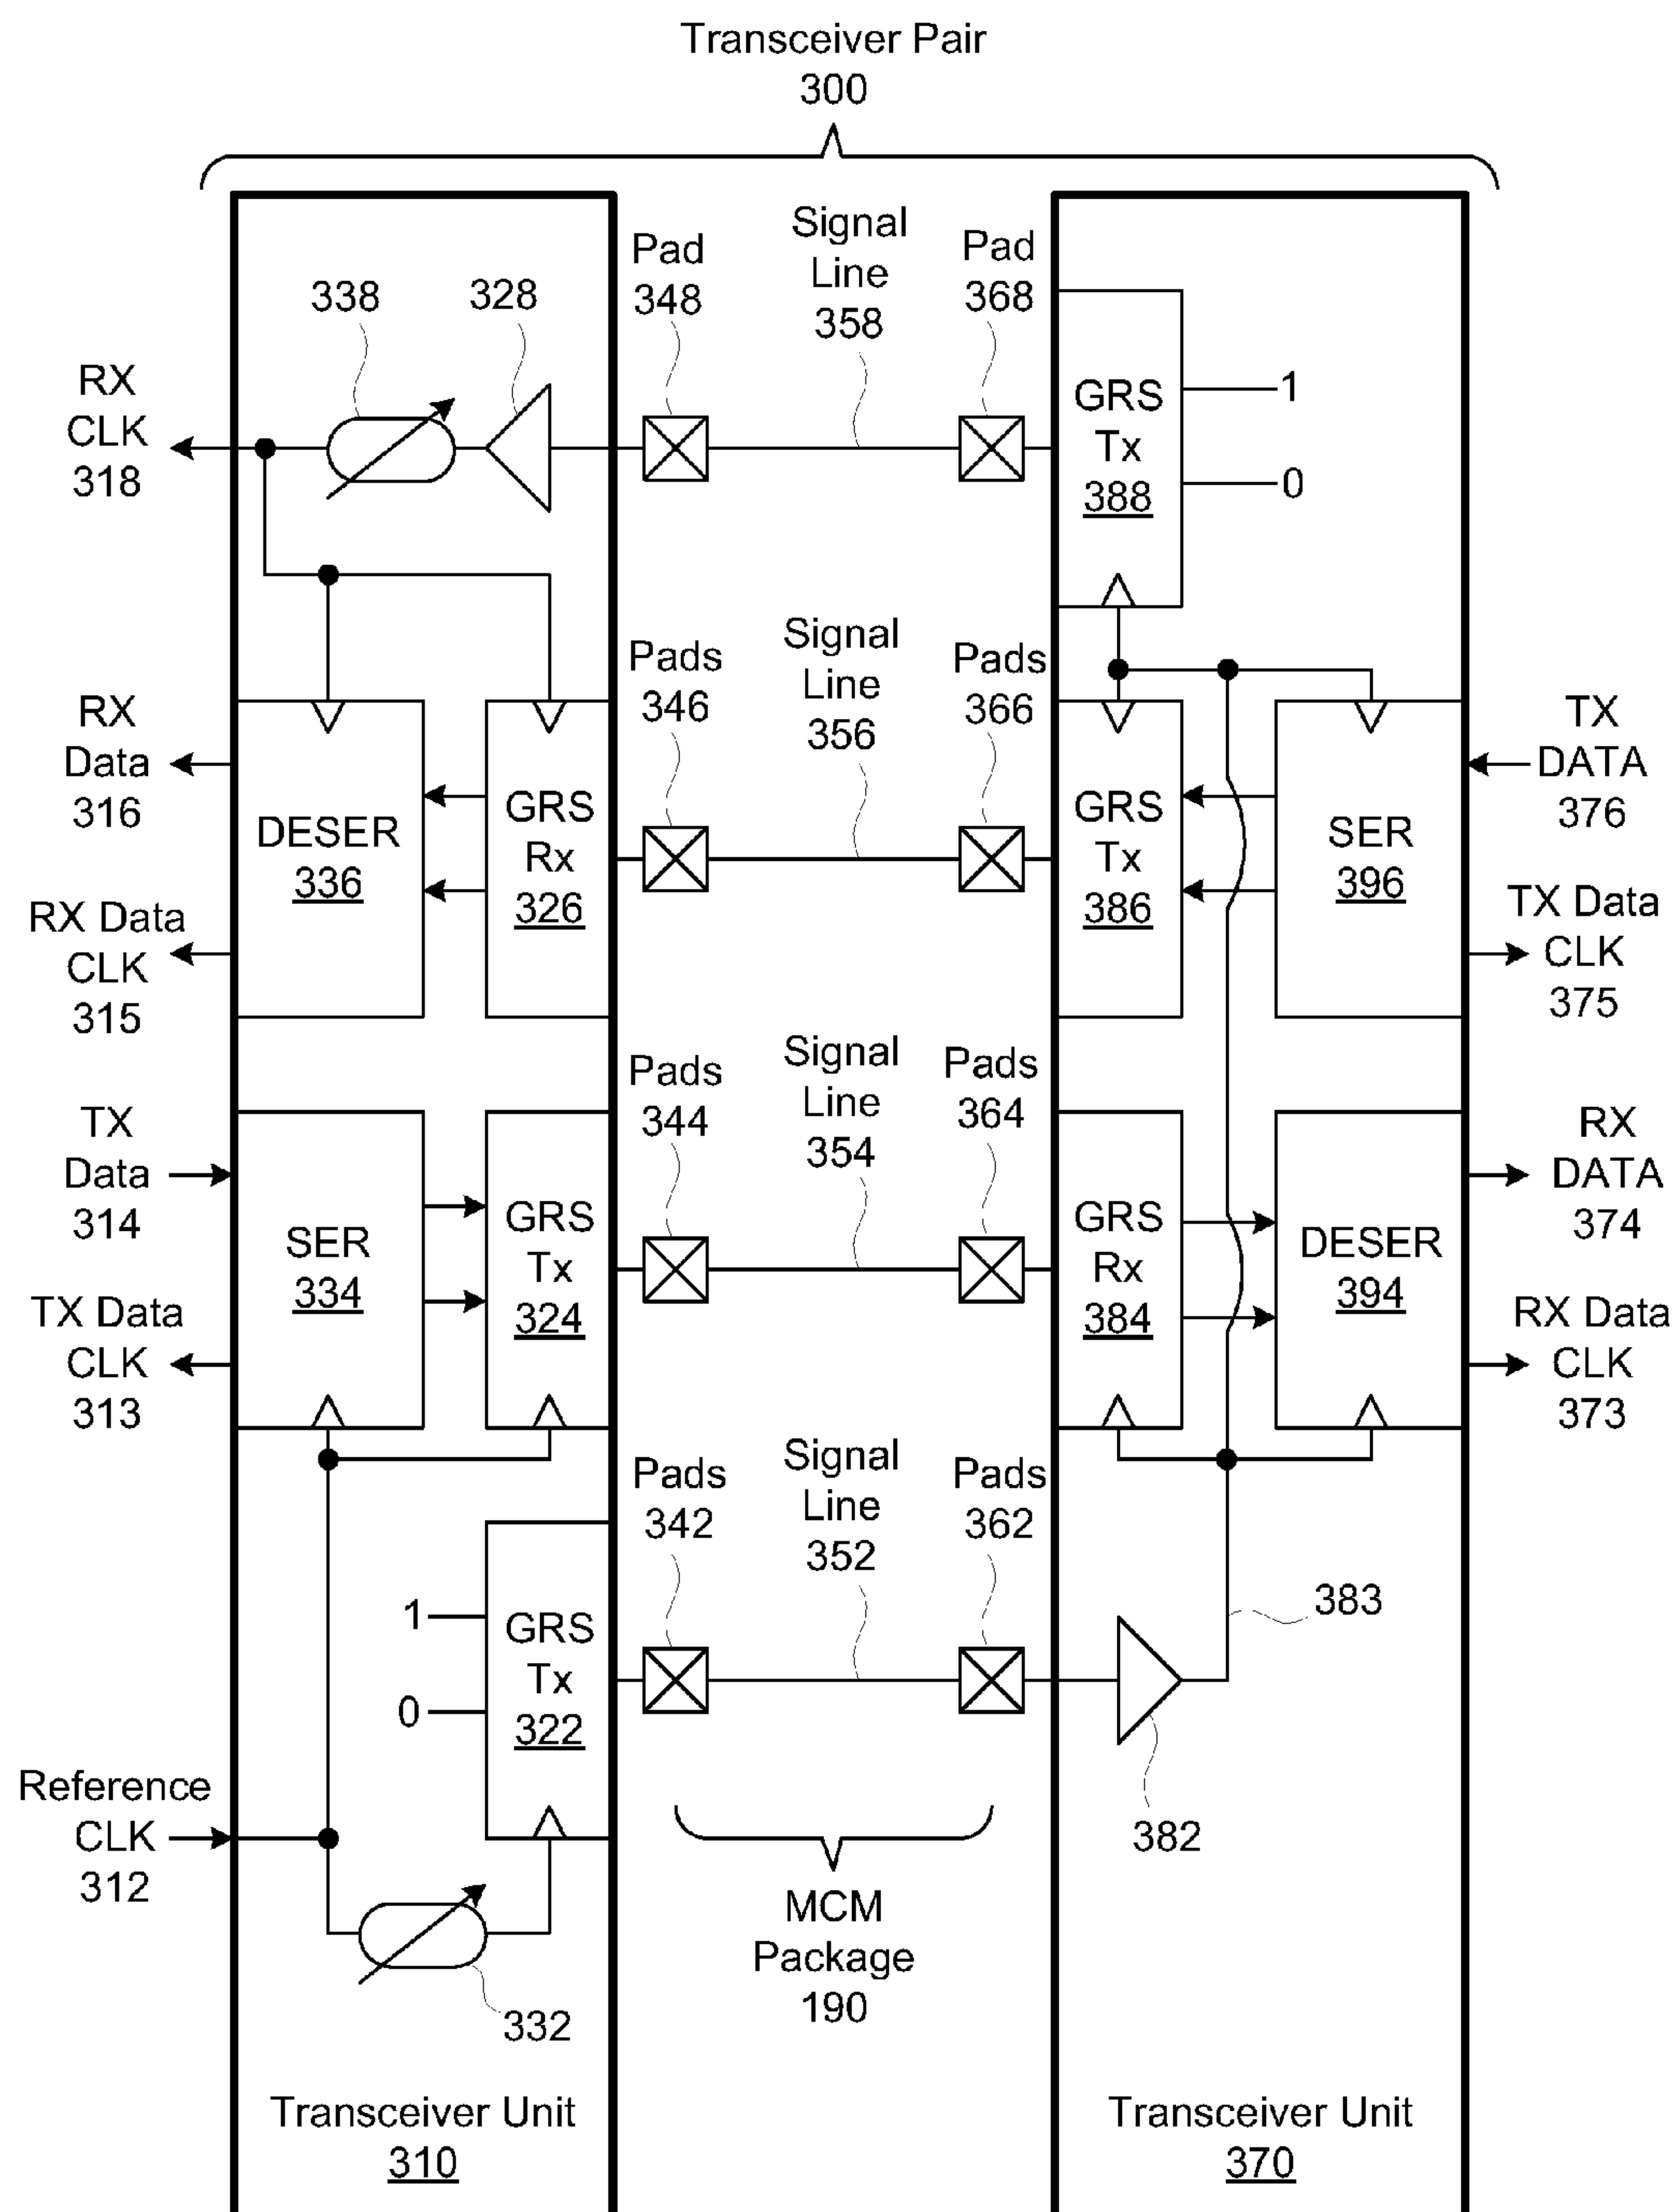
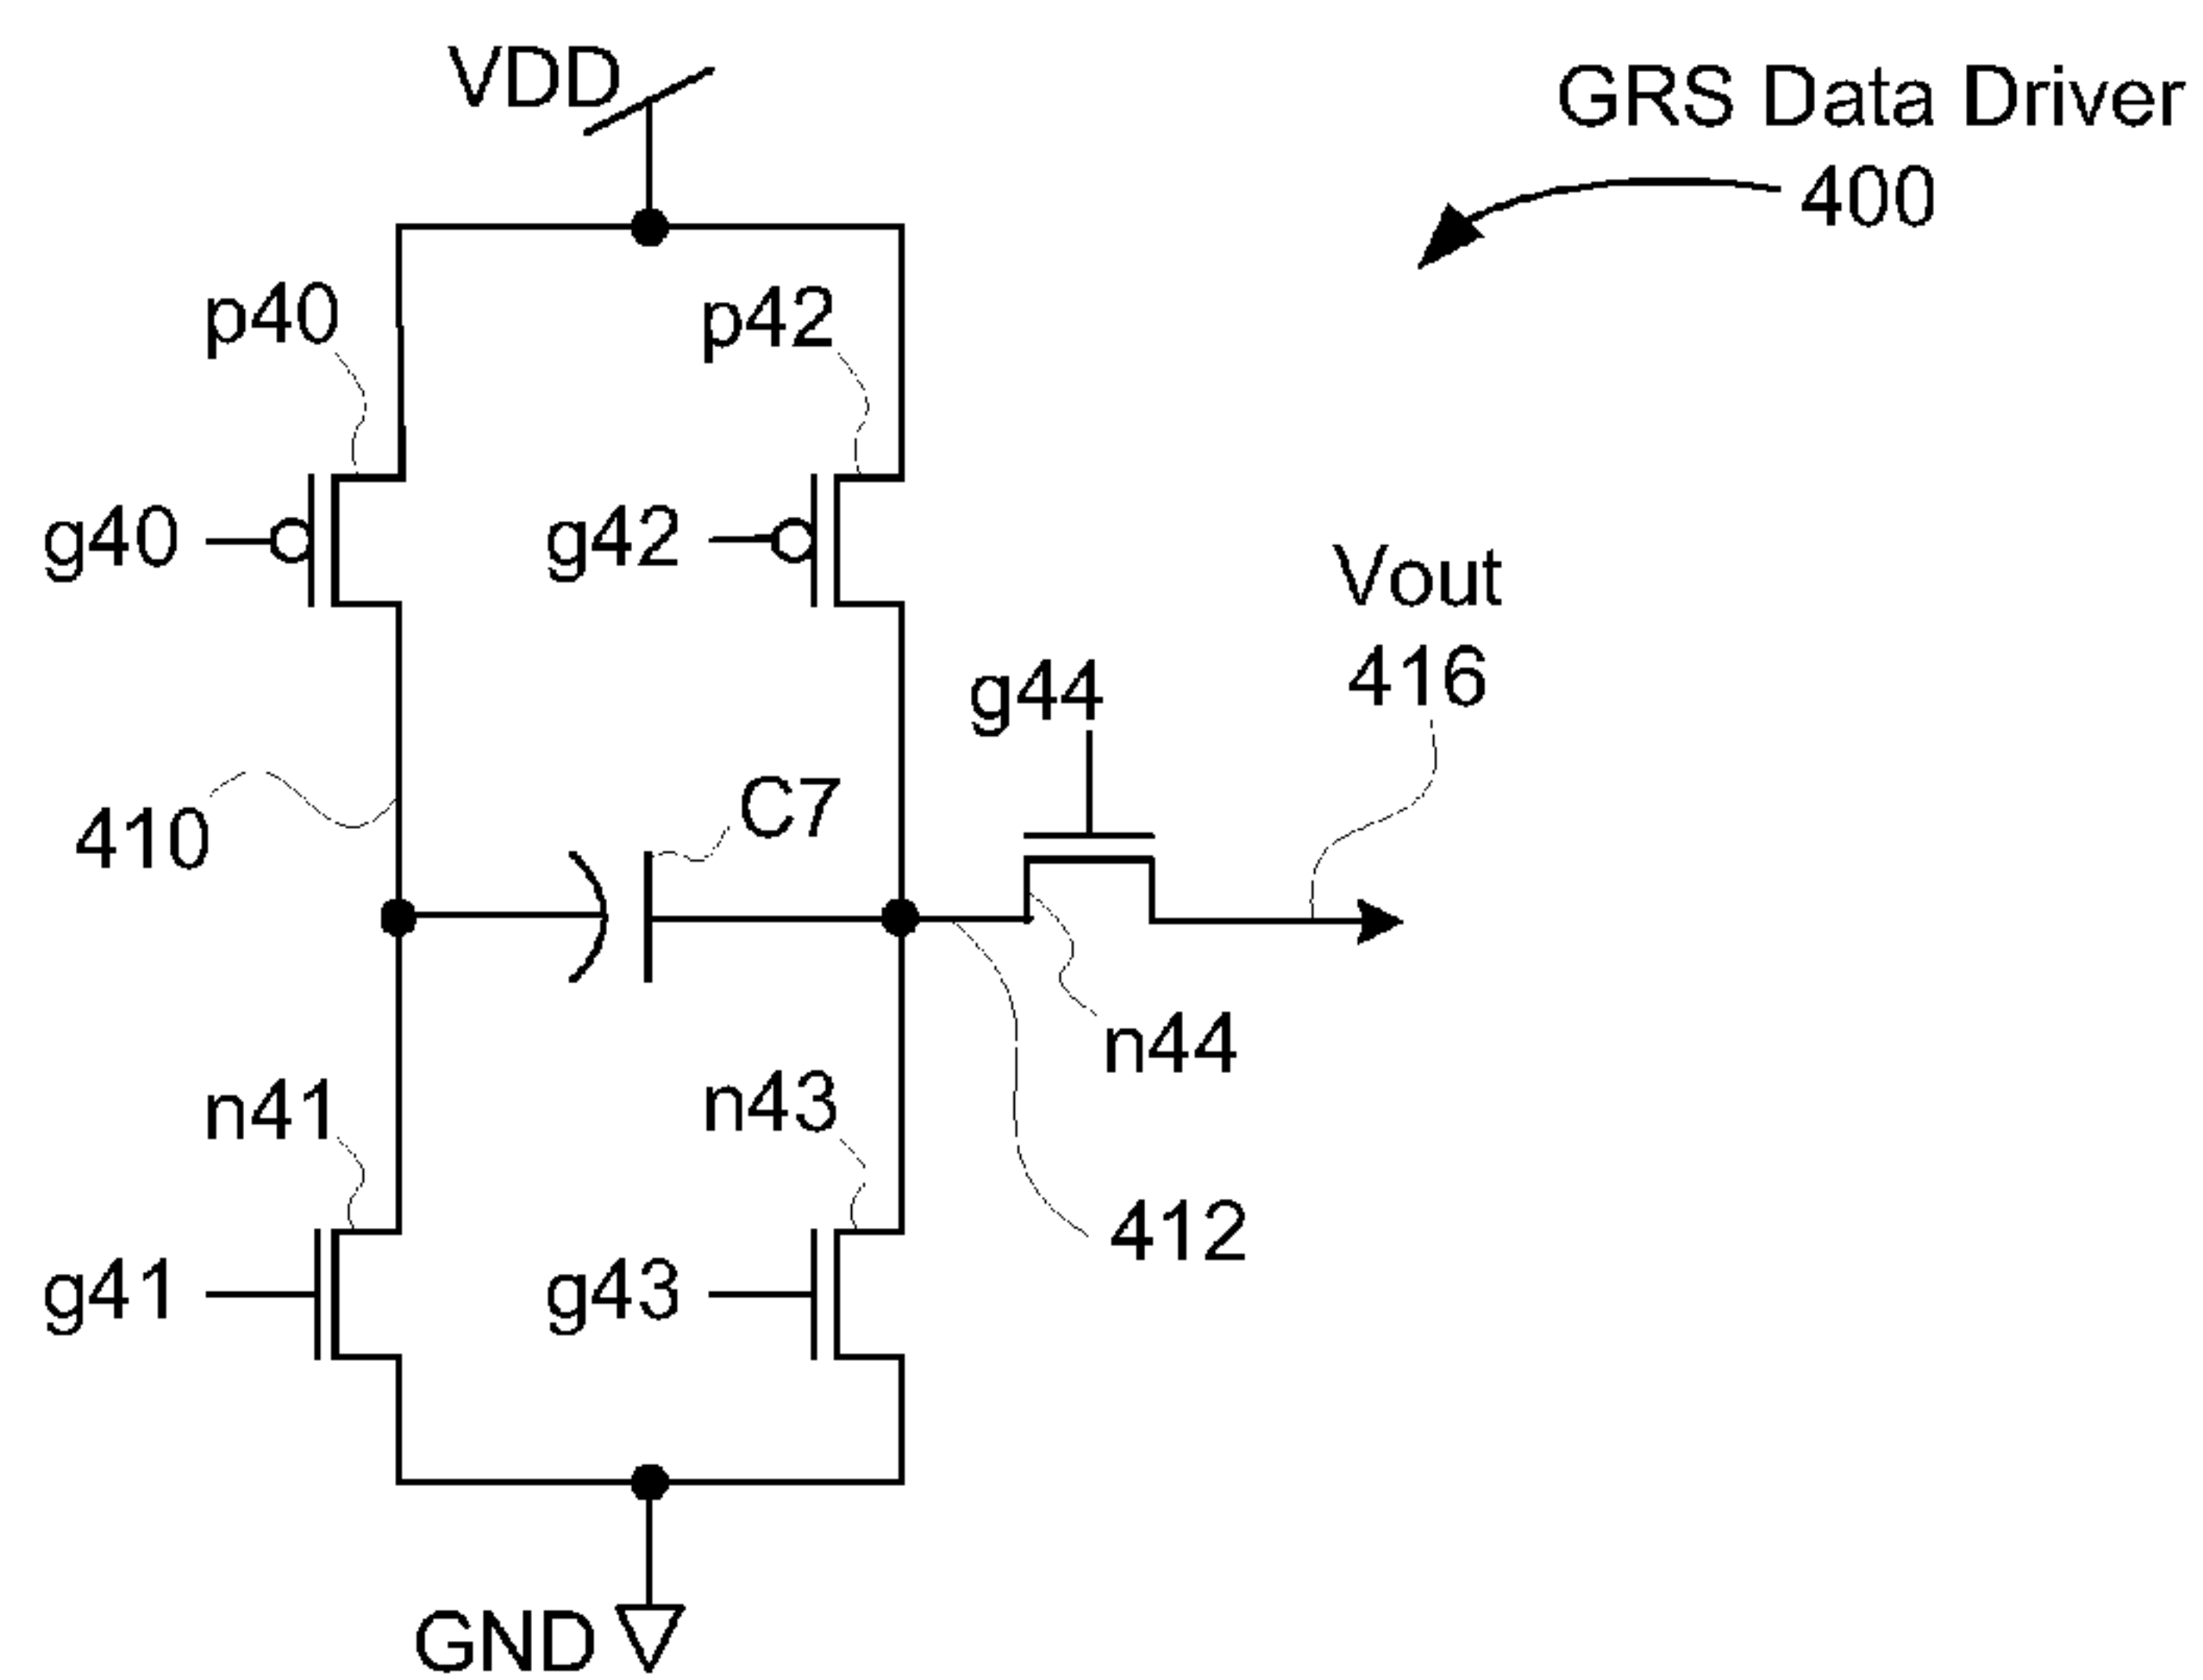
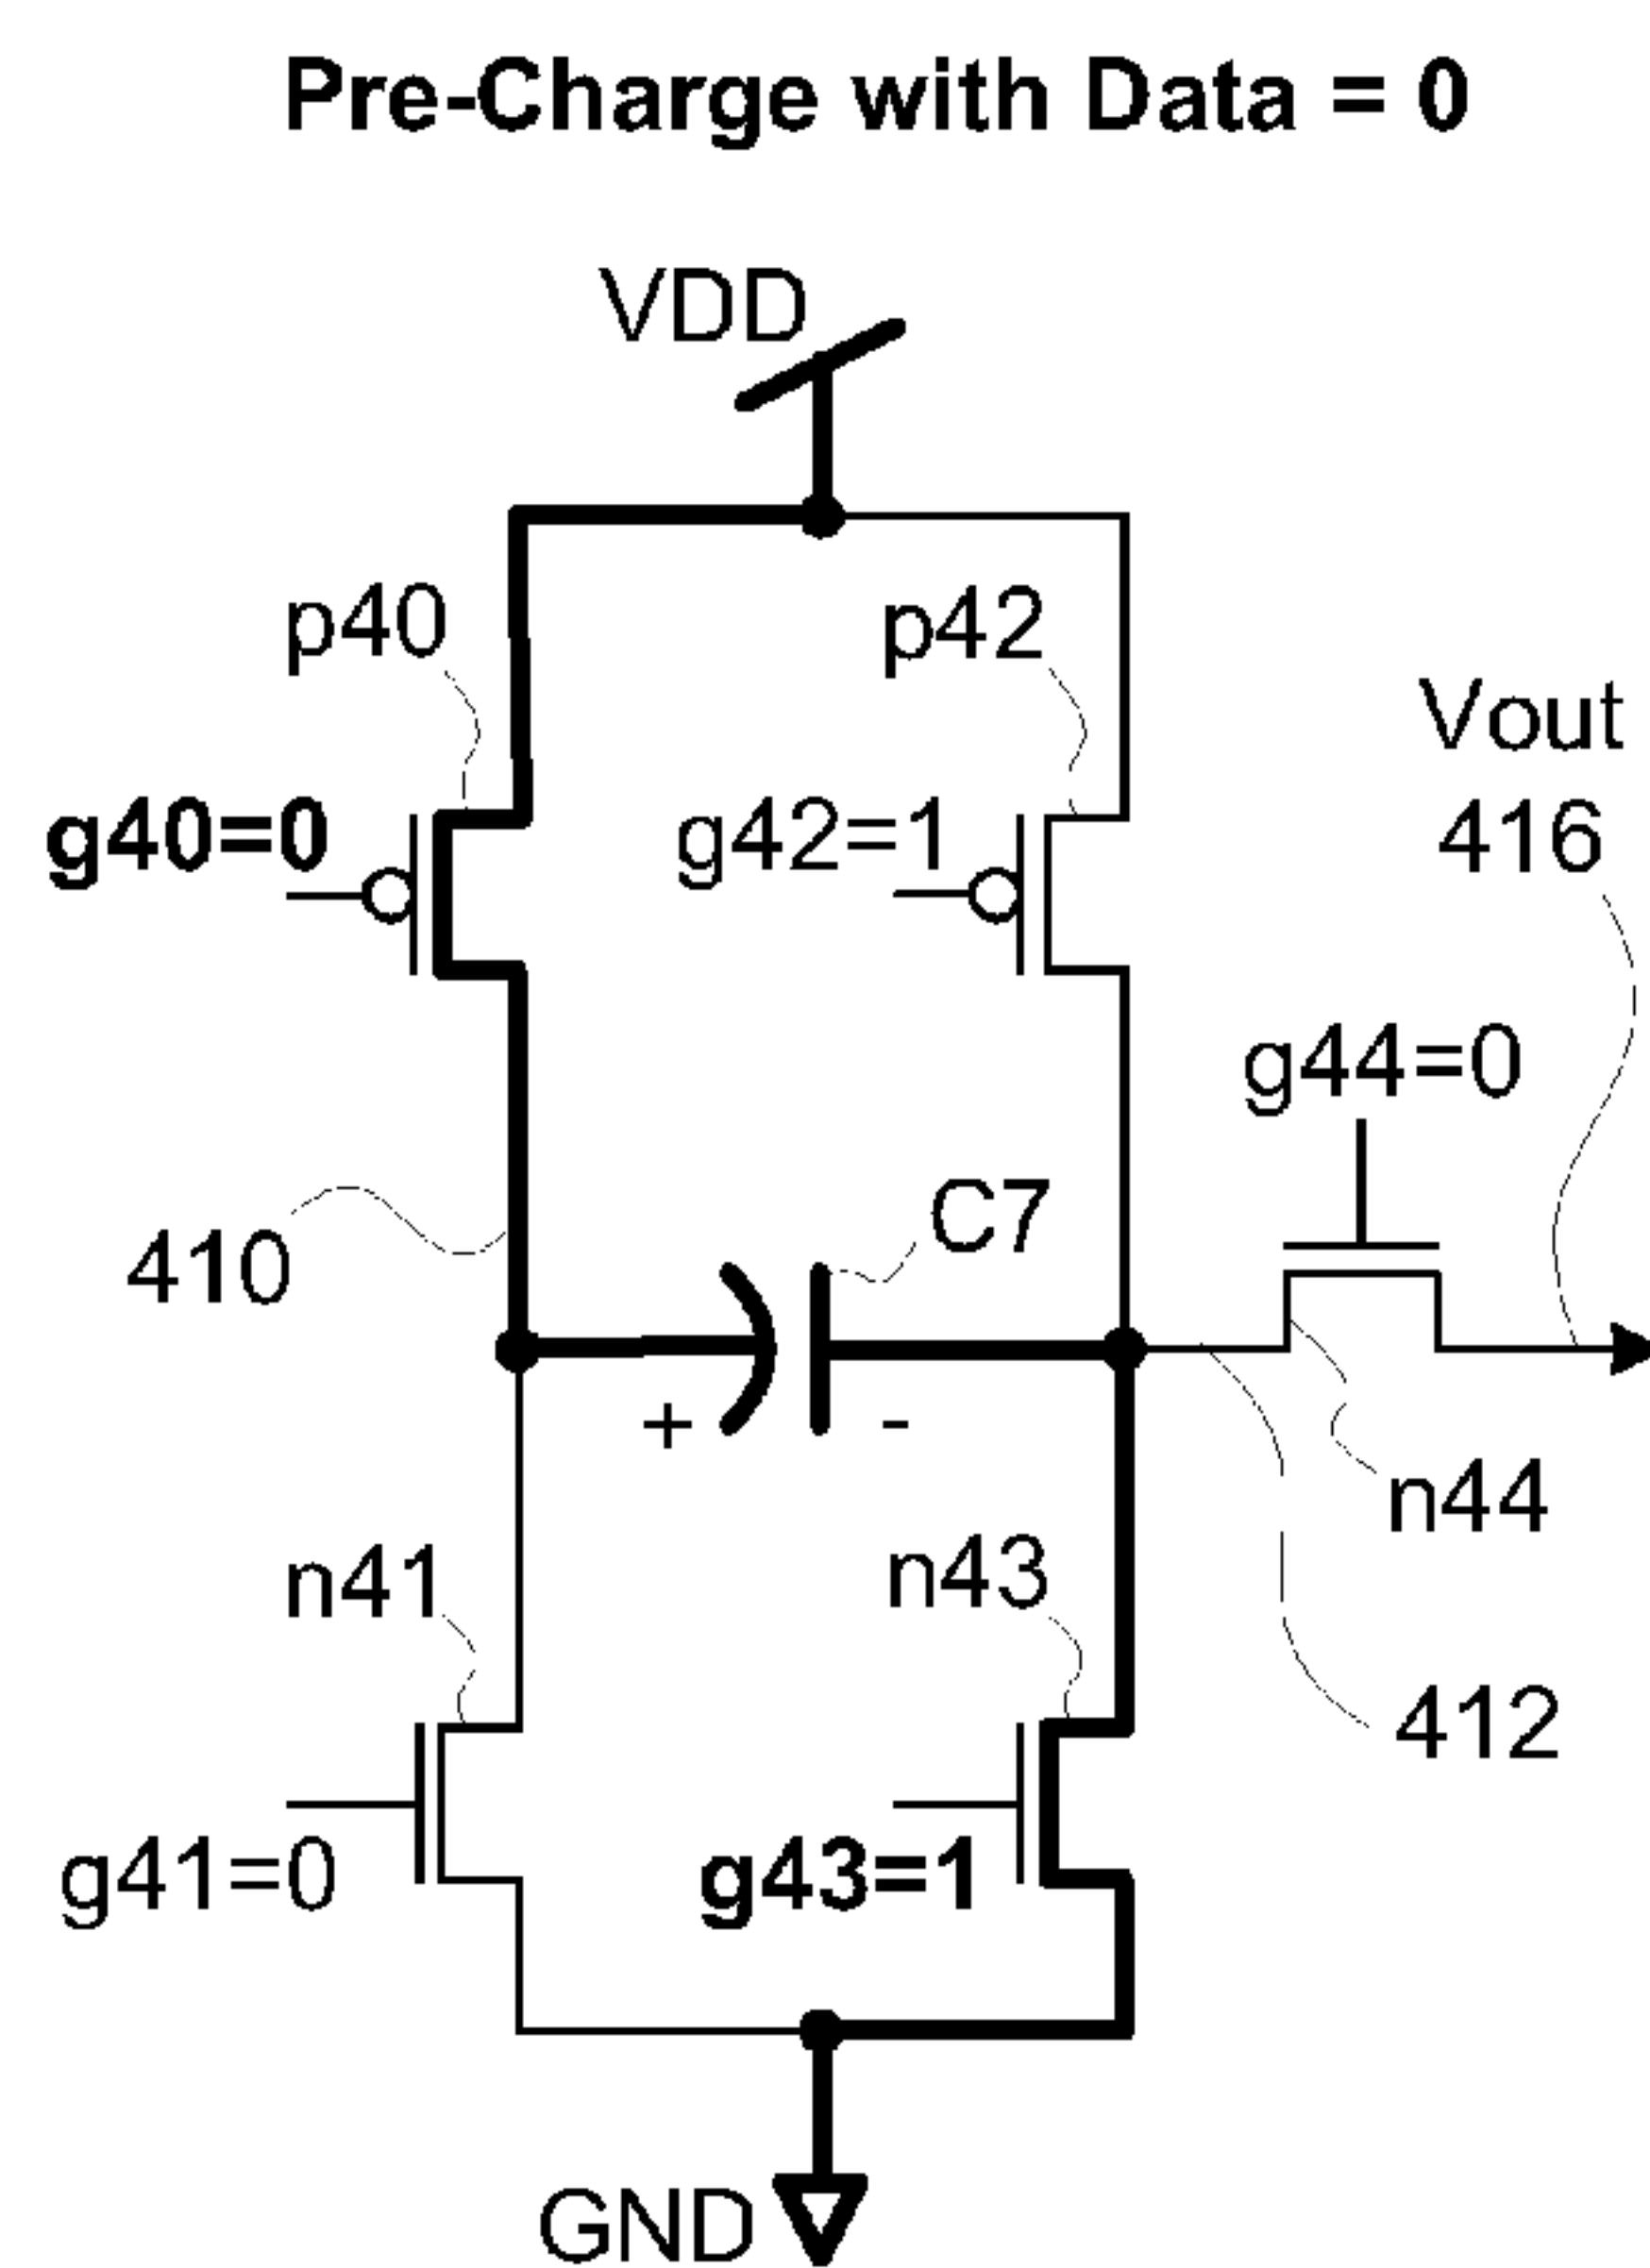


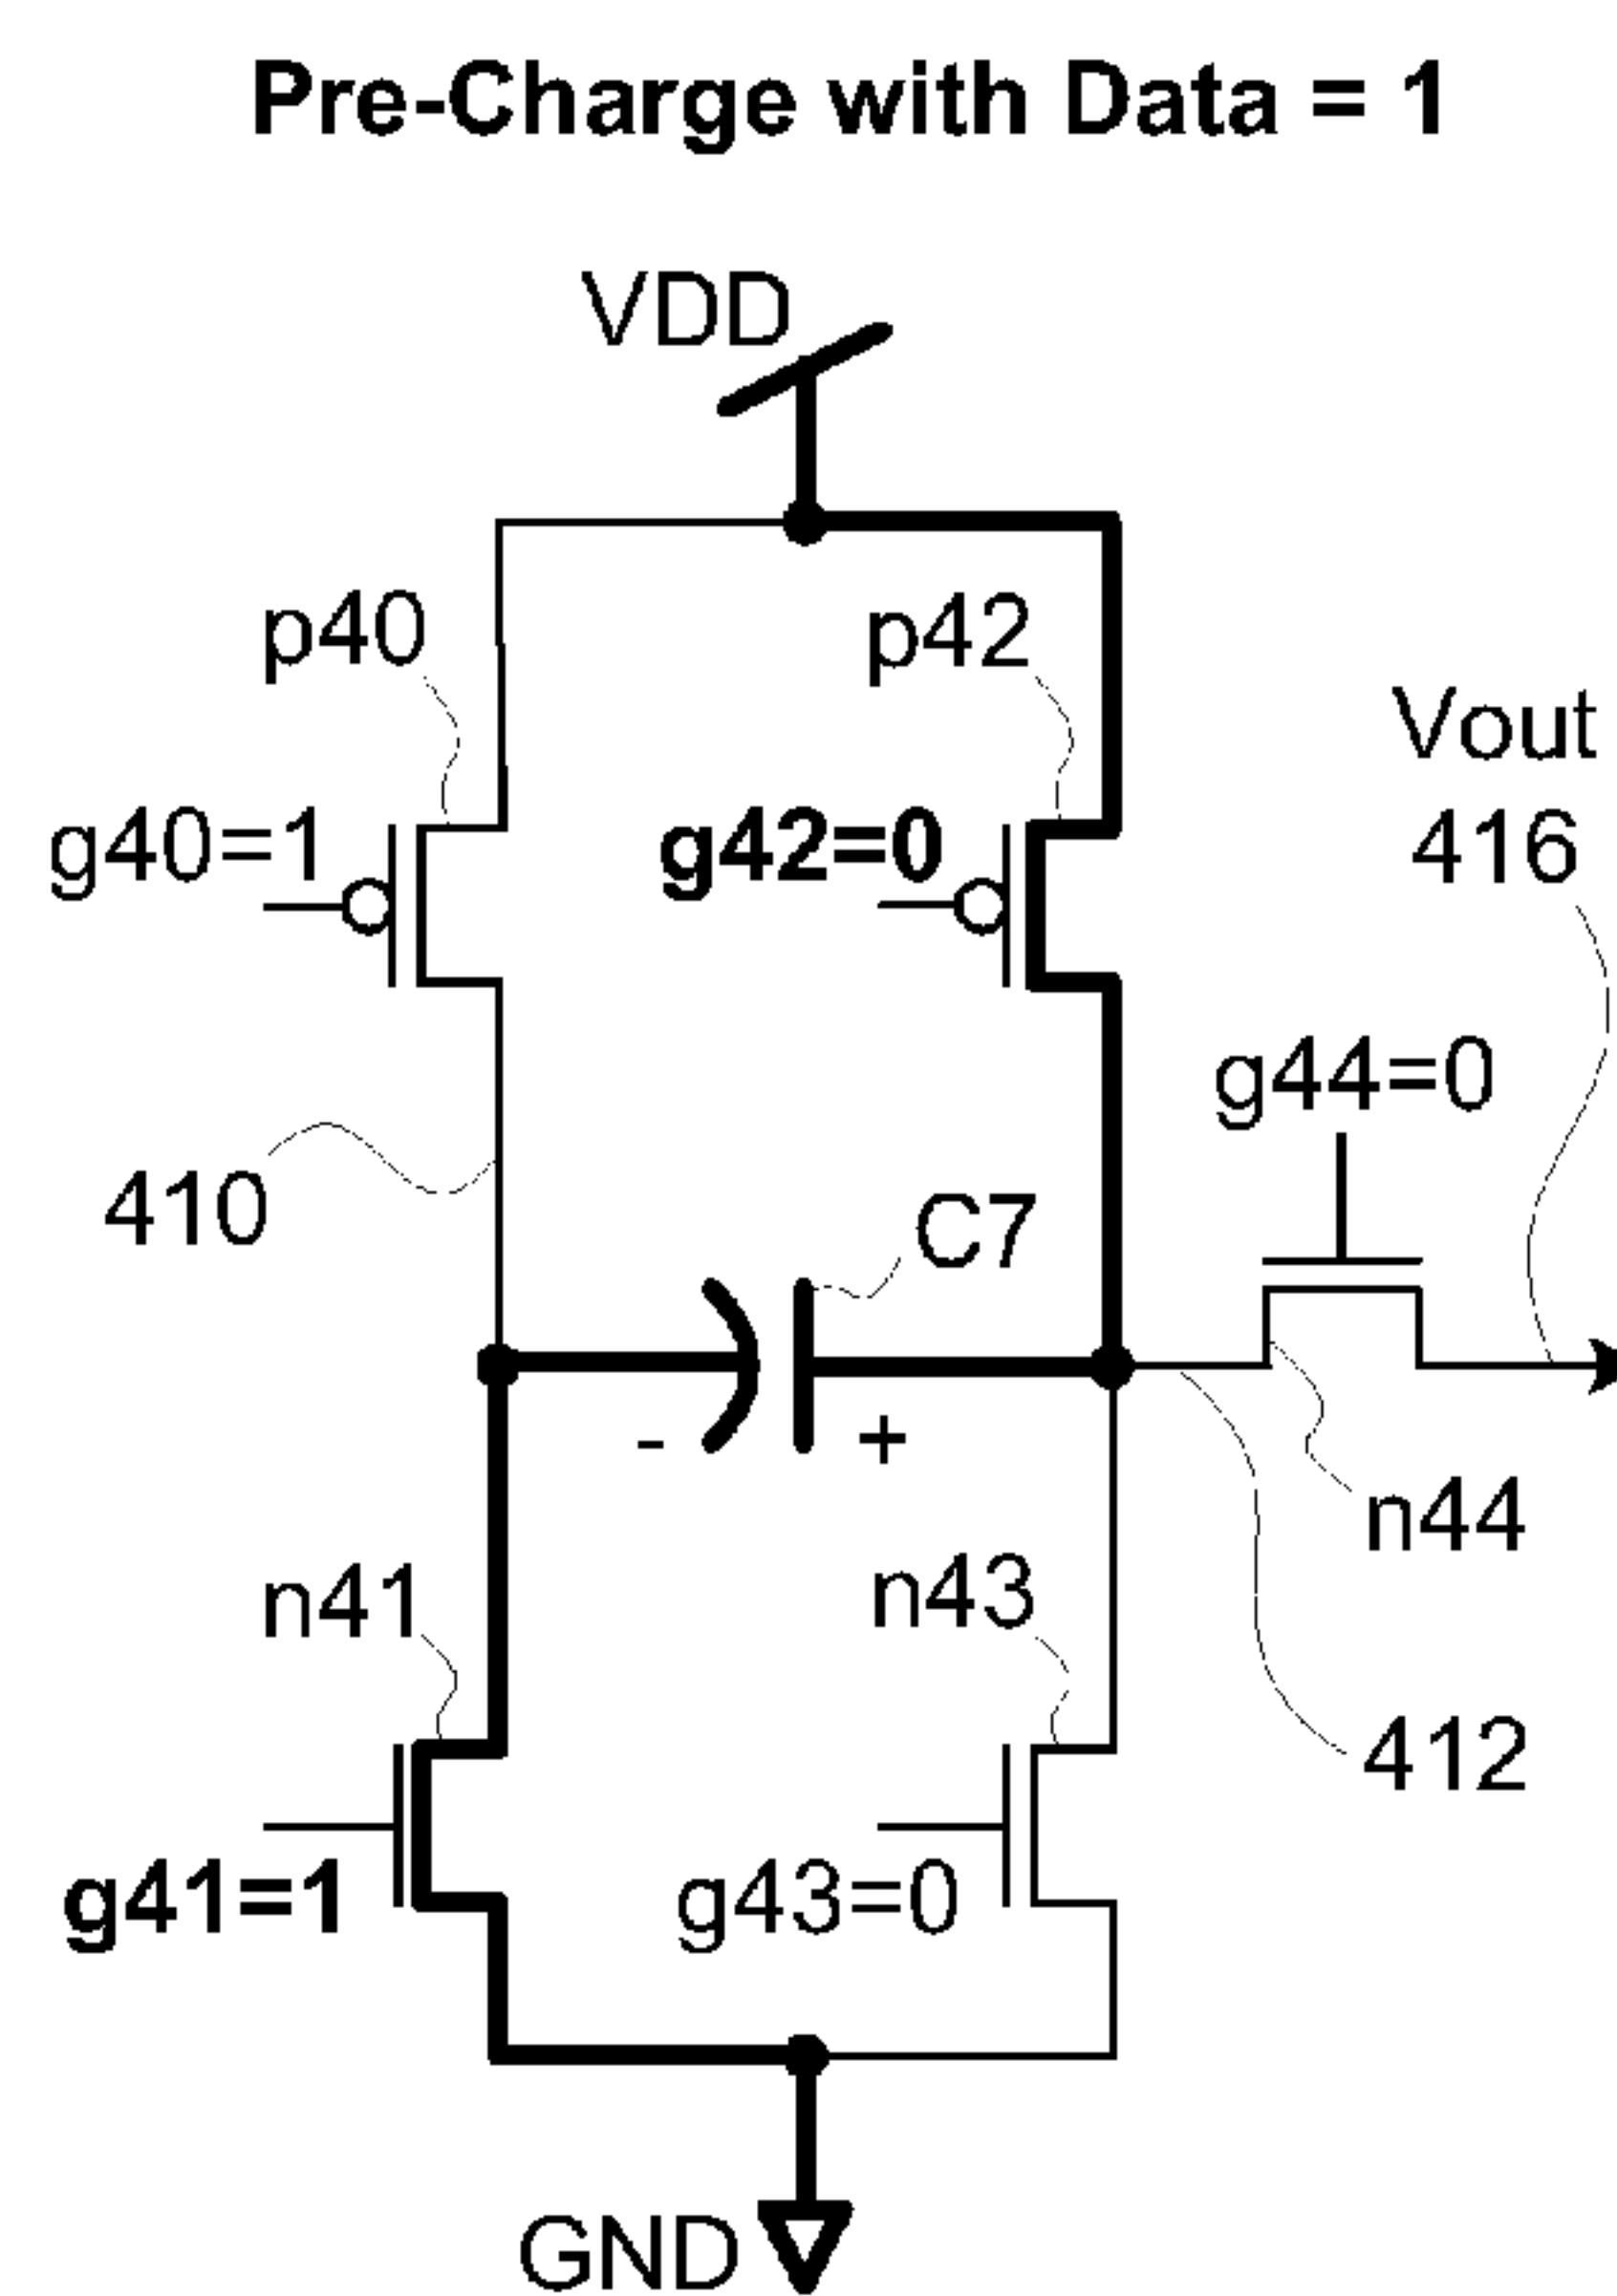
Fig. 3



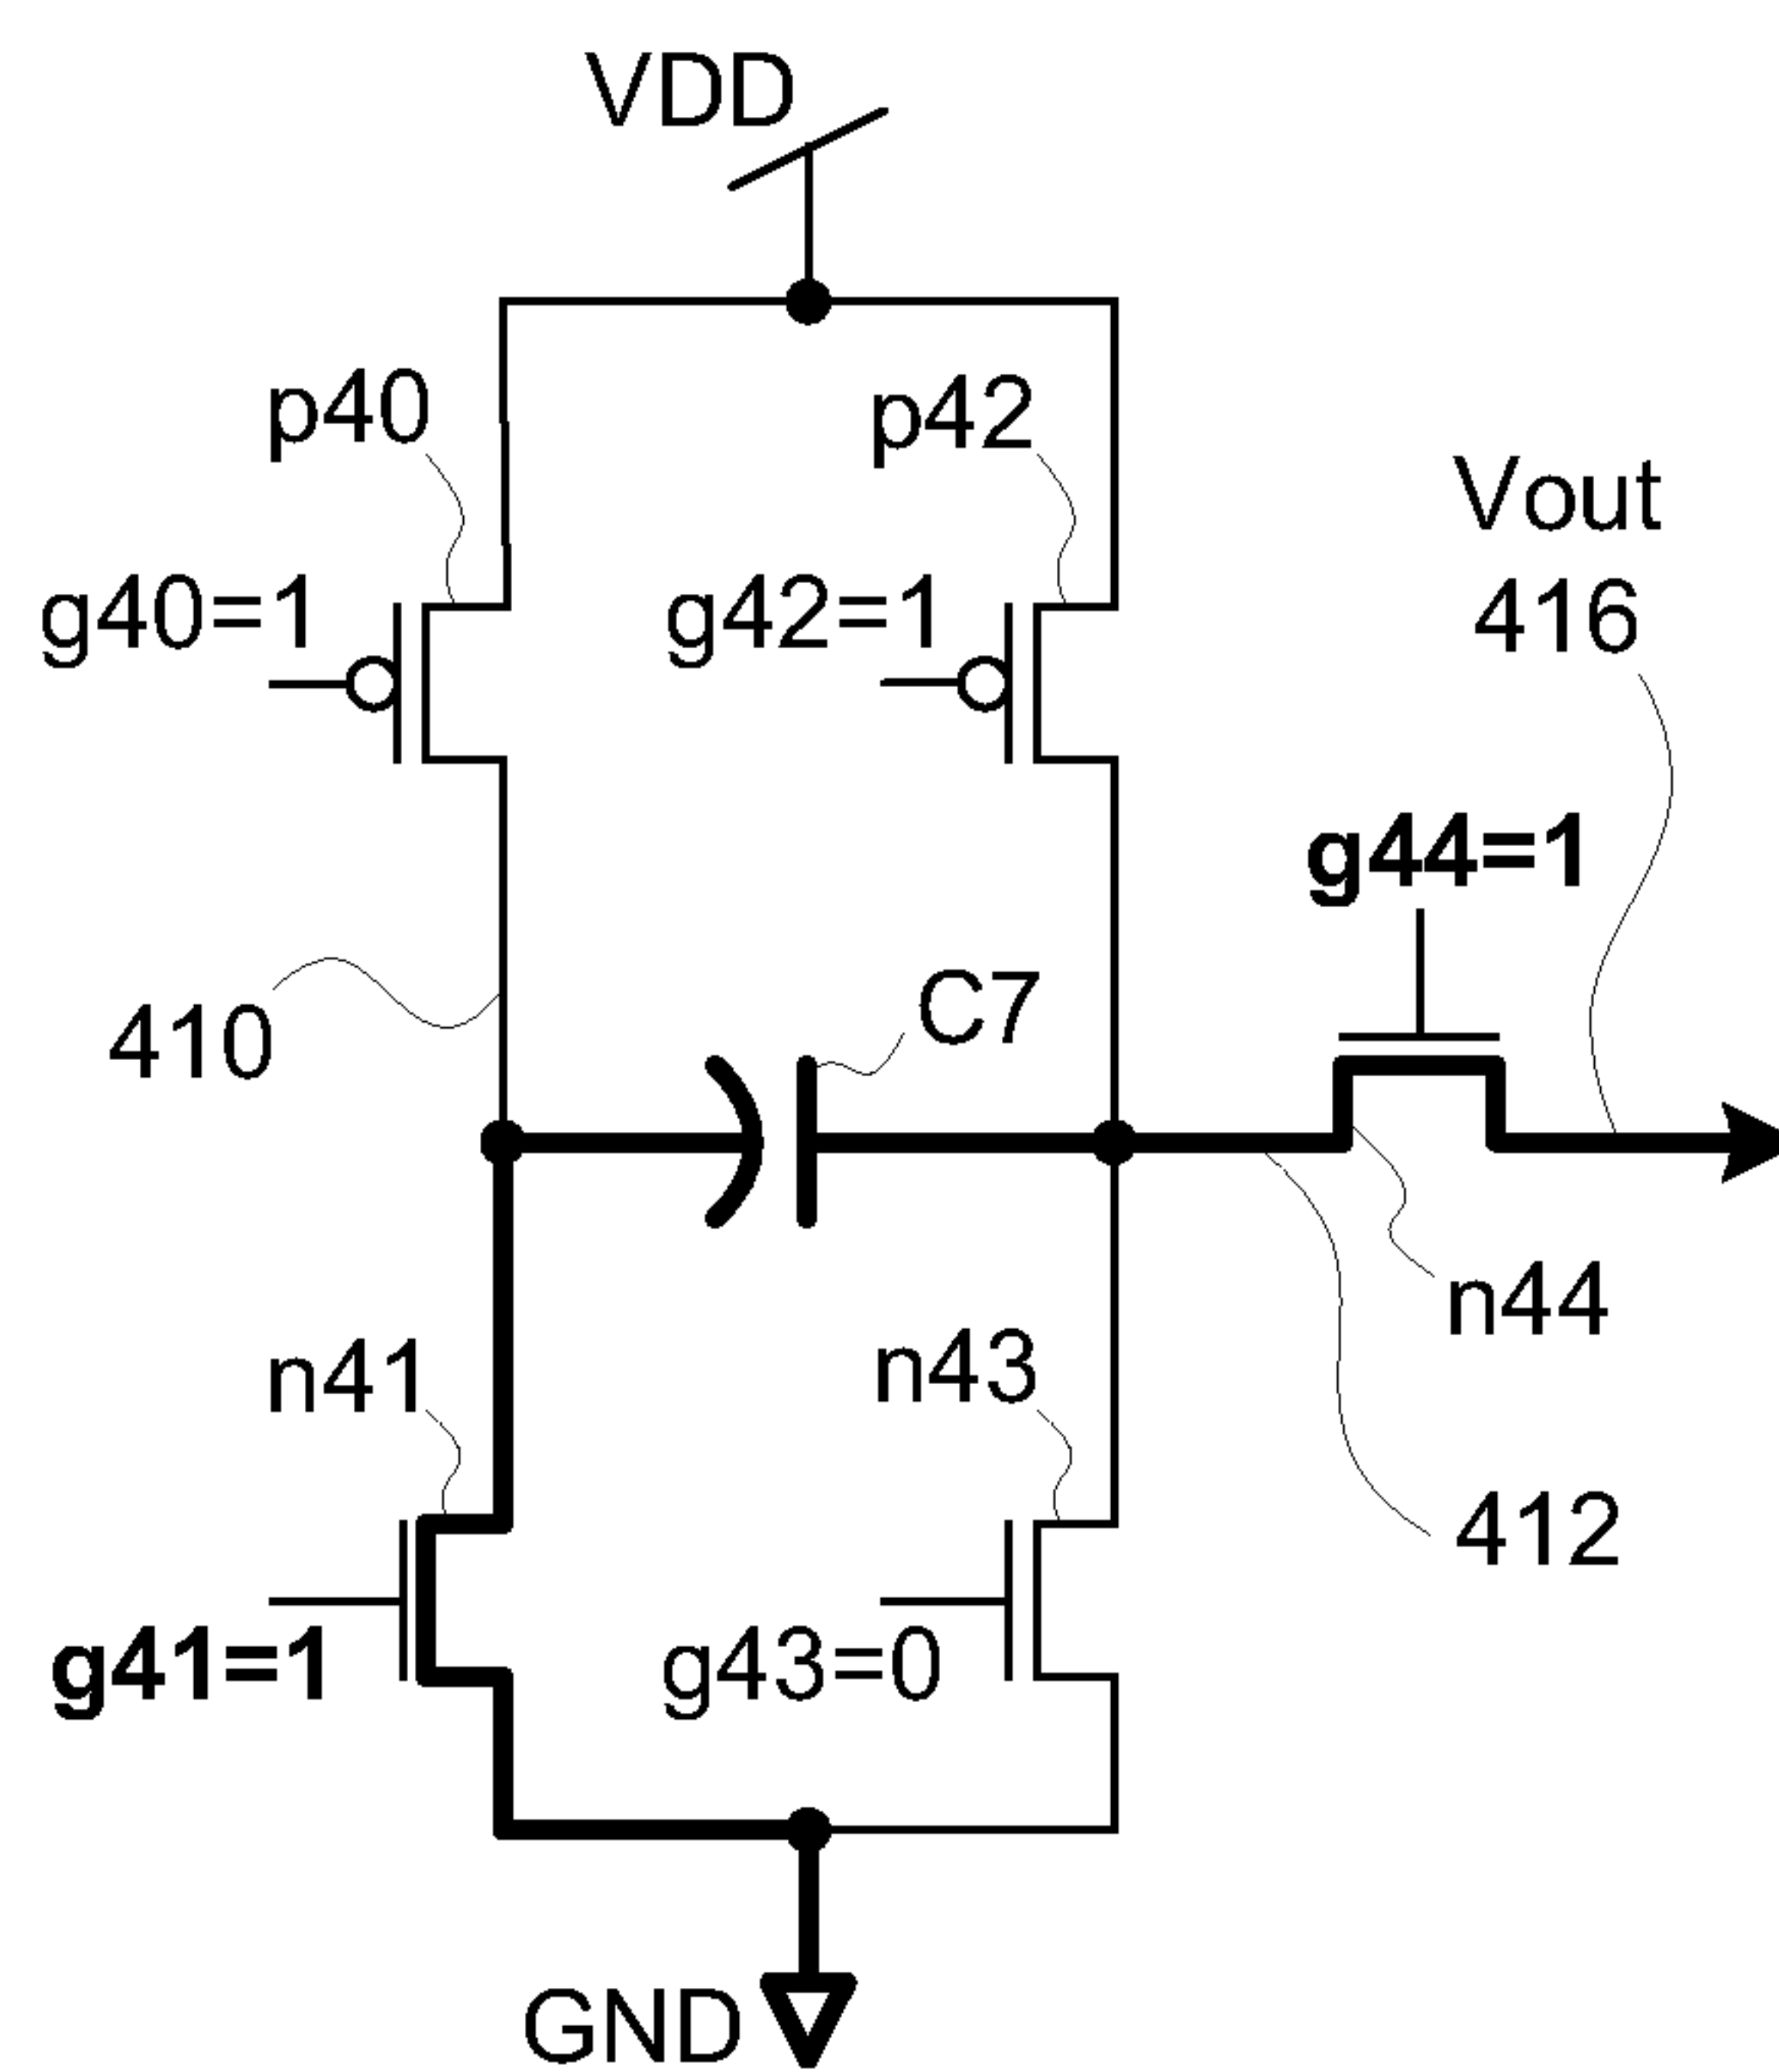
**Fig. 4A**



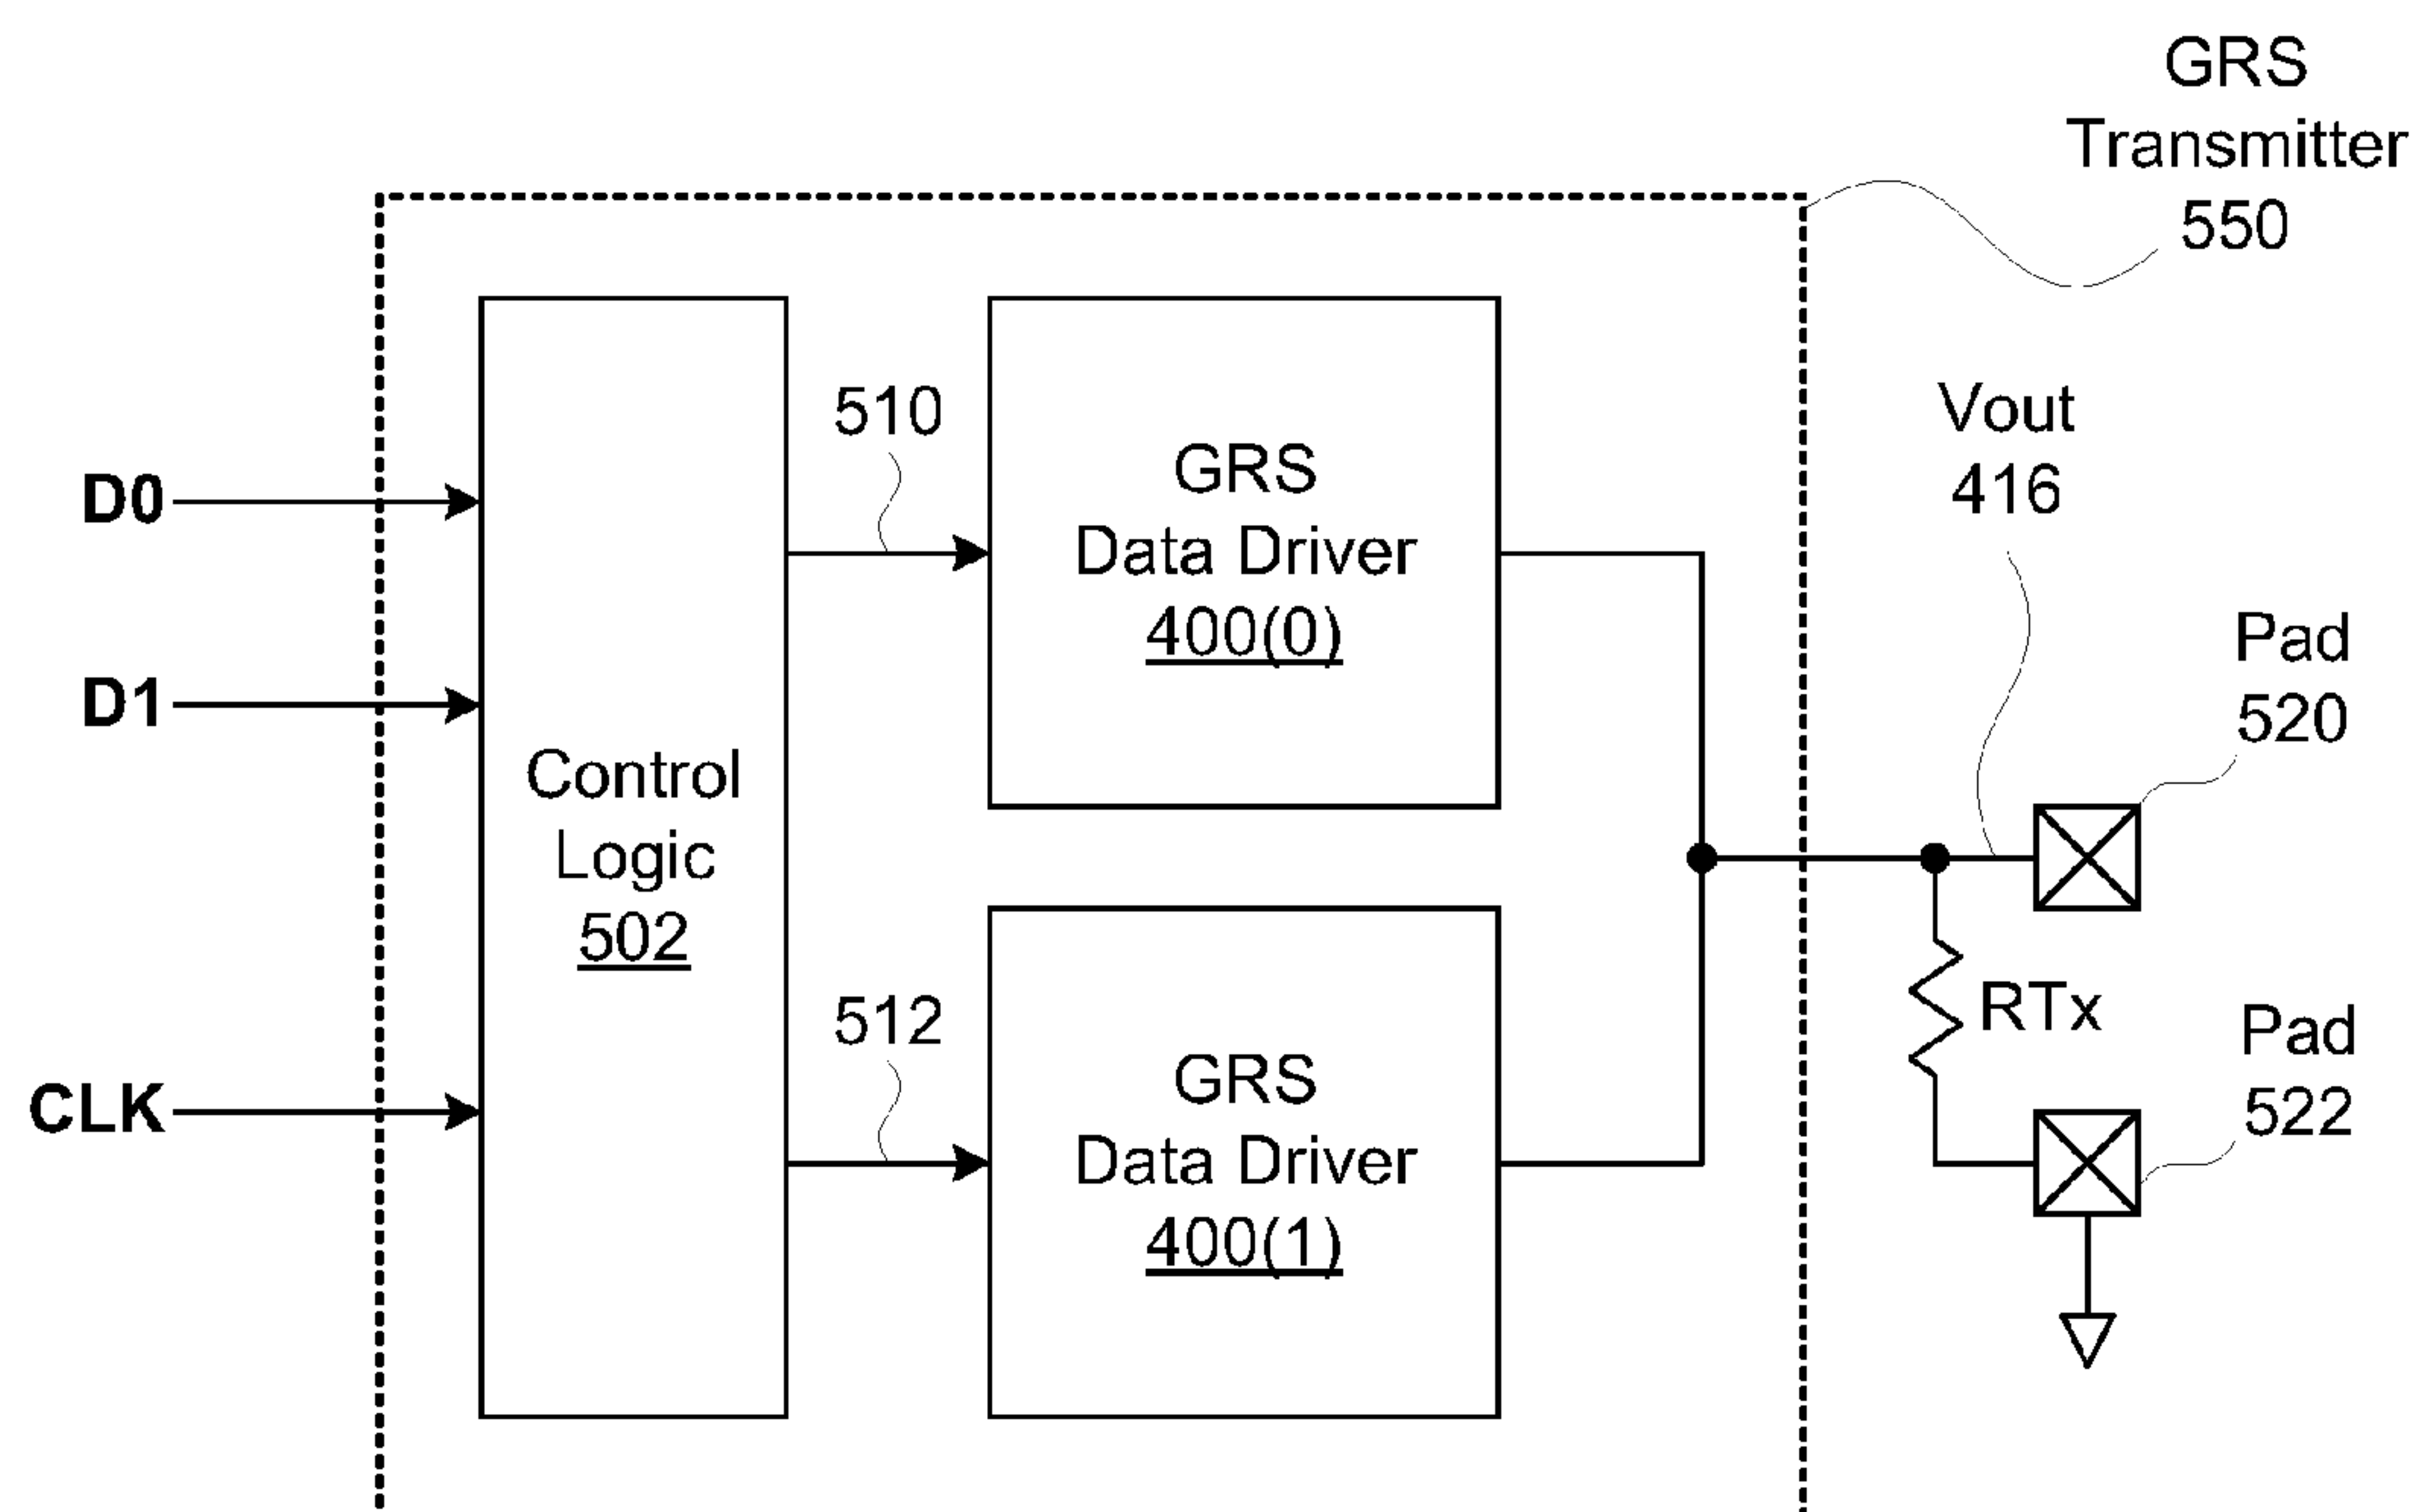
**Fig. 4B**



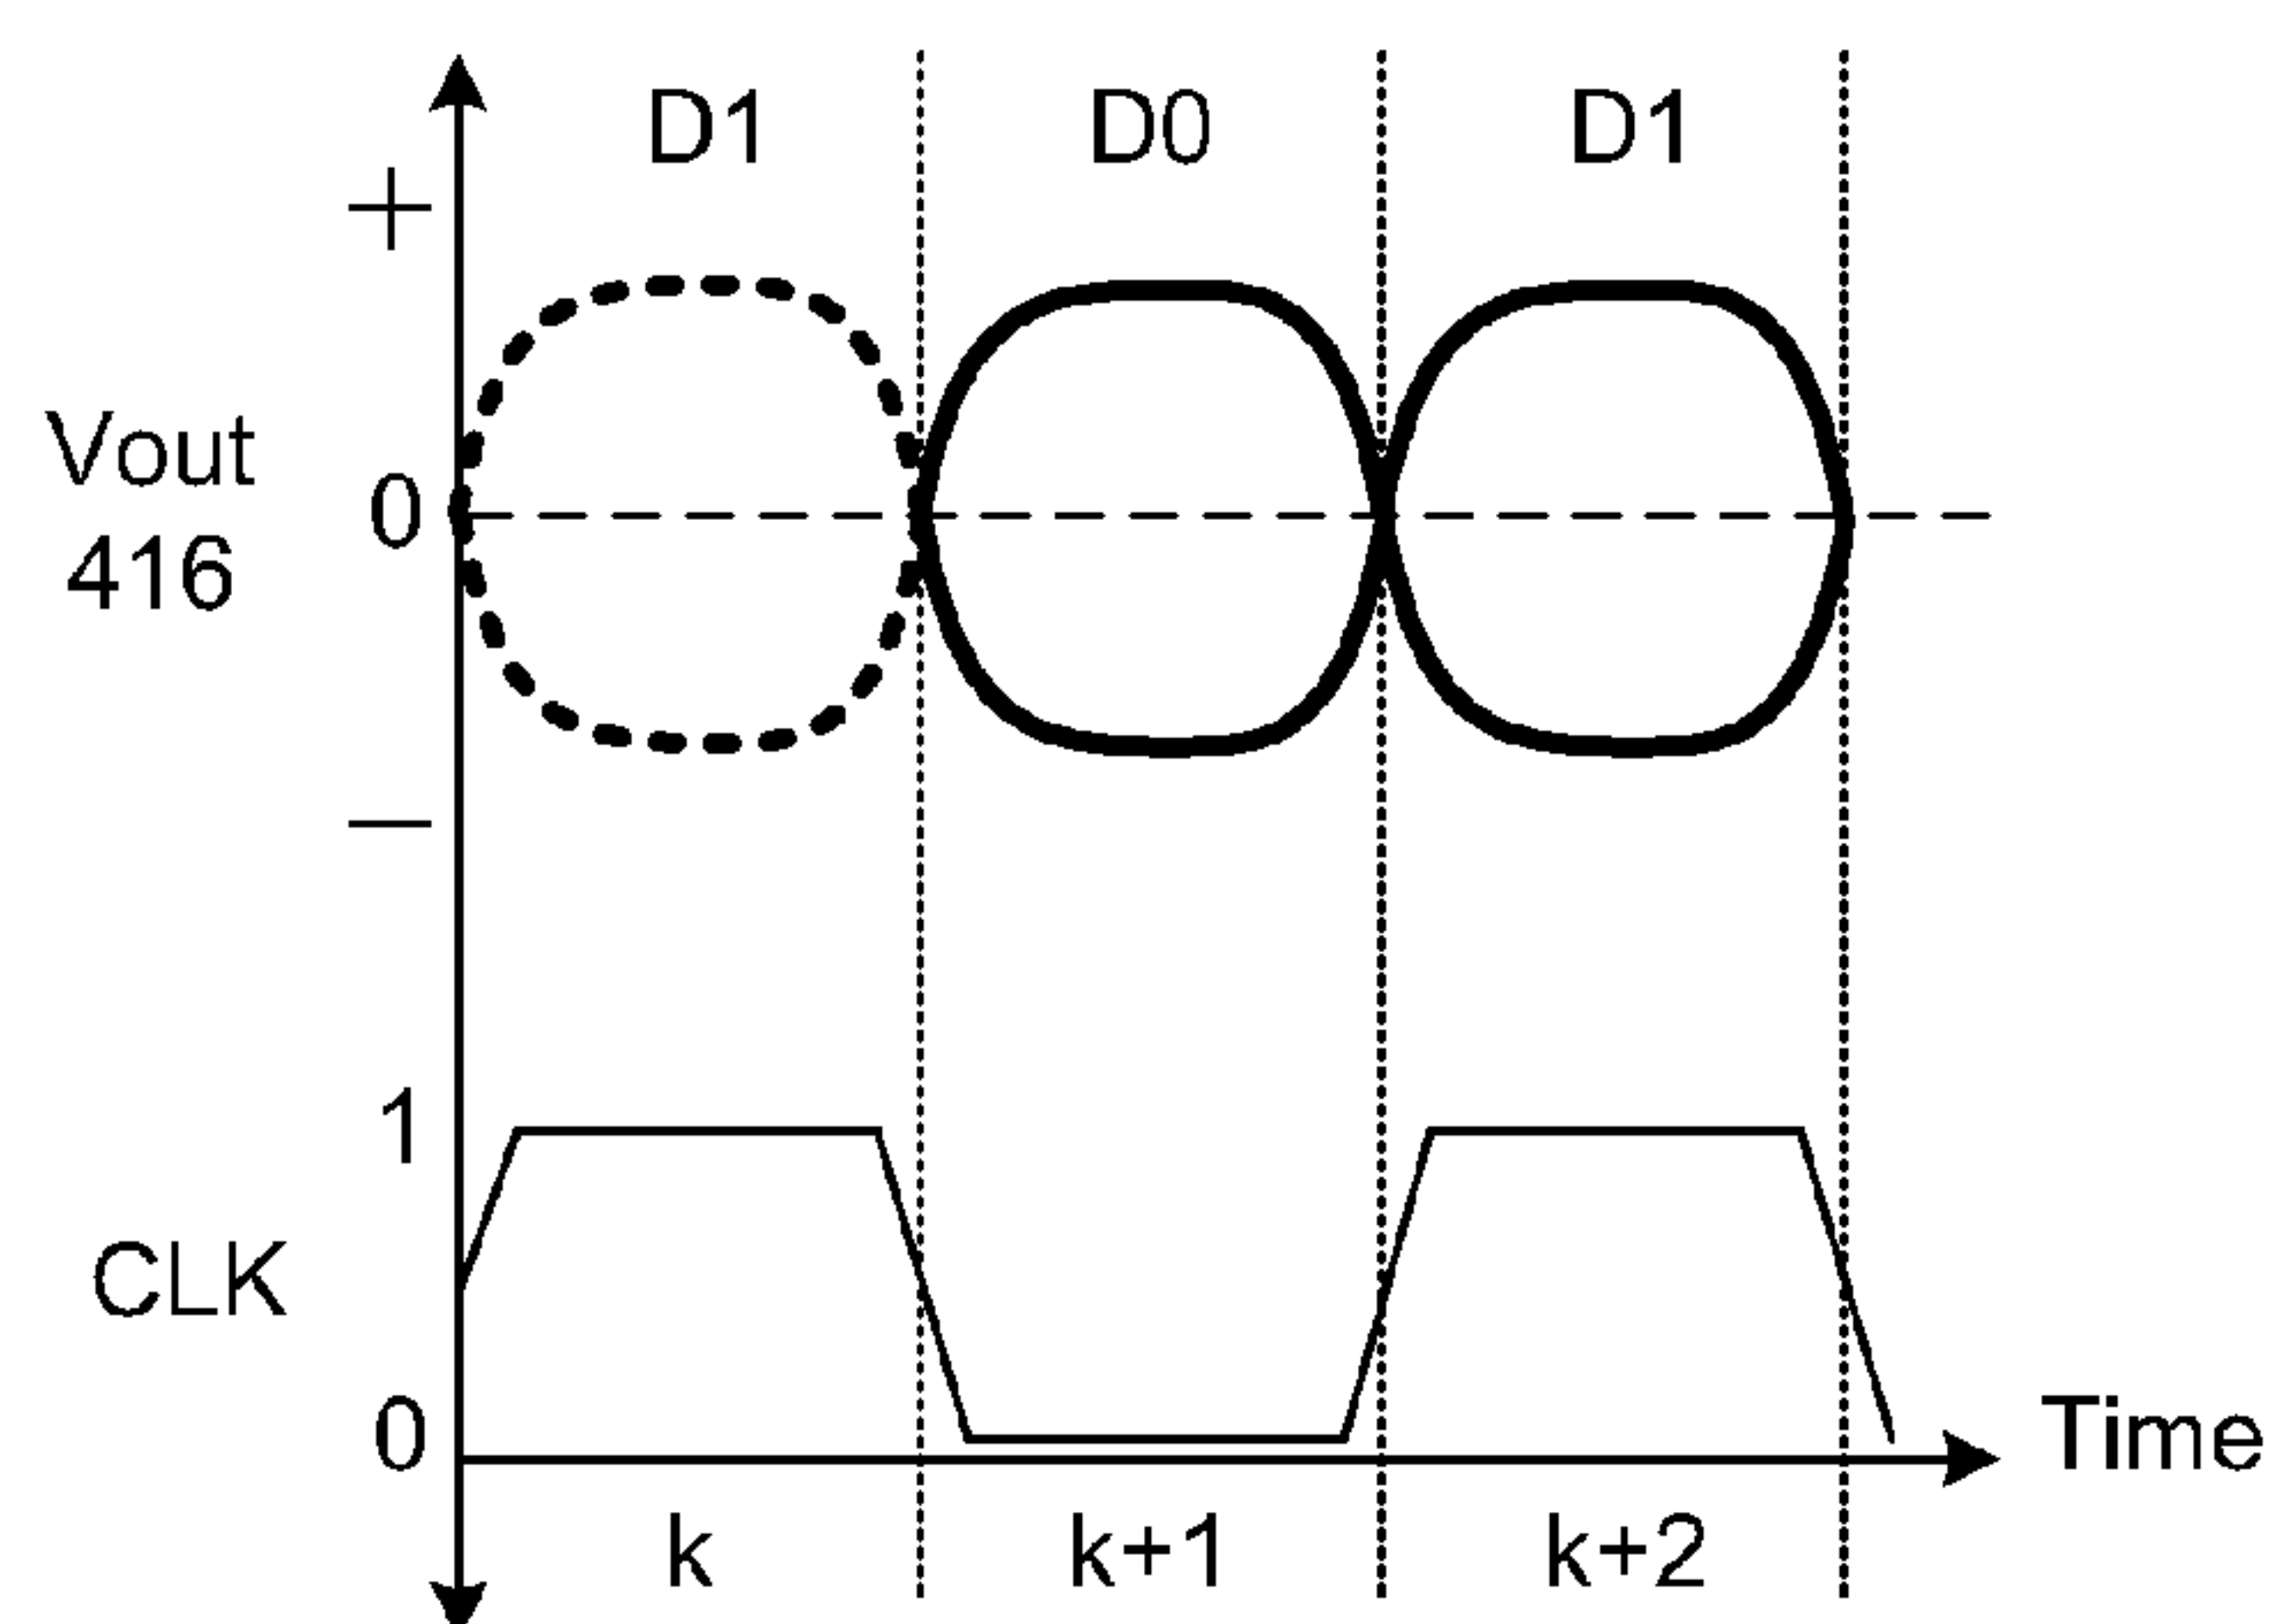
**Fig. 4C**



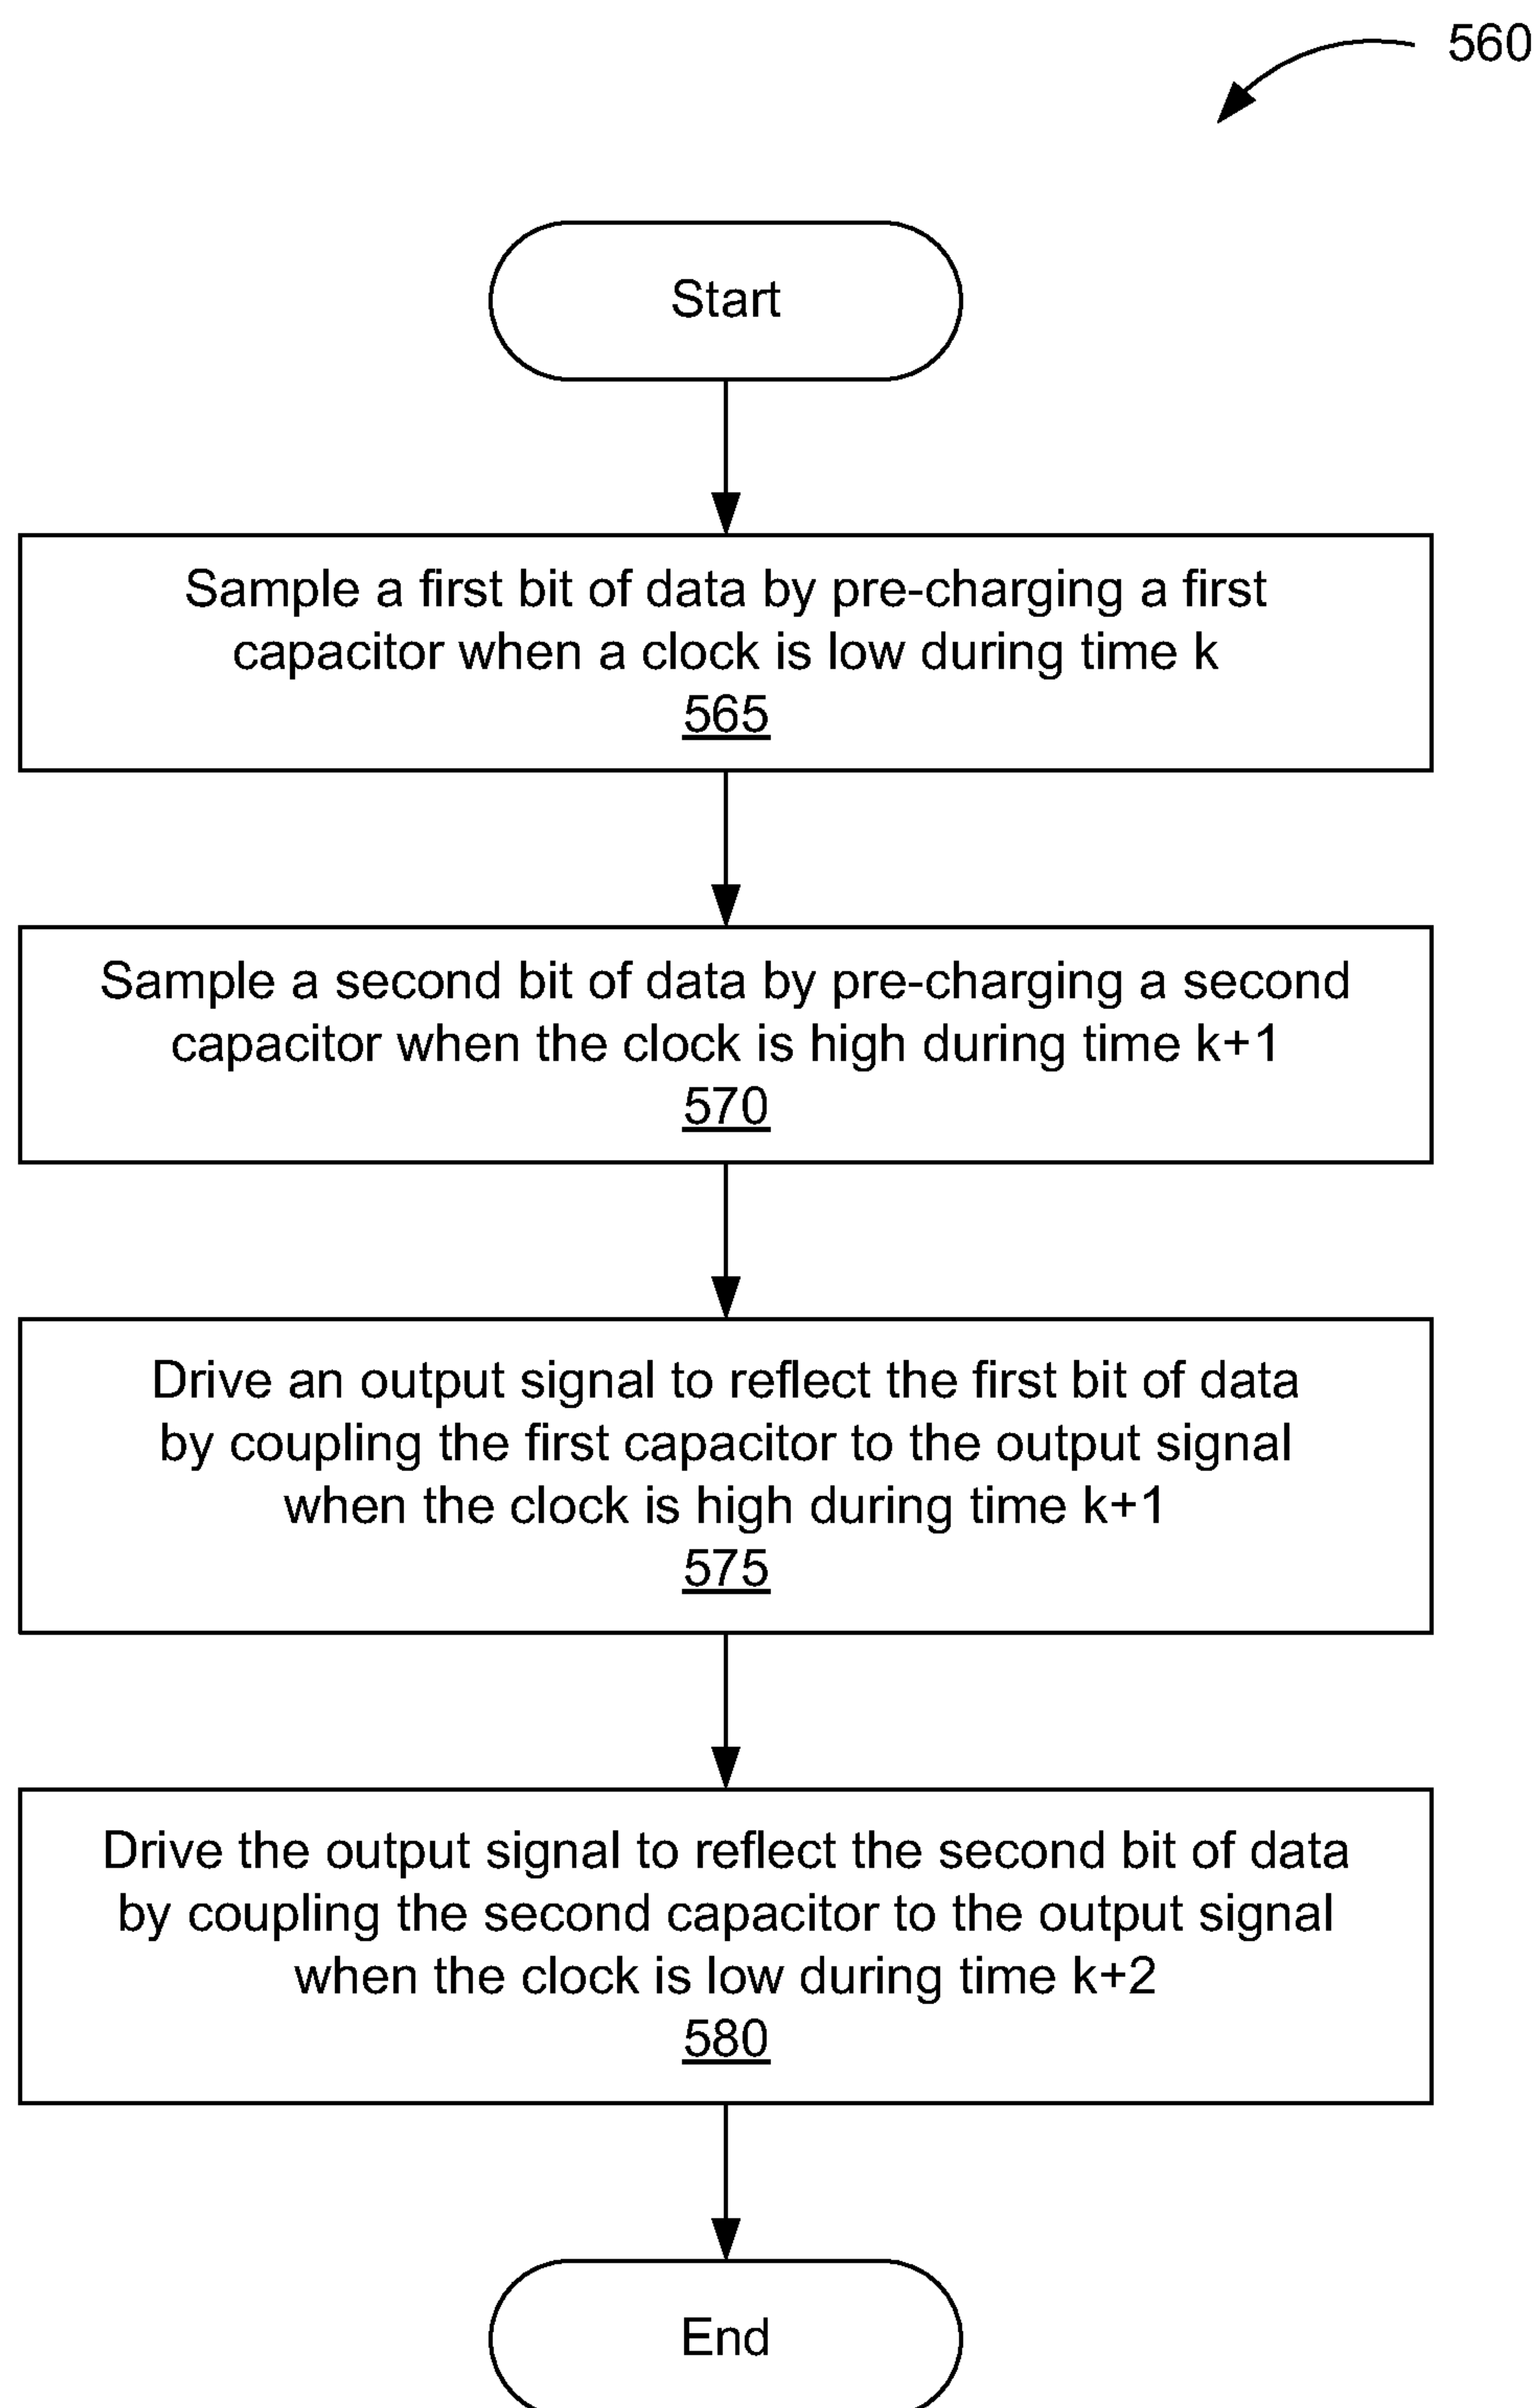
*Fig. 4D*



**Fig. 5A**



**Fig. 5B**

**Fig. 5C**

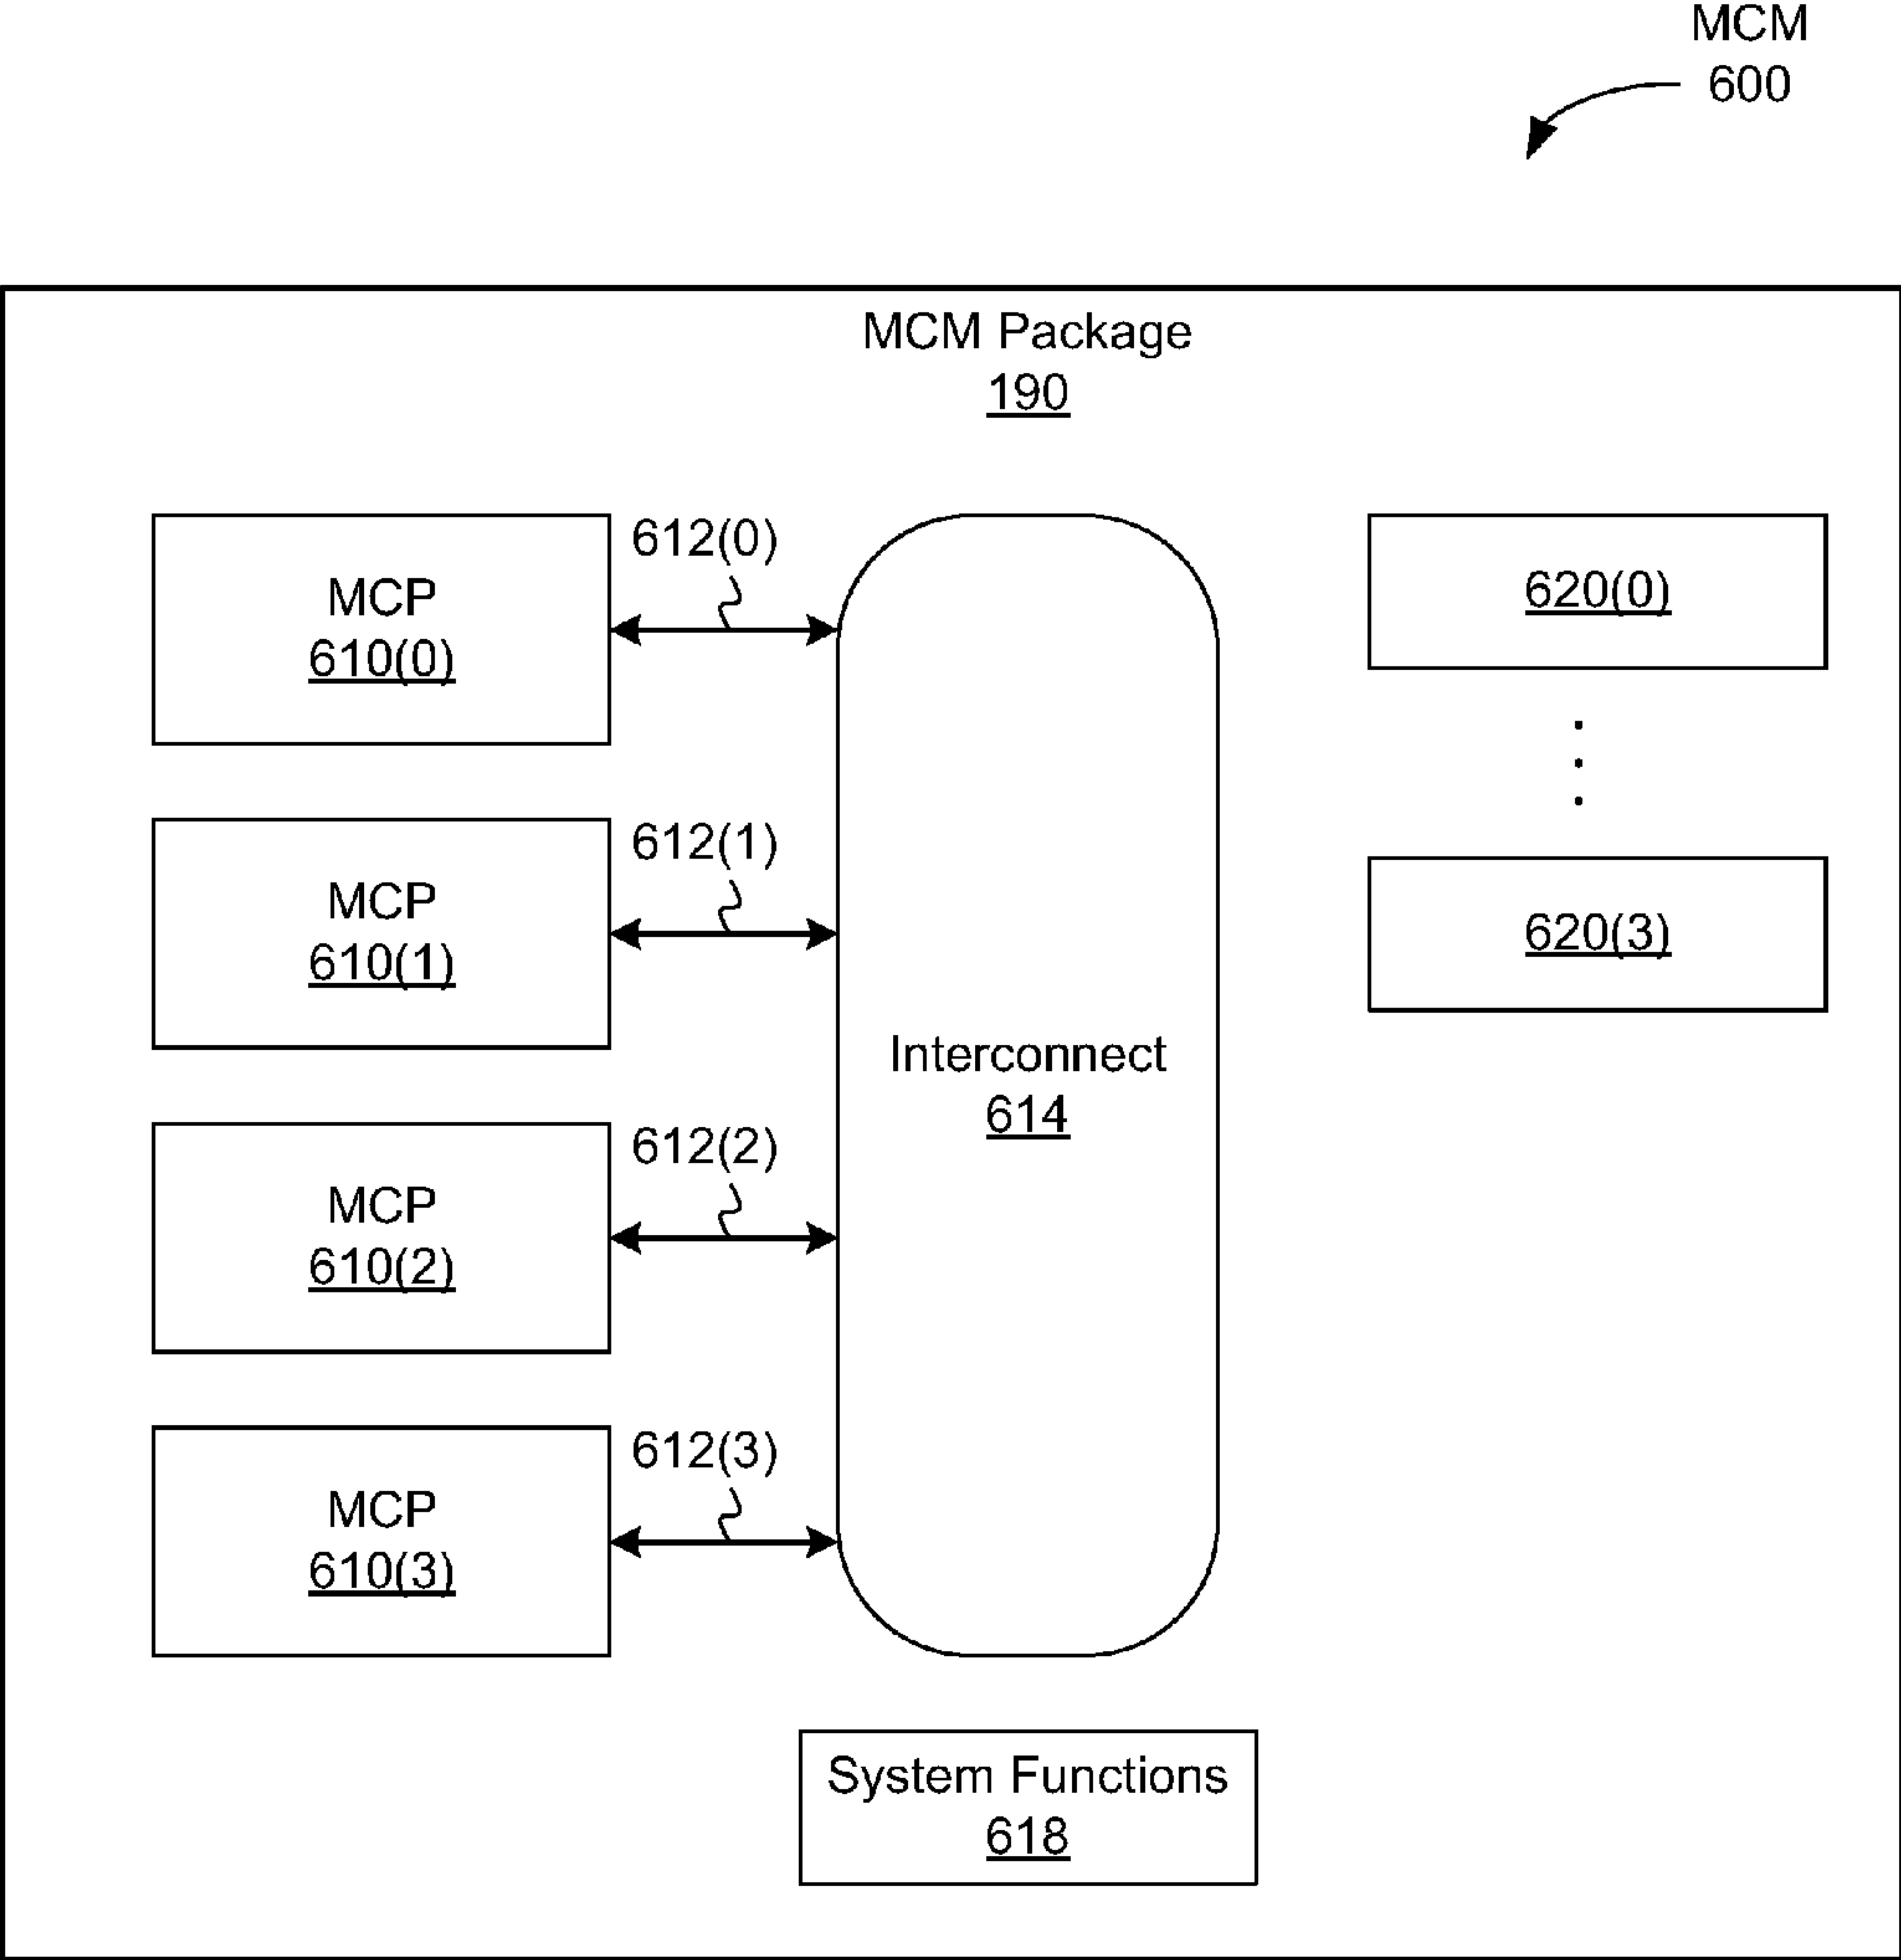
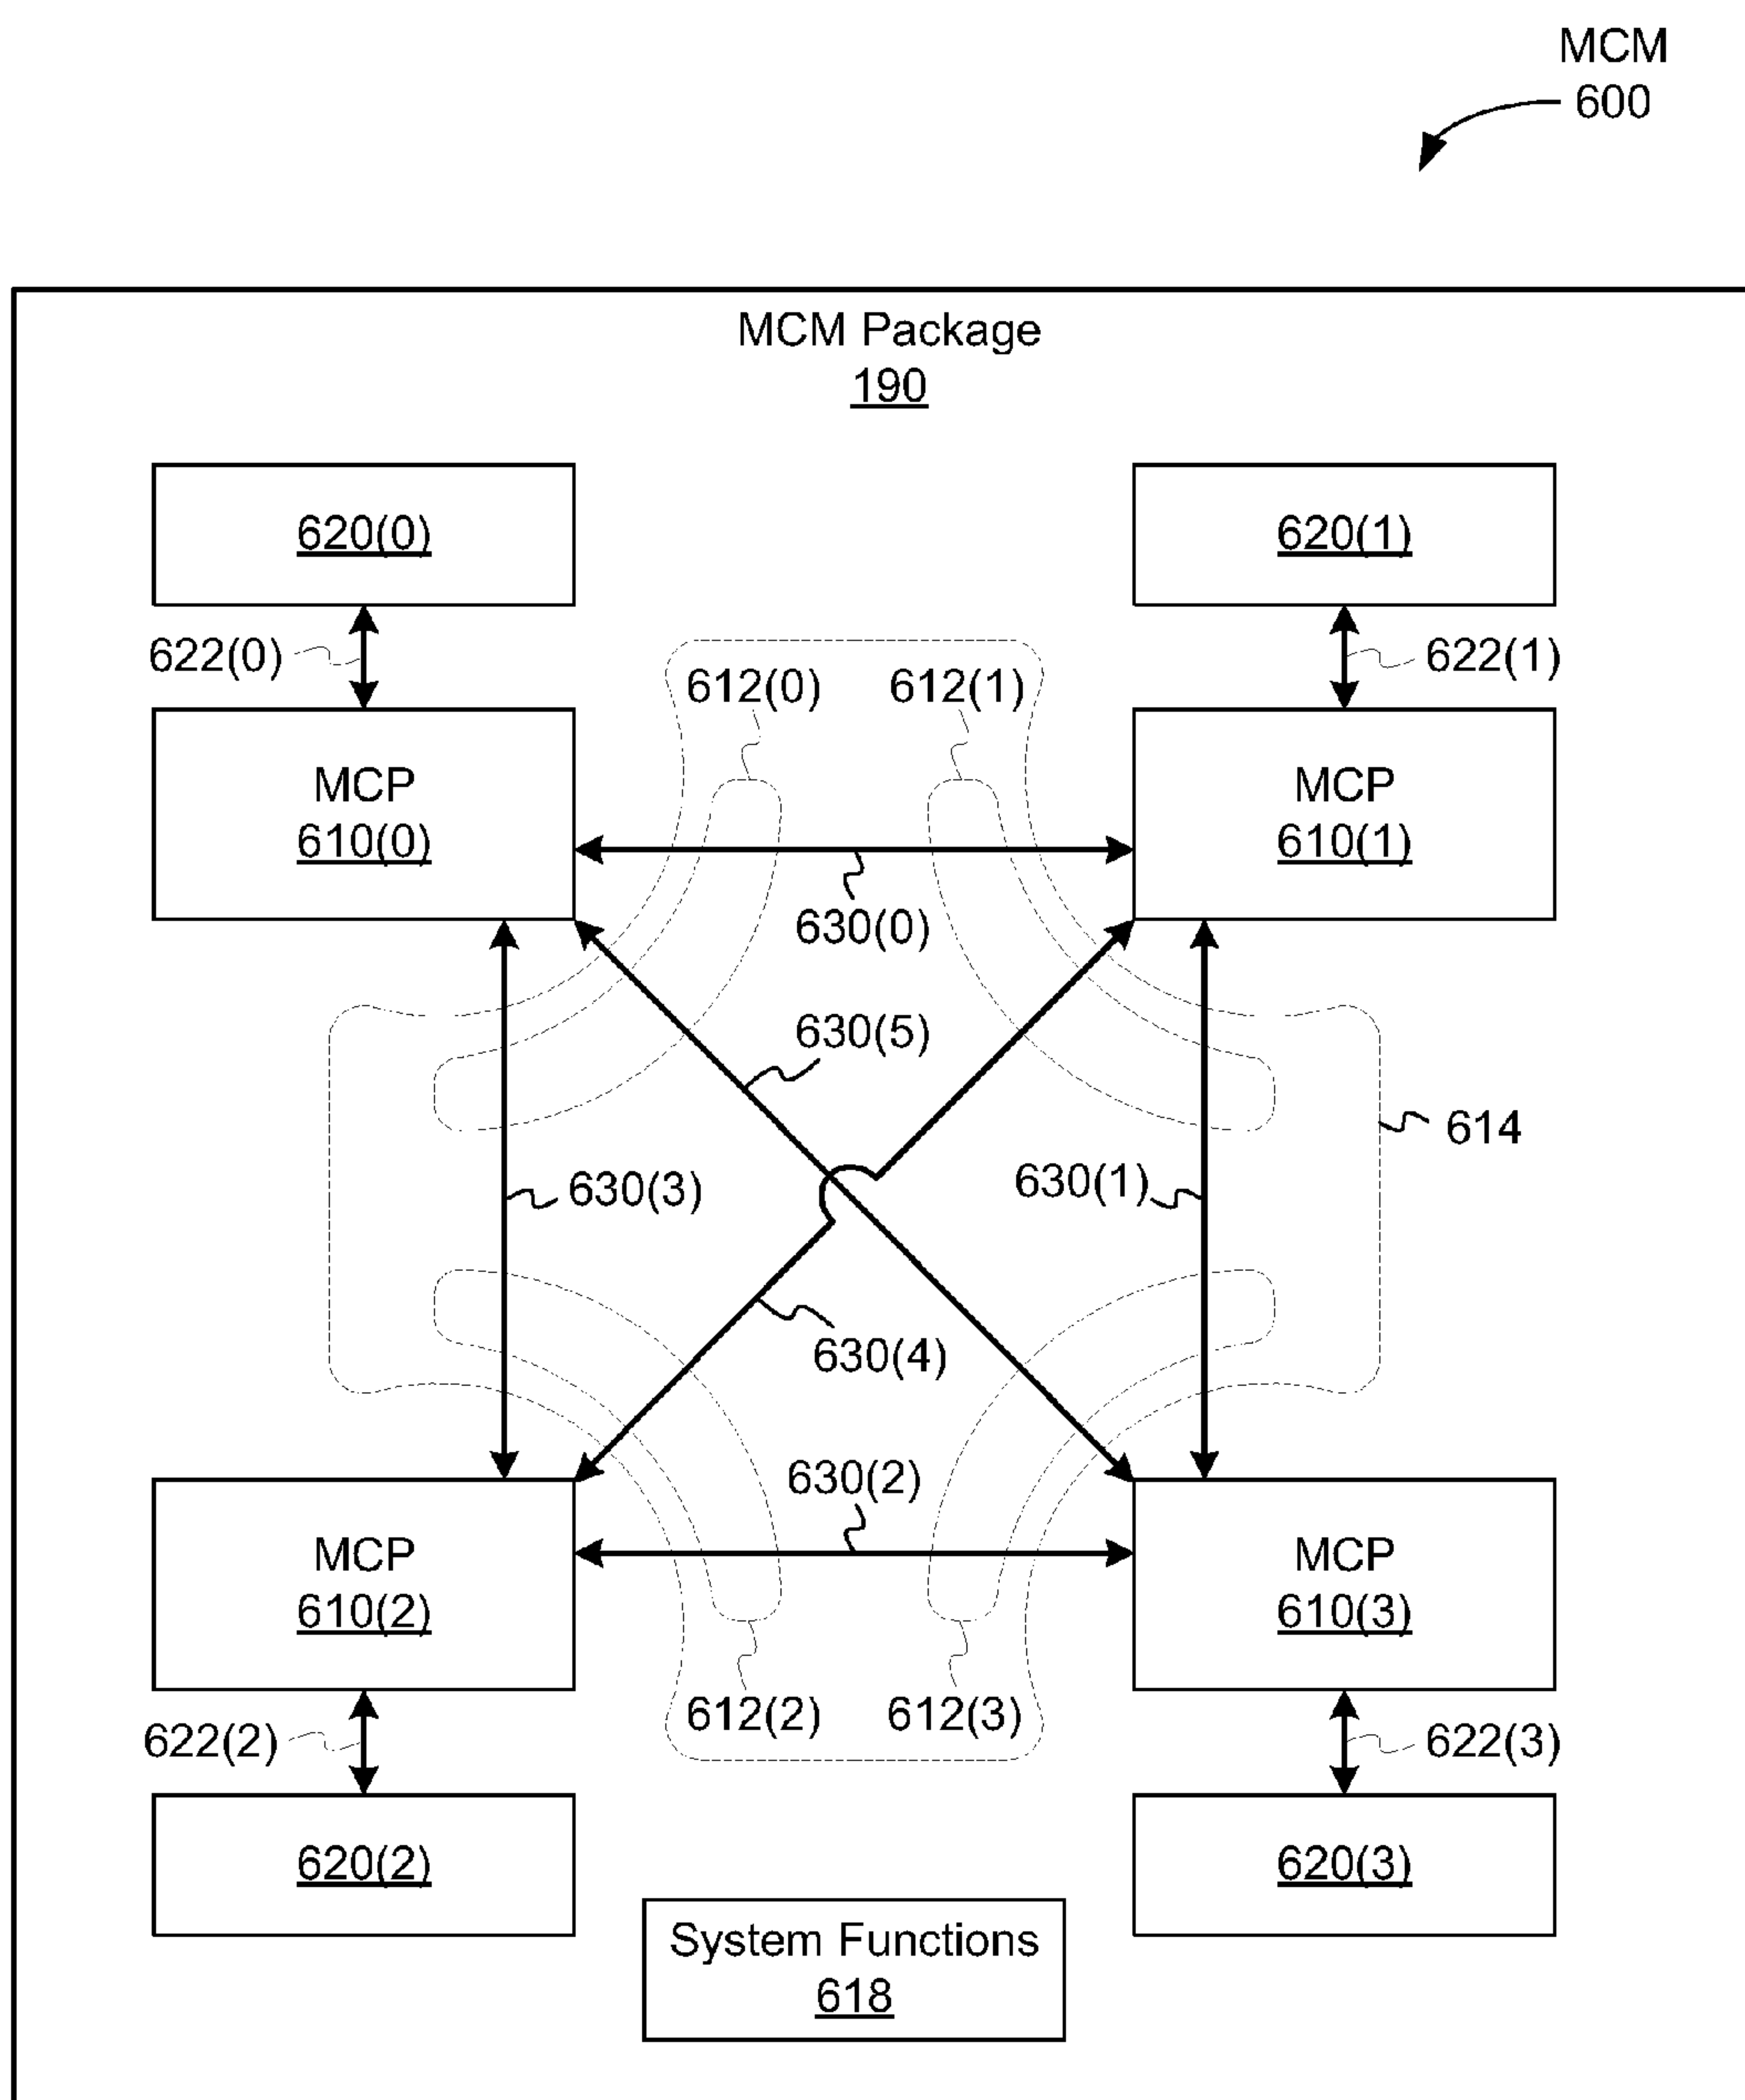


Fig. 6A





**Fig. 6B**

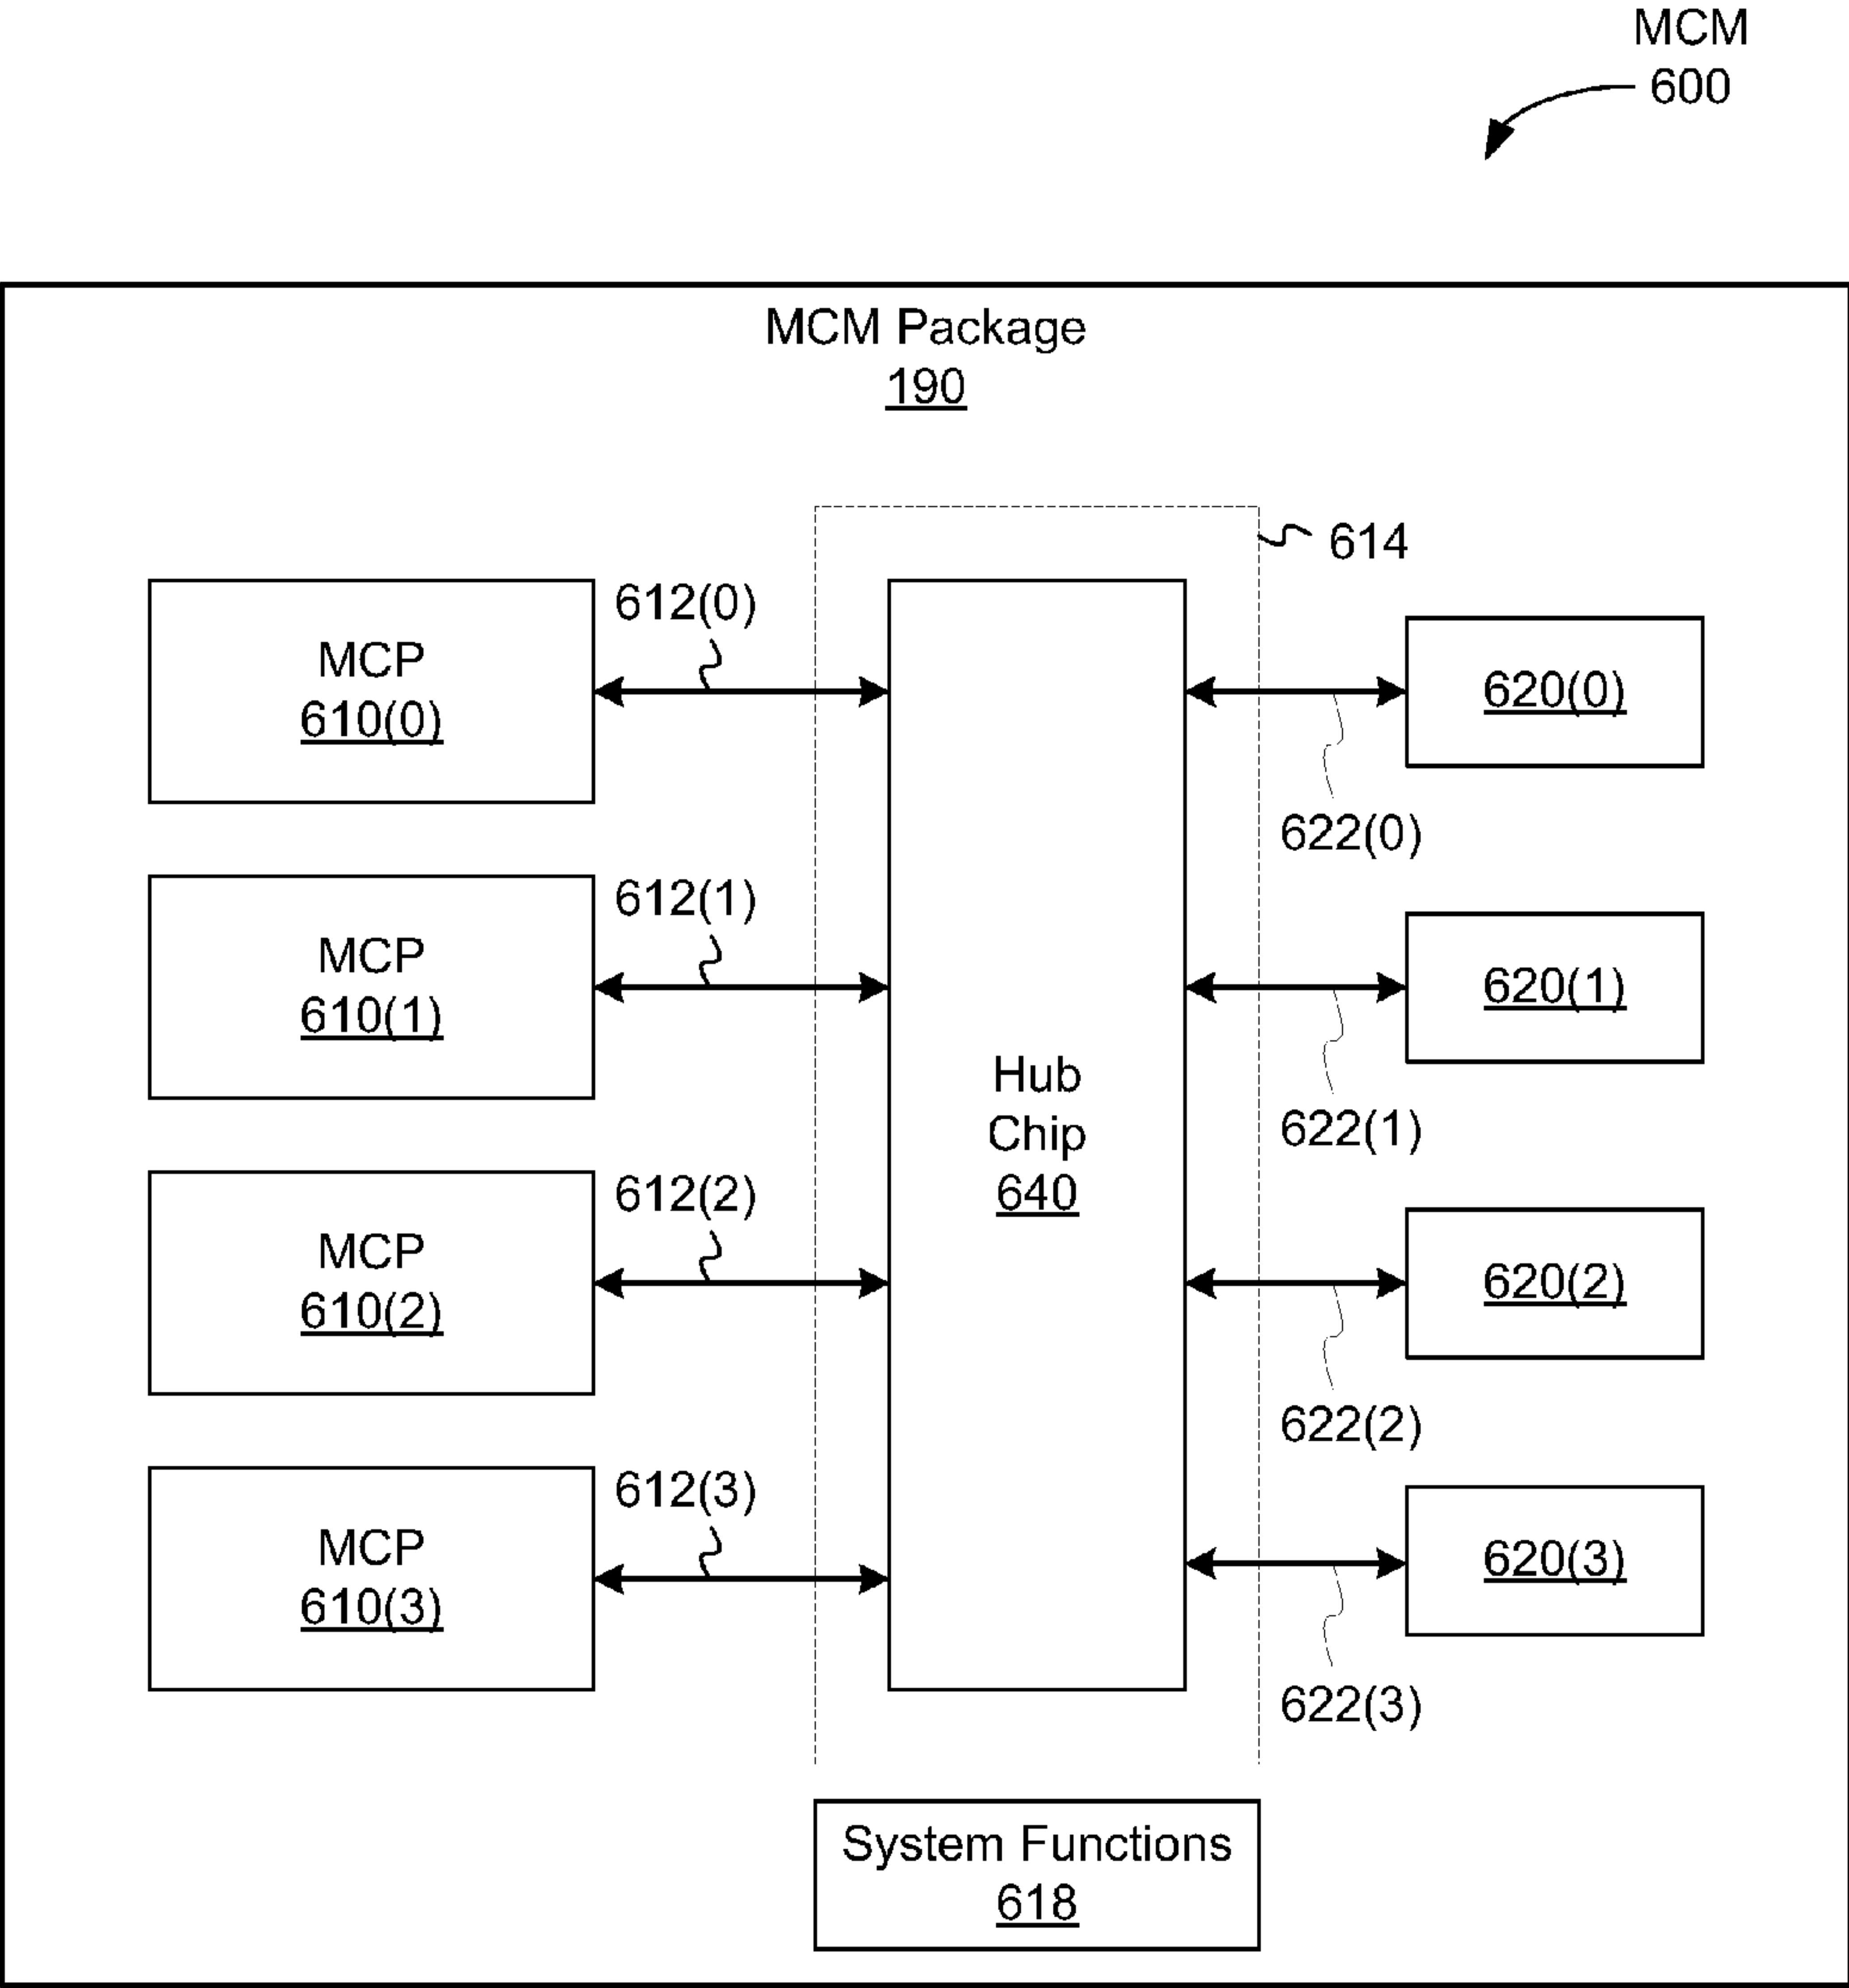


Fig. 6C

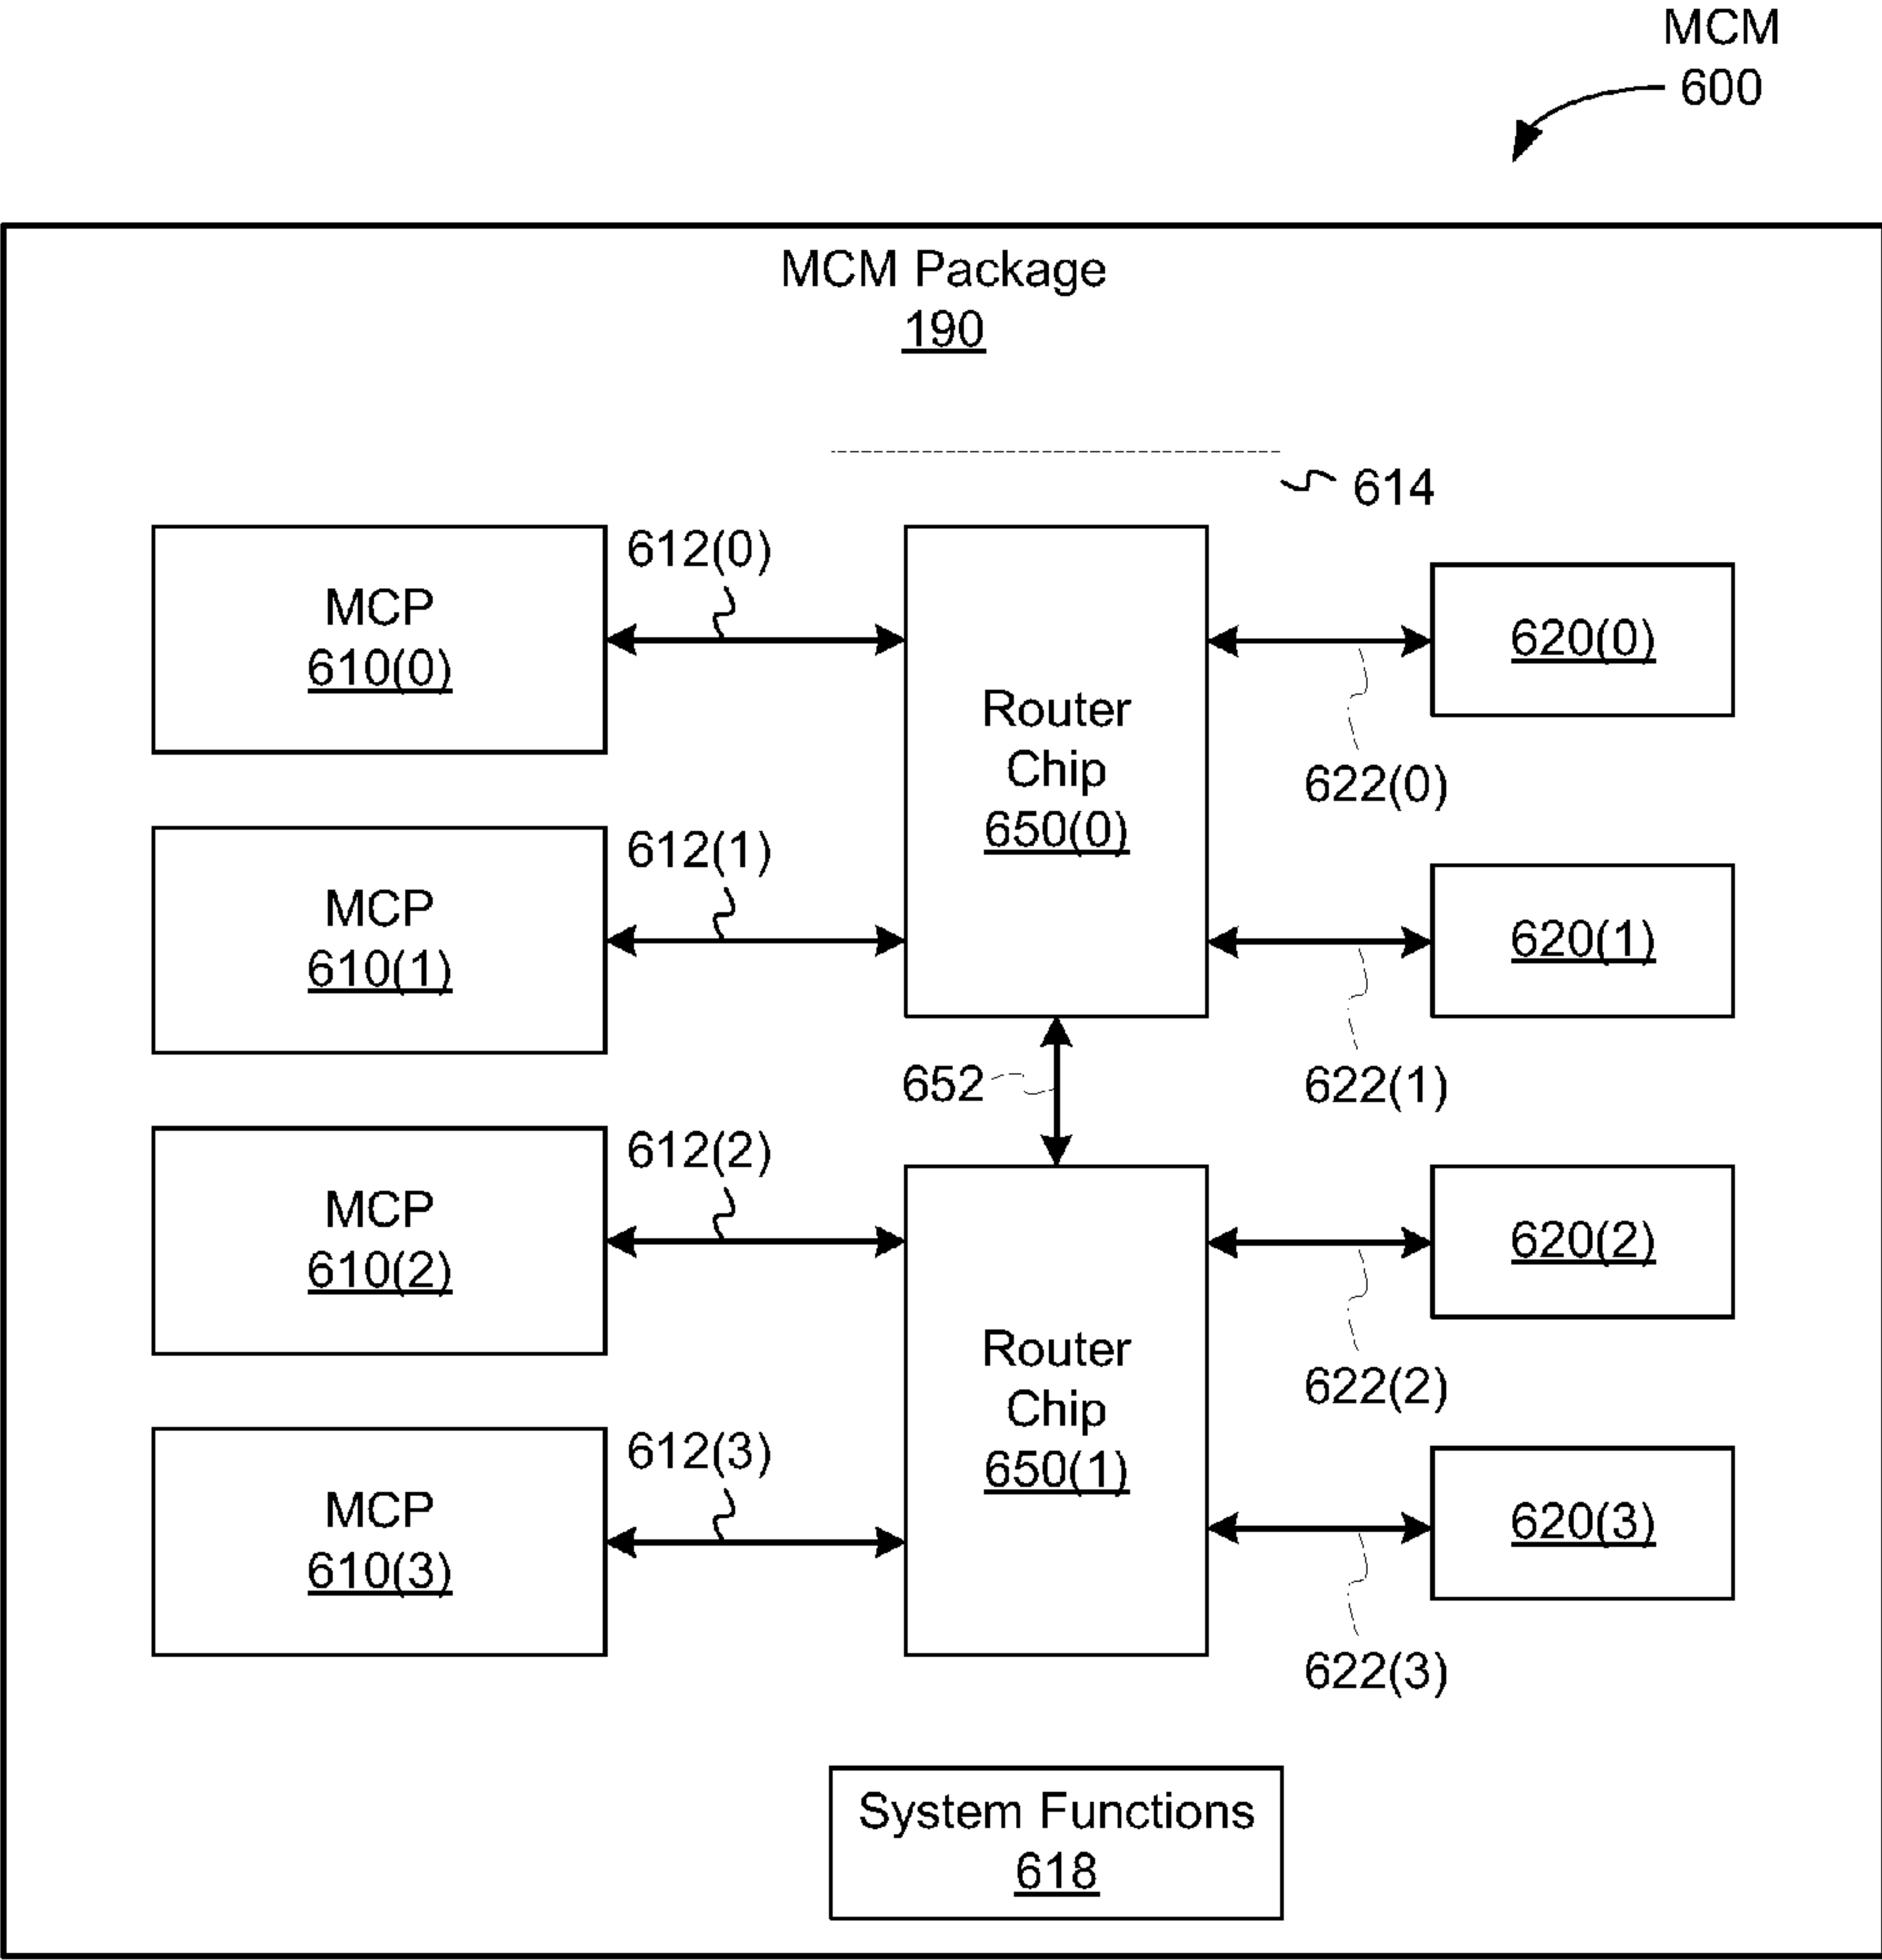
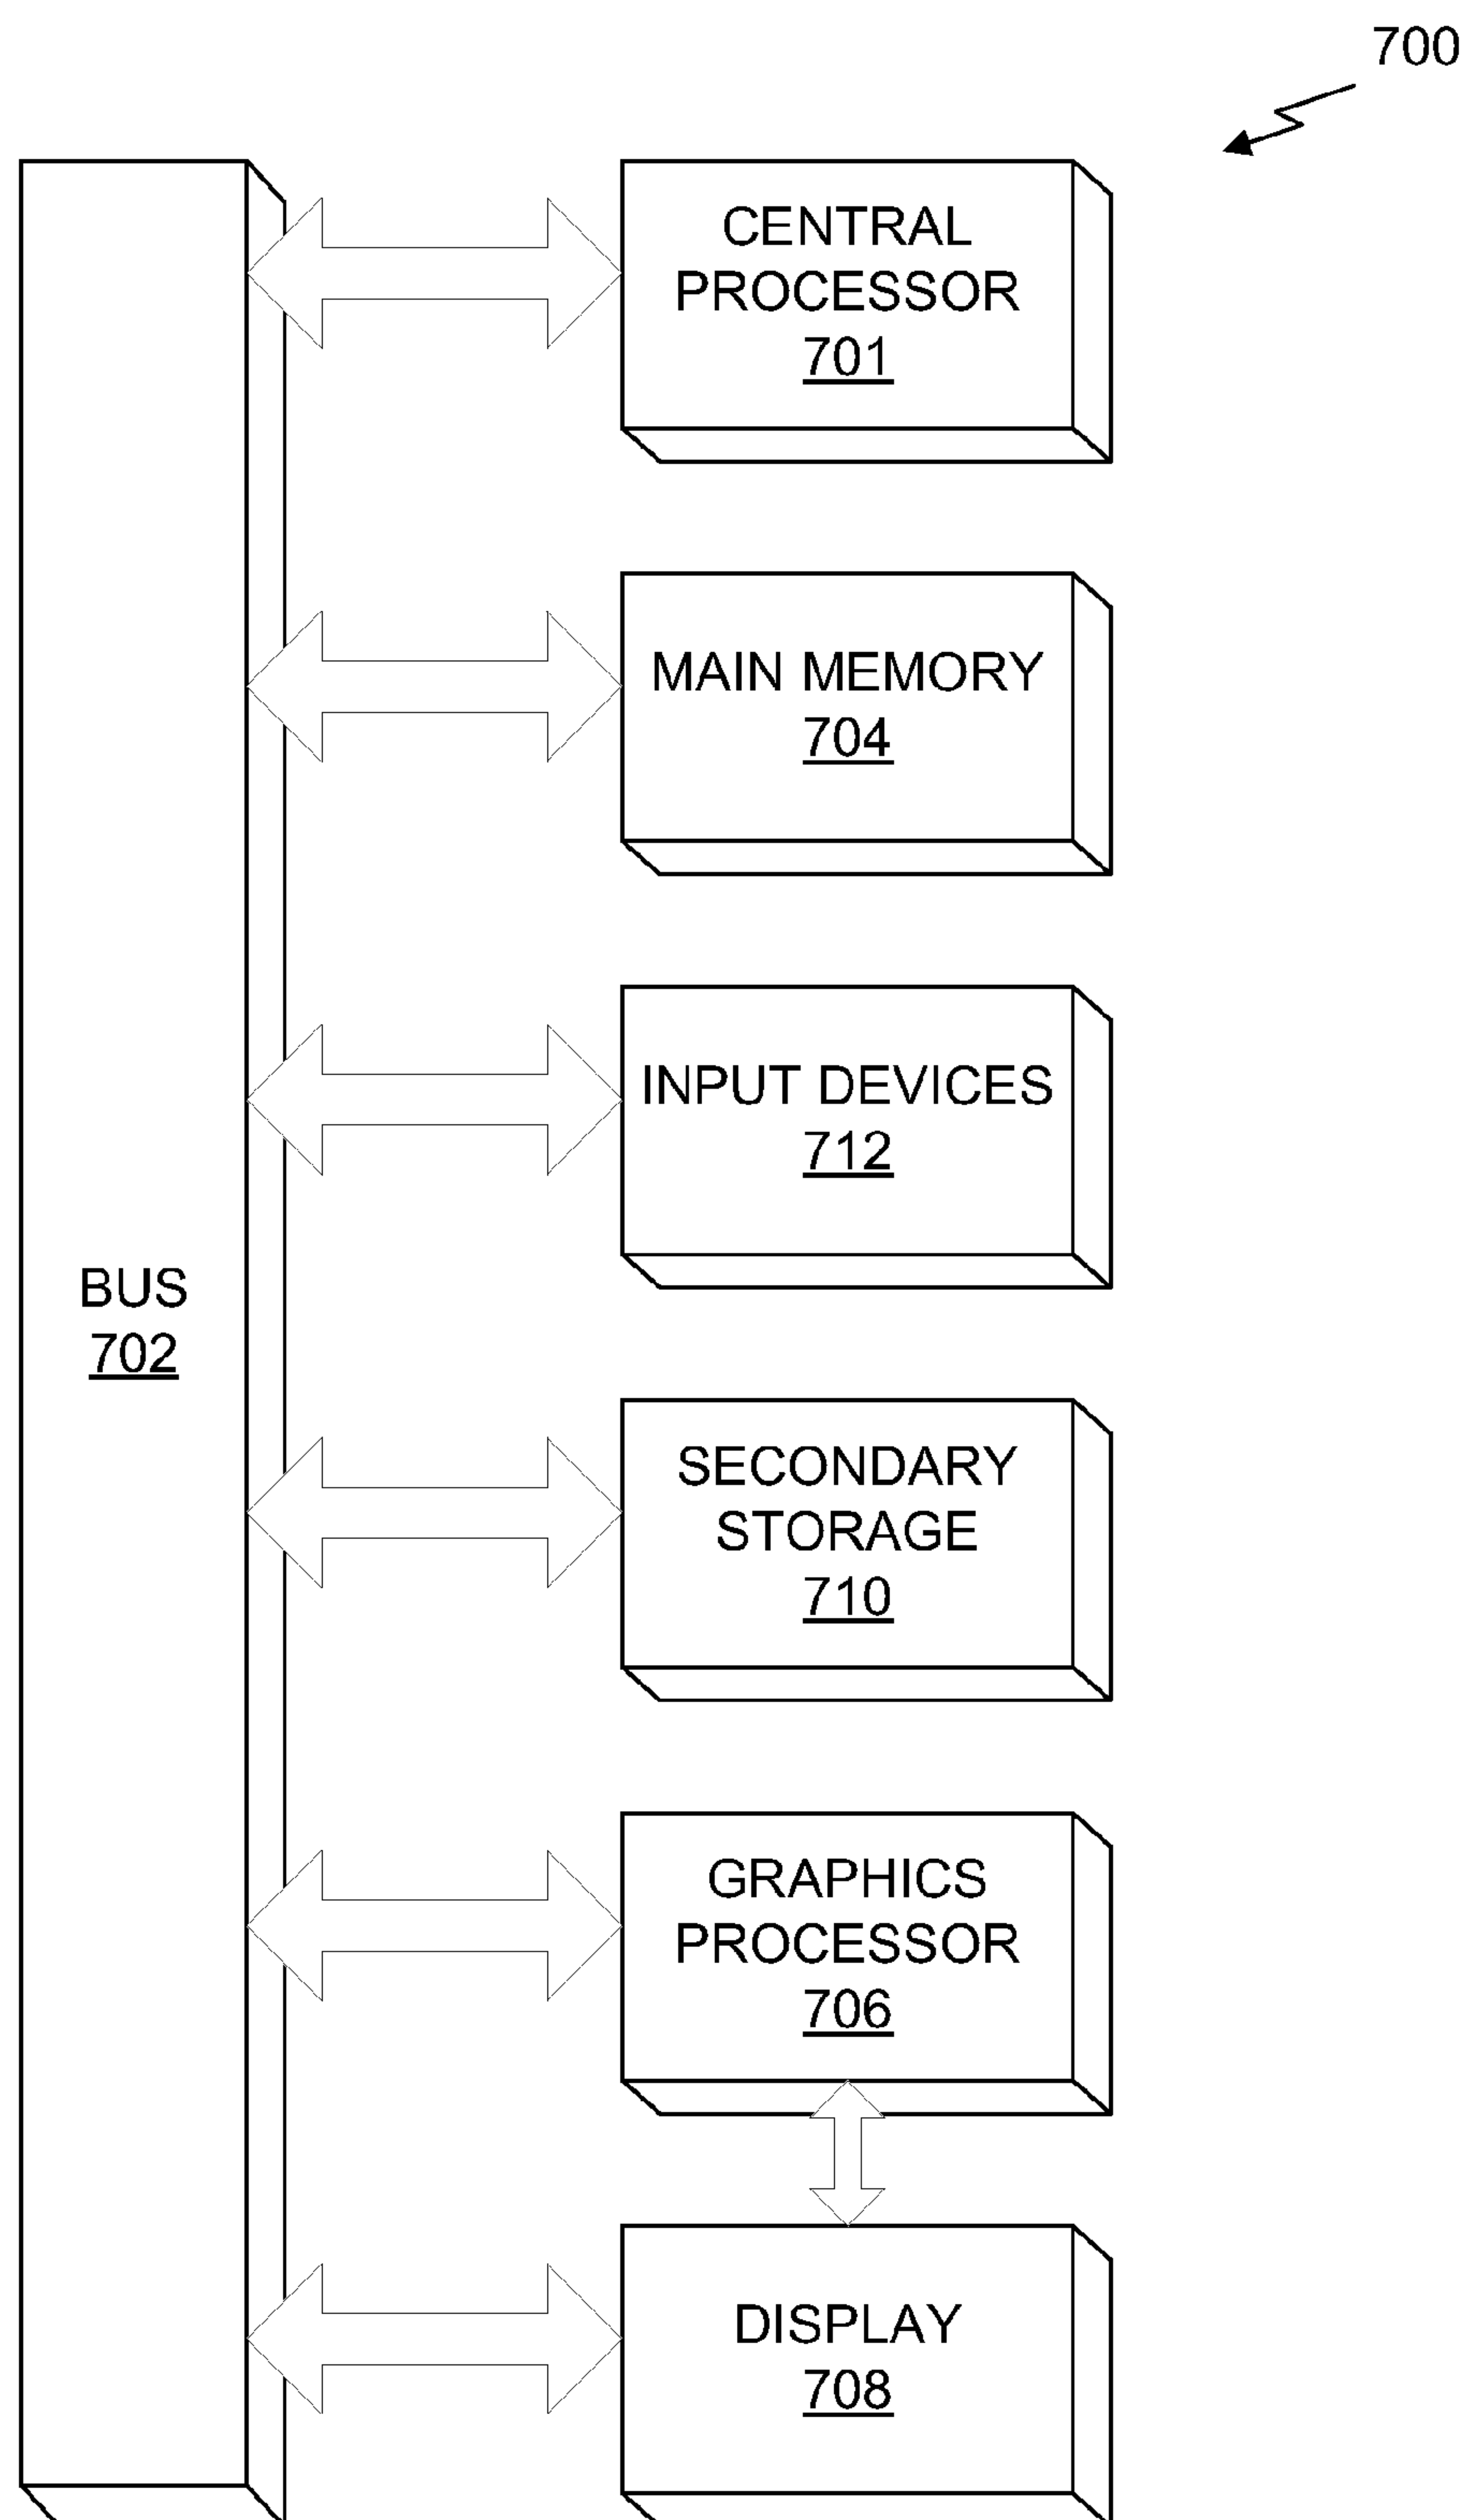


Fig. 6D



**Fig. 7**



# ON-PACKAGE MULTIPROCESSOR GROUND-REFERENCED SINGLE-ENDED INTERCONNECT

## CLAIM OF PRIORITY

[0001] This application is a continuation-in-part of U.S. application Ser. No. 13/844,570 (Attorney Docket No. NVIDP811/SC-13-0072-US1), filed Mar. 15, 2013, the entire contents of which is incorporated herein by reference.

## FIELD OF THE INVENTION

[0002] The present invention relates to multiprocessor architecture, and more specifically to an on-package multiprocessor ground-referenced single-ended interconnect.

## BACKGROUND

[0003] Sequential generations of computing systems typically require increasing degrees of performance and integration. A typical computing system includes a central processing unit (CPU), a graphics processing unit (GPU), a high-capacity memory subsystem, and set of interface subsystems. The set of interface subsystems may be configured to communicate with other devices, including devices that provide user interaction, devices that provide physical measurement, and devices that provide connectivity to storage systems and other computing systems.

[0004] Conventional computing systems typically achieve higher degrees of performance and integration by implementing an increasing number of processing cores on a single die or "chip." Additional cache memory may also be added to each processing core and as a resource shared by multiple processing cores. Measures of die area for multi-core devices have increased over time, as more CPU cores, GPU cores, on-chip cache memory, and additional interface blocks are integrated into a single processor chip. One advantage of integrating multiple processing cores and other subsystems onto a single die is that high-performance may be achieved by scaling conventional design techniques and leveraging advances in fabrication technology that enable greater circuit density.

[0005] However, one disadvantage of simply integrating more processing cores onto a single chip is that manufacturing cost for the chip typically increases disproportionately with respect to die area, increasing marginal cost associated with each additional processor core. More specifically, manufacturing cost for a given chip is typically a strong function of die area for the chip. In many cases, die area associated with highly-integrated multi-core processors is well above a characteristic cost knee, leading to disproportionate cost inefficiencies associated with multi-core processors. Alternatively, a computing system may be build from a plurality of independently packaged processing devices; however conventional chip-to-chip signaling techniques do not efficiently support multiprocessing performance targets commonly associated with highly-integrated multi-core devices.

[0006] Thus, there is a need for improving signaling and/or other issues associated with the prior art.

## SUMMARY

[0007] A system of interconnected chips comprising a multi-chip module (MCM) is disclosed. The system includes a first processor chip, a second processor chip, and an MCM package configured to include the first processor chip, the

second processor chip, and an interconnect circuit. The first processor chip is configured to include a first ground-referenced single-ended signaling (GRS) interface circuit. A first set of electrical traces fabricated within the MCM package and configured to couple the first GRS interface circuit to the interconnect circuit. The second processor chip is configured to include a second GRS interface circuit. A second set of electrical traces fabricated within the MCM package and configured to coupled the second GRS interface circuit to the interconnect circuit.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0008] FIG. 1A illustrates a ground-referenced single-ended signaling (GRS) system that implements a GRS transmitter based on a flying capacitor charge pump, in accordance with one embodiment;

[0009] FIG. 1B illustrates operation of a data driver in a pre-charge state and in two different data-dependent drive states, in accordance with one embodiment;

[0010] FIG. 1C illustrates a GRS system that implements a GRS transmitter based on a dual-capacitor charge pump, in accordance with one embodiment;

[0011] FIG. 1D illustrates operation of a data driver in a pre-charge state, in accordance with one embodiment;

[0012] FIG. 1E illustrates operation of a data driver in different data-dependent drive states, in accordance with one embodiment;

[0013] FIG. 1F illustrates operation of a ground-referenced single-ended data driver based on a flying capacitor charge pump, in accordance with one embodiment;

[0014] FIG. 1G illustrates operation of a ground-referenced single-ended data driver based on a dual capacitor charge pump, in accordance with one embodiment;

[0015] FIG. 2A illustrates an exemplary ground-referenced single-ended receiver, in accordance with one embodiment;

[0016] FIG. 2B illustrates an exemplary ground-referenced single-ended receiver, configured to demultiplex incoming data, in accordance with one embodiment;

[0017] FIG. 3 illustrates an exemplary transceiver pair, configured to implement ground-referenced single-ended signaling, in accordance with one embodiment;

[0018] FIG. 4A illustrates a ground-referenced single-ended data driver comprising a CMOS circuit, in accordance with one embodiment;

[0019] FIG. 4B illustrates a ground-referenced single-ended data driver in a pre-charge state associated with driving a data value of zero, in accordance with one embodiment;

[0020] FIG. 4C illustrates a ground-referenced single-ended data driver in a pre-charge state associated with driving a data value of one, in accordance with one embodiment;

[0021] FIG. 4D illustrates a ground-referenced single-ended data driver in a drive state, in accordance with one embodiment;

[0022] FIG. 5A illustrates a ground-referenced single-ended transmitter comprising two instances of a ground-referenced single-ended data driver, in accordance with one embodiment;

[0023] FIG. 5B illustrates timing for a ground-referenced single-ended transmitter comprising two ground-referenced single-ended data drivers, in accordance with one embodiment;

[0024] FIG. 5C illustrates a flow chart of a method for generating a ground-referenced single-ended signal, in accordance with one embodiment;



**[0025]** FIG. 6A illustrates a multiprocessor system implemented as a multi-chip module, in accordance with one embodiment;

**[0026]** FIG. 6B illustrates a directly-connected multiprocessor system implemented as a multi-chip module, in accordance with one embodiment;

**[0027]** FIG. 6C illustrates a hub-connected multiprocessor system implemented as a multi-chip module, in accordance with one embodiment;

**[0028]** FIG. 6D illustrates a network-connected multiprocessor system implemented as a multi-chip module, in accordance with one embodiment; and

**[0029]** FIG. 7 illustrates an exemplary system in which the various architecture and/or functionality of the various previous embodiments may be implemented.

#### DETAILED DESCRIPTION

**[0030]** A technique is provided for high-speed, single-ended signaling between different chips comprising a system-on-package device. A ground-referenced driver transmits a pulse having a polarity determined by a corresponding logic state. The pulse traverses a signal path and is received by a ground-referenced amplifier, which amplifies the pulse for interpretation as a conventional logic signal. Sets of ground-referenced drivers and ground-referenced amplifiers implement high-speed interfaces configured to interconnect different chips comprising the system-on-package device. The high-speed communication enabled by ground-referenced signaling advantageously improves bandwidth between different chips within the system-on-package device, enabling higher performance and higher density systems than provided by conventional signaling techniques.

**[0031]** Embodiments of the present invention implement a system comprising a plurality of different processor chips, one or more memory chips, and feature-specific chips coupled to a multi-chip package. Interconnections between the different chips are routed through the multi-chip package. At least one of the interconnections is configured to implement a ground-referenced single-ended signaling (GRS) link, described below.

**[0032]** A GRS data driver implements a charge pump driver configured to transmit a ground-referenced pulse on an associated signal line. In one implementation, a pulse of positive charge indicates a logical one, while a pulse of negative charge indicates a logical zero. The charge pump driver eliminates simultaneous switching noise (SSN) commonly associated with single-ended signaling by forcing transient signal current and ground current to be locally balanced, and by drawing a constant amount of charge from the power supply each half clock cycle, independent of the data being transmitted. The pulse is received and amplified by a common gate amplifier stage configured to use a local ground signal as an input reference. This configuration provides substantial immunity to common mode noise, the dominant source of transmission errors in single-ended signaling. A second amplifier stage translates a given received pulse to full-swing logic voltages, allowing the received pulse to be properly interpreted as one or two logic states by conventional logic circuitry. In one embodiment, a GRS receiver comprises a common gate amplifier stage, the second amplifier stage, and two storage elements, such as flip-flops, configured to capture received data during alternate clock phases.

**[0033]** A GRS transceiver includes a GRS data driver and a GRS receiver. The GRS transceiver transmits outbound data

through the GRS data driver and receives inbound data through the GRS receiver. An isochronous GRS transceiver may also transmit clocking information having a fixed phase relationship to the outbound data and receives clocking information having a fixed phase relationship to the inbound data. A GRS interconnect includes two different GRS transceivers, coupled through an electrical trace that is manufactured within a common multi-chip module package.

**[0034]** FIG. 1A illustrates a ground-referenced single-ended signaling (GRS) system **100** that implements a GRS transmitter **110** based on a flying capacitor charge pump, in accordance with one embodiment. GRS system **100** includes GRS transmitter **110**, a transmission path comprising a signal line **105** and a ground network **107**, and a GRS receiver **130**. In one embodiment, GRS transmitter **110** comprises two data drivers **112**, **114**. Input data signals **D0** and **D1** are presented to GRS transmitter **110** based on a clock signal **CLK**. Data driver **112** is configured to capture a logic state associated with input **D0** and drive output signal **Vout 116** onto signal line **105** with a pulse corresponding to the logic state of input **D0** while **CLK** is low. Similarly, data driver **114** is configured to capture a logic state associated with input **D1** and drive output signal **Vout 116** onto signal line **105** with a pulse corresponding to the logic state of **D1** while **CLK** is high. A sequence of pulses is formed along signal line **105** corresponding to a sequence of input data from inputs **D0** and **D1**. The sequence of pulses is referenced to ground with a voltage swing that may be lower than conventional logic voltage swings. GRS receiver **130** is configured to amplify an incoming sequence of pulses from signal line **105** and translate the pulses to a conventional logic voltage swing so the pulses may be properly interpreted as logic signals on amplifier output signal **132**. For example, the sequence of pulses along signal line **105** may have a nominal amplitude of plus or minus one-hundred millivolts, while amplifier output signal **132** may have a corresponding voltage swing of twelve hundred millivolts to zero volts with respect to ground if logic coupled to amplifier output signal **132** operates on a twelve hundred millivolt positive supply rail.

**[0035]** In one embodiment, GRS transmitter **110** is fabricated on a transmitter chip and GRS receiver **130** is fabricated on a receiver chip distinct from the transmitter chip. Pads **120** comprise bonding pads configured to couple output signal **Vout 116** from the transmitter chip to signal line **105**, which is fabricated as an impedance-controlled trace within a multi-chip module (MCM) package **190**. Pads **122** comprise bonding pads configured to couple a local ground signal within the transmitter chip to ground network **107**, fabricated within MCM package **190**. Similarly, pads **124** comprise bonding pads configured to couple signal line **105** to an input signal for GRS receiver **130** within the receiver chip, and pads **126** comprise bonding pads configured to couple ground network **107** to a local ground within the receiver chip. A termination resistor **RTx** is coupled between output signal **Vout 116** and the local ground within the transmitter chip to absorb incoming signals, such as reflections or induced noise signals. A termination resistor **RRx** is coupled across inputs to GRS receiver **130** to similarly absorb incoming signals at the receiver chip.

**[0036]** Data driver **112** comprises capacitor **C0**, and switches **S01** through **S06**. Switch **S01** enables a first node of capacitor **C0** to be coupled to a positive supply rail, while switch **S02** enables a second node of capacitor **C0** to be coupled to a local ground net. Switches **S01** and **S02** are



active (closed) during a pre-charge state for data driver **112**, defined when CLK is equal to a logical “1” value. Switch **S03** enables the first node of capacitor **C0** to be coupled to GND, while switch **S06** enables the second node of capacitor **C0** to be coupled to GND. Switch **S04** enables the first node of capacitor **C0** to be coupled to Vout **116**, while switch **S05** enables the second node of capacitor **C0** to be coupled to Vout **116**. When CLK is equal to a logical “0” value, switches **S04** and **S06** are active when data driver **112** is driving a logical “1” value to Vout **116**, or switches **S03** and **S05** are active when data driver **112** is driving a logical “0” value to Vout **116**. Data driver **114** comprises a substantially identical circuit topology, with an inverted sense for CLK, so that data driver **114** is in a pre-charge state when CLK is equal to a logical “0” value and driving Vout **116** when CLK is equal to a logical “1” value.

[0037] In one embodiment, switches **S01** through **S06** and switches **S11** through **S16** are fabricated using monolithic complementary metal-oxide semiconductor (CMOS) devices, such as enhancement mode n-channel and p-channel field-effect transistors. Any technically feasible logic circuit topologies may be implemented to drive switches **S01-S06** and switches **S11-S16** into individually active or inactive states without departing the scope and spirit of embodiments of the present invention.

[0038] FIG. 1B illustrates operation of a data driver **112** in a pre-charge state and in two different data-dependent drive states, in accordance with one embodiment. As shown, when CLK is equal to a logical “1” value, data driver **112** is in a pre-charge state, whereby switches **S01** and **S02** are active and capacitor **C0** charges to a voltage corresponding approximately to a positive supply rail, such as a “VDD” supply rail. All of switches **S03-S06** are inactive (open) during the pre-charge state. When CLK is equal to a logical “0” value, two of switches **S03-S06** are configured to couple capacitor **C0** to Vout **116** to transmit a pulse having a polarity corresponding to a logical value for **D0**. To drive a logical “0” value, switches **S03** and **S05** are driven active, thereby coupling a negative charge relative to ground onto Vout **116**. To drive a logical “1” value, switches **S04** and **S06** are driven active, thereby coupling a positive charge relative to ground onto Vout **116**.

[0039] FIG. 1C illustrates a GRS system **102** that implements a GRS transmitter **150** based on a dual-capacitor charge pump, in accordance with one embodiment. GRS system **102** includes GRS transmitter **150**, a transmission path comprising a signal line **105** and a ground network **107**, and a GRS receiver **130**. In one embodiment, GRS transmitter **150** comprises two data drivers **152** and **154**. Operation of GRS system **102** is substantially identical to the operation of GRS system **100** described above in FIGS. 1A and 1B, with the exception of the internal topology and operation of data drivers **152** and **154**.

[0040] Data driver **152** comprises capacitors **C0A** and **C0B**, as well as switches **S0A** through **S0H**. Switch **S0A** enables a first node of capacitor **C0A** to be coupled to a positive supply rail, while switch **S0C** enables the first node to be coupled to a local ground net. Switch **S0B** enables a second node of capacitor **C0A** to be coupled to Vout **116**, while switch **S0D** enables the second node to be coupled to the local ground net. Similarly, switch **S0E** enables a first node of capacitor **C0B** to be coupled to the positive supply rail, while switch **S0G** enables the first node to be coupled to the local ground net. Switch **S0F** enables a second node of capacitor **C0B** to be

coupled to Vout **116**, while switch **S0H** enables the second node to be coupled to the local ground net.

[0041] A pre-charge state for data driver **152** is defined when CLK is equal to a logical “1” value. During the pre-charge state, switches **S0A**, **S0D**, **S0G**, and **S0H** are driven active, pre-charging capacitor **C0A** to a voltage corresponding to the positive supply rail relative to the local ground net, and pre-charging capacitor **C0B** to have approximately no charge. When CLK is equal to a logical “0” value, either capacitor **C0A** is coupled to Vout **116** to generate a negative pulse or capacitor **C0B** is coupled to Vout **116** to generate a positive pulse, as described below in conjunction with FIG. 1E. Data driver **154** comprises a substantially identical circuit topology, with an inverted sense for CLK, so that data driver **154** is in a pre-charge state when CLK is equal to a logical “0” value and driving Vout **116** when CLK is equal to a logical “1” value.

[0042] In one embodiment, switches **S0A** through **S0H** and switches **S1A** through **S1H** are fabricated using monolithic CMOS devices, such as enhancement mode n-channel and p-channel FETs. Any technically feasible logic circuit topologies may be implemented to drive switches **S0A-S0H** and switches **S1A-S1H** into individually active or inactive states without departing the scope and spirit of embodiments of the present invention.

[0043] FIG. 1D illustrates operation of data driver **152** in a pre-charge state, in accordance with one embodiment. As shown, when CLK is equal to a logical “1” value, switch **S0A** is active, coupling a first node of capacitor **C0A** to a positive supply rail, and switch **S0D** is active, coupling a second node of capacitor **C0A** to a local ground net. At the same time, switch **S0G** is active, coupling a first node of capacitor **C0B** to ground, and switch **S0H** is active, coupling a second node of capacitor **C0B** to ground. By the end of this pre-charge state, capacitor **C0B** is substantially discharged.

[0044] FIG. 1E illustrates operation of data driver **152** in different data-dependent drive states, in accordance with one embodiment. As shown, when CLK is equal to a logical “0” value and **D0** is equal to a logical “0” value, switches **S0C** and **S0B** are configured to couple capacitor **C0A** to Vout **116** to transmit a pulse having a negative polarity. Alternatively, when CLK is equal to a logical “0” value and **D0** is equal to a logical “1” value, switches **S0E** and **S0F** are configured to couple capacitor **C0B** to Vout **116** to transmit a pulse having a positive polarity. Here, the positive supply rail is assumed to have adequate high-frequency capacitive coupling to the local ground net to force transient return current through the local ground net in conjunction with driving Vout **116** with a positive pulse.

[0045] More illustrative information will now be set forth regarding various optional architectures and features with which the foregoing framework may or may not be implemented, per the desires of a designer or user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the exclusion of other features described.

[0046] FIG. 1F illustrates operation of a ground-referenced single-ended data driver **162** based on a flying capacitor charge pump, in accordance with one embodiment. One or more instances of data driver **162** may be configured to operate as data drivers within a GRS transmitter. For example, an instance of data driver **162** may be configured to operate in



place of data driver 112 within GRS transmitter 110 of FIG. 1A. Similarly, an instance of data driver 162 may be configured to operate in place of data driver 114.

[0047] Data driver 162 includes capacitor C2, and switches S20, S21, S22, S23, and S24, configured to pre-charge capacitor C2 during a pre-charge phase, and discharge capacitor C2 into Vout 116 during a data output phase. In one embodiment, a first instance of data driver 162 is configured to operate in a pre-charge phase when a clock signal is in a logical “0” state and a data output phase when the clock signal is in a logical “1” state. A second instance of data driver 162 is configured to operate in a pre-charge phase when the clock signal is in a logical “1” state and a data output phase when the clock signal is in a logical “0” state.

[0048] When each instance of data driver 162 is in the pre-charge phase, if D0 is in a logical “1” state, then switches S22 and S21 are active, while switches S20, S23, and S24 are inactive. While in the pre-charge phase, if D0 is in a logical “0” state, then switches S20 and S23 are active, while switches S21, S22, and S24 are inactive. During a data output phase, switches S21 and S24 are active, while switches S20, S22, and S23 are inactive. In sum, flying capacitor C2 is pre-charged with either a positive or negative polarity charge during the pre-charge phase. The charge is then discharged through ground and Vout 116 during the data output phase.

[0049] FIG. 1G illustrates operation of a ground-referenced single-ended data driver 172 based on a dual capacitor charge pump, in accordance with one embodiment. One or more instances of data driver 172 may be configured to operate as data drivers within a GRS transmitter. For example, an instance of data driver 172 may be configured to operate in place of data driver 112 within GRS transmitter 110 of FIG. 1A. Similarly, an instance of data driver 162 may be configured to operate in place of data driver 114.

[0050] Data driver 172 includes capacitors C3, C4, and switches S30, S31, S32, S33, S40, S41, and S42, configured to pre-charge capacitors C3 and C4 during a pre-charge phase, and discharge one of capacitors C3, C4 into Vout 116 during a data output phase. In one embodiment, a first instance of data driver 172 is configured to operate in a pre-charge phase when a clock signal is in a logical “0” state and a data output phase when the clock signal is in a logical “1” state. A second instance of data driver 172 is configured to operate in a pre-charge phase when the clock signal is in a logical “1” state and a data output phase when the clock signal is in a logical “0” state.

[0051] When each instance of data driver 172 is in the pre-charge phase, switches S30, S33, S40, and S41 are active, and switches S31, S32, and S42 are inactive. During the data output phase, if D0 is in a logical “0” state, then switches S31 and S32 are active, allowing capacitor C3 to discharge a negative polarity charge into Vout 116. At the same time, switches S30, S33, and S40-S42 are inactive. During the data output phase, if D0 is in a logical “1” state, then switches S41 and S42 are active, allowing capacitor C4 to discharge a positive polarity charge into Vout 116. At the same time, switches S40 and S30-S33 are inactive.

[0052] FIG. 2A illustrates an exemplary GRS receiver 130, in accordance with one embodiment. As shown, GRS receiver 130 receives input signals Vin 264 and GRef 266, and generates amplifier output signal 132. In one embodiment, an arriving pulse at Vin 264 having a positive voltage with respect to GRef 266 represents a logical “1” and an arriving pulse at Vin 264 having a negative voltage with respect to GRef 266 rep-

resents a logical “0”. GRS receiver 130 amplifies a differential voltage between input signals Vin 264 and GRef 266 to generate a corresponding difference signal 262. In one embodiment, GRS receiver 130 is designed to bias difference signal 262 to be centered about a switching threshold for inverter inv3, which amplifies difference signal 262 to generate amplifier output signal 132 according to conventional logic voltage levels.

[0053] In one embodiment, GRS receiver 130 comprises resistors R1 through R4, inverters inv1 through inv3, capacitor C5, and field-effect transistors n1 and n2. Resistors R2 and R4 may be implemented as variable resistors, using any technically feasible technique. One exemplary implementation of a variable resistor provides digital control of a resistance value and comprises a set of n-channel FETs connected in a parallel configuration. Each n-channel FET is controlled by a different digital control signal from a control word used to establish the resistance value. If the control word is defined to be a binary number, a corresponding resistance value for the set of n-channel FETs may be monotonic if the n-channel FETs are sized appropriately. In a practical implementation, resistors R2 and R4 are tuned to balance the termination of incoming pulses and current injected into Vin 264 and GRef 266 by GRS receiver 130. A monotonic mapping from a binary code word to a resistance value simplifies any required digital trimming needed to achieve balanced termination. Any technically feasible technique may be implemented to adjust resistors R2 and R4 to achieve balanced termination.

[0054] Resistors R1 and R3 may also be implemented using any technically feasible technique. For example, resistors R1 and R3 may be implemented as p-channel FETs that are biased appropriately. Inverters inv1 and inv2 provide gain, while capacitor C5 serves to stabilize a loop formed by inverters inv1 and inv2, in conjunction with resistor R1 and FET n1.

[0055] FIG. 2B illustrates an exemplary GRS receiver unit 270, configured to demultiplex incoming data, in accordance with one embodiment. GRS receiver unit 270 comprises a GRS receiver 130, and storage elements configured to capture and store the logic state of amplifier output signal 132 on alternating clock phases to demultiplex input data represented as arriving pulses on input signal Vin 264, referenced to input signal GRef 266. Each output signal D0 284 and D1 282 presents captured input data at half the frequency of the arriving data pulses.

[0056] In one embodiment, the storage elements comprise a positive edge triggered flip-flop 274 and a negative edge triggered flip-flop 272. As shown, positive edge triggered flip-flop 274 is configured to capture D0 during the rising edge of a clock signal CLK 268, while negative edge triggered flip-flop 272 is configured to capture D1 during a falling edge of CLK 268. Such a configuration assumes that CLK 268 and amplifier output signal 132 transition together and that flip-flops 272 and 274 require more setup time than hold time. In alternative embodiments, D0 is captured on a falling edge of CLK 268, while D1 is captured on a rising edge of CLK 268. In other alternative embodiments, the storage elements comprise level-sensitive latches rather than flip-flops.

[0057] FIG. 3 illustrates an exemplary transceiver pair 300, configured to implement GRS signaling, in accordance with one embodiment. As shown, the transceiver pair 300 includes transceiver unit 310 coupled to transceiver unit 370 through signal lines 352, 354, 356, and 358. Signal lines 352, 354, 356, and 358 may be manufactured as controlled-impedance traces embedded within an MCM package 190. Transceiver



**310** is configured to receive a reference clock **312** operating at one half the data transmission rate for the signal lines. Adjustable phase delay **332** may introduce an adjustable phase delay prior to transmitting reference clock **312** to GRS transmitter **322**, GRS transmitter **324**, and serializer **334**.

[0058] As shown, the GRS transmitter **322** is configured to transmit a sequential “01” pattern to the GRS receiver **382** through pads **342**, signal line **352**, and pads **362**. In one embodiment, this “01” pattern is transmitted at substantially the same phase as data transmitted from the GRS transmitter **324** to GRS receiver **384** through pads **344**, signal line **354**, and pads **364**. Serializer **334** receives transmit data **314** at a lower frequency than reference clock **312**, but at a correspondingly wider parallel width. For example, if reference clock **312** is configured to operate at 10 GHz, and serializer **334** is configured to multiplex a sixteen bit word into two bits for transmission through GRS transmitter **324**, then sixteen bit words may arrive at a rate of 10 GHz divided by eight or 1.25 GHz. Here, a transmission data clock **313** may be generated by serializer **334** to operate at 1.25 GHz for timing transfers of arriving transmit data **314**. In this example, reference clock **312** has a 100 pS period and each distinct bit transmitted by GRS transmitters **322** and **324** has a unit interval of 50 pS.

[0059] GRS receiver **382** receives a phase-delayed version of reference clock **312** through signal line **352** and generates a local reference clock **383**, which may be coupled to GRS receiver **384** for capturing arriving pulses on signal line **354**. Local reference clock **383** may also be coupled to deserializer **394** for capturing and demultiplexing data from GRS receiver **384**. Extending the above example, GRS receiver **384** may capture arriving pulses on alternating clock phases of local reference clock **383**, operating at 10 GHz, to generate two bits every 100 pS. Deserializer **394** is configured to demultiplex sequential data comprising two bits from GRS receiver **384** and to generate corresponding sixteen-bit words at a rate of 1.25 GHz. The sixteen-bit words are presented as receive data **374**. Deserializer **394** may generate receiver data clock **373** to reflect appropriate clocking for receive data **374**. Receive data **374** represents a local copy of transmit data **314**. In one embodiment, deserializer **394** is configured to align arriving data along word boundaries. Persons skilled in the art will understand that serialization and deserialization of parallel data may require alignment of the parallel data along word boundaries and that well-known techniques in the art may be implemented by transceiver unit **370** or associated logic without departing the scope and spirit of embodiments of the present invention.

[0060] Serializer **396** captures arriving transmit data **376** and serializes the data for transmission by GRS transmitter **386** through signal line **356**. In one embodiment, serializer **396** generates transmit data clock **375** based on local reference clock **383** as a clocking reference for arriving transmit data **376**. GRS receiver **326** captures the data arriving from signal line **356** and deserializer **336** demultiplexes the data into words, presented as receive data **316**. GRS transmitter **388** is configured to transmit a sequential “01” pattern to GRS receiver **328** through pads **368**, signal line **358**, and pads **348**. In one embodiment, this “01” pattern is transmitted at substantially the same phase as data transmitted from GRS transmitter **386** to GRS receiver **326** through pads **366**, signal line **356**, and pads **346**. GRS receiver **328** and adjustable phase delay **338** generate receive clock **318** based on the sequential

“01” pattern. In one embodiment, receive data clock **315** is generated by deserializer **336** to reflect appropriate clocking for receive data **316**.

[0061] Determining a proper phase delay value for adjustable phase delay **332** and adjustable phase delay **338** may be performed using any technically feasible technique. For example, phase delay values for adjustable phase delay **332** and adjustable phase delay **338** may be swept over a range of phase delay values during a link training phase, whereby phase delays corresponding to a substantially minimum bit error rate during training are determined and used for normal link operation.

[0062] Although an isochronous clocking model is illustrated herein for transmitting data between transceiver unit **310** and transceiver unit **370**, any technically feasible clocking model may be implemented without departing the scope and spirit of embodiments of the present invention.

[0063] FIG. 4A illustrates a GRS data driver **400** comprising a CMOS circuit, in accordance with one embodiment. As shown, the CMOS circuit illustrates a circuit topology that may be used to implement the data driver **162** of FIG. 1F using CMOS circuit elements. Specifically, switches S20 and S22 are implemented as p-channel FET p40, and p-channel FET p42, respectively; and switches S21, S23, and S24 are implemented as n-channel FET n41, n-channel FET n43, and n-channel FET n44, respectively. A reference node **410** is coupled to a capacitor C7, p-channel FET p40 and n-channel FET n41. An output node **412** is coupled to an opposing side of capacitor C7, as well as to p-channel FET p42, n-channel FET n43, and n-channel FET n44.

[0064] Control signal g40 is coupled to a gate node of p-channel FET p40. When control signal g40 is driven to a logical 0 level, p-channel FET p40 turns on, pulling node **410** to a voltage level associated with VDD. Control signal g41 is coupled to a gate node of n-channel FET n41. When control signal g41 is driven to a logical 1 level, n-channel FET n41 turns on, pulling node **410** to a voltage level associated with GND. Similarly, p-channel FET p42 responds to control signal g42, selectively pulling node **412** to VDD, while n-channel FET n43 responds to control signal g43, selectively pulling node **412** to GND. Control signal g44 is coupled to a gate node of n-channel FET n44. When control signal g44 is driven to a logical 0 level, n-channel FET n44 substantially isolates node **412** from node Vout **416**. However, when control signal g44 is driven to a logical 1 level, n-channel FET n44 forms a low impedance path between node **412** and Vout **416**. As described below in conjunction with FIG. 4D, this low impedance path facilitates driving Vout **416** with an appropriate signal.

[0065] GRS data driver **400** operates primarily in three different states, including a first pre-charge state for subsequently driving a data value of zero, a second pre-charge state for subsequently driving a data value of one, and a drive state for driving a signal line, such as signal line **105**, with a signal corresponding to a preceding pre-charge state. These states are illustrated below in FIGS. 4B-4D. Transitions between pre-charge states and the drive state are orchestrated by control signals g40 through g44.

[0066] FIG. 4B illustrates GRS data driver **400** in the first pre-charge state that is associated with driving a data value of zero, in accordance with one embodiment. As shown, in the first pre-charge state, control signal g40 is set to zero, to turn on p-channel FET p40, thereby coupling node **410** to VDD. At the same time, control signal g43 is set to one (1), to turn on



n-channel FET **n43**, thereby coupling node **412** to GND. Also, control signal **g42** is set to one to turn off p-channel FET **p42**, and control signals **g41** and **g44** are set to zero to turn off n-channel FET **n41** and n-channel FET **n44**, respectively. In this first pre-charge state, capacitor **C7** is charged with a positive charge on node **410** and a negative charge on node **412**, which is electrically isolated from node Vout **416**.

[0067] FIG. 4C illustrates GRS data driver **400** in the second pre-charge state that is associated with driving a data value of one, in accordance with one embodiment. As shown, in the second pre-charge state, control signal **g42** is set to zero, to turn on p-channel FET **p42**, thereby coupling node **412** to VDD. At the same time, control signal **g41** is set to one, to turn on n-channel FET **n41**, thereby coupling node **410** to GND. Also, control signal **g40** is set to one to turn off p-channel FET **p40**, and control signals **g43** and **g44** are set to zero to turn off n-channel FET **n43** and n-channel FET **n44**, respectively. In this second pre-charge state, capacitor **C7** is charged with a negative charge on node **410** and a positive charge on node **412**, which is electrically isolated from node Vout **416**.

[0068] FIG. 4D illustrates GRS data driver **400** in a drive state, in accordance with one embodiment. As shown, control signal **g41** is set to one, coupling node **410** to GND and control signal **g44** is set to one, coupling node **412** to node Vout **416**. Control signals **g40** and **g42** are set to one, to turn off p-channel FET **p40** and p-channel FET **p42**, respectively. Additionally, control signal **g43** is set to zero, to turn off n-channel FET **n43**. In this state, capacitor **C7** discharges into node Vout **416**. If a negative charge has been accumulated in capacitor **C7** in a previous pre-charge state, then **C7** discharges the negative charge into node Vout **416** with respect to GND. Otherwise, if a positive charge has been accumulated in capacitor **C7** in a previous pre-charge state, then **C7** discharges a positive charge into node Vout **416** with respect to GND. Current passing through node Vout **416** is substantially balanced with a corresponding ground current passing through GND.

[0069] Capacitor **C7** may be implemented using any technically feasible technique without departing the scope and spirit of embodiments of the present invention. In one embodiment, the capacitor **C7** is implemented using n-channel FETs. For example, a gate node of a first n-channel FET may be coupled to node **412** of FIG. 4A to form a back-to-back metal-oxide transistor capacitor. Additionally, source and drain nodes of the first n-channel FET may be coupled to node **410**. A gate node of a second n-channel FET may be coupled to node **410**, while source and drain nodes of the second n-channel FET may be coupled to node **412**. Gate capacitance is relatively area-efficient compared to other capacitor structures available within a CMOS process. However, gate capacitance varies significantly with charge polarity. To compensate for polarity-dependent gate capacitance, two n-channel devices are symmetrically configured to store charge in opposite polarities. In this way, a positive pulse discharged into node Vout **416** has a substantially equal magnitude relative to a negative pulse discharged into Vout **416**.

[0070] In another embodiment, the capacitor **C7** may be implemented using traces in adjacent metal layers. For example, traces in sequential metal layers may be configured to provide plate capacitance ( $C_p$ ) and edge capacitance ( $C_e$ ) between nodes **410** and **412**. Unlike gate capacitance, plate and edge capacitance between metal structures embedded within conventional dielectric materials are stable with

respect to polarity. However, a capacitor formed using metal layer traces may require more die area compared to a capacitor formed using gate capacitance for an equivalent capacitance value. While two parallel traces on two adjacent layers may be used to implement the capacitor **C7**, one skilled in the art will understand that such a metal-oxide-metal (MOM) capacitor can be realized using more than two layers and more than two adjacent traces on each layer.

[0071] FIG. 5A illustrates a GRS transmitter **550** comprising two instances of a GRS data driver **400**, in accordance with one embodiment. As shown, GRS transmitter **550** receives data input signals **D0** and **D1** that are synchronized to clock signal **CLK**. Control logic **502** receives signals **D0**, **D1** and **CLK**, and, in response, generates driver control signals **510** and driver control signals **512**. In one embodiment, driver control signals **510** comprise control signals **g40** through **g44** for instance **400(0)** of GRS data driver **400**, and driver control signals **512** comprise control signals **g40** through **g44** for instance **400(1)** of GRS data driver **400**.

[0072] In one embodiment, when **CLK** is in a logical one state, control logic **502** configures instance **400(0)** to operate in a pre-charge state. If **D0** is in a logical zero state, then instance **400(0)** enters the pre-charge state associated with driving a data value of zero, illustrated previously in FIG. 4B. Here, driver control signals **510** are generated such that **g40=0**, **g41=0**, **g42=1**, **g43=1**, and **g44=0**. If, instead, **D0** is in a logical one state, then instance **400(0)** enters the pre-charge state associated with driving a data value of one, illustrated previously in FIG. 4C. Here, driver control signals **510** are generated such that **g40=1**, **g41=1**, **g42=0**, **g43=0**, and **g44=0**. When **CLK** is in a logical zero state, control logic **502** configures instance **400(0)** to operate in the drive state, illustrated previously in FIG. 4D. Here, driver control signals **510** are generated such that **g40=1**, **g41=1**, **g42=1**, **g43=0**, and **g44=1**.

[0073] When **CLK** is in a logical zero state, control logic **502** configures instance **400(1)** to operate in a pre-charge state. If **D1** is in a logical zero state, then instance **400(1)** enters the pre-charge state associated with driving a data value of zero, illustrated previously in FIG. 4B. Here, driver control signals **512** are generated such that **g40=0**, **g41=0**, **g42=1**, **g43=1**, and **g44=0**. If, instead, **D1** is in a logical one state, then instance **400(1)** enters the pre-charge state associated with driving a data value of one, illustrated previously in FIG. 4C. Here, driver control signals **512** are generated such that **g40=1**, **g41=1**, **g42=0**, **g43=0**, and **g44=0**. When **CLK** is in a logical one state, control logic **502** configures instance **400(1)** to operate in the drive state, illustrated previously in FIG. 4D. Here, driver control signals **510** are generated such that **g40=1**, **g41=1**, **g42=1**, **g43=0**, and **g44=1**.

[0074] Each instance **400(0)**, **400(1)** is coupled to a common Vout **416** signal, which is further coupled to a pads **520**. In one embodiment, Vout **416** is coupled to pads **522** through resistor **RTx**. Pads **522** are coupled to a circuit ground node, corresponding to GND in FIGS. 4A-4D.

[0075] In one embodiment, GRS transmitter **550** is configured to replace GRS transmitter **110** of FIG. 1A. Here, pads **520** couple Vout **416** to signal line **105** and pads **522** couple GND to ground network **107**. In such a configuration, GRS receiver **130** receives data from GRS transmitter **550**. In certain embodiments, GRS transmitter **550** comprises GRS Tx **322**, GRS Tx **324**, GRS Tx **386**, and GRS Tx **388** of FIG. 3.

[0076] FIG. 5B illustrates timing for a GRS transmitter **550**, in accordance with one embodiment. As shown, one bit of data from input **D0** is transmitted to Vout **416** during time



k+1 when CLK is in a logical zero state, and one bit of data from input D1 is transmitted to Vout 416 during time k+2 when CLK is in a logical one state. In one embodiment, inputs D0 and D1 are synchronous to and are updated on the rising edge of CLK. In such an embodiment, instance 400(1) is in a data driving state when inputs D0 and D1 change in response to a rising edge of CLK going into time k. On the rising edge of CLK going into time k, instance 400(0) enters a pre-charge state, thereby sampling data on D0. On the falling edge of CLK exiting time k and entering time k+1, instance 400(0) enters a data driving state and drives the captured data from D0 onto Vout 416. On the falling edge of CLK going into time k+1, instance 400(1) enters a pre-charge state, thereby sampling data on D1. On the rising edge of CLK exiting time k+1 and entering time k+2, instance 400(1) enters a data driving state and drives the captured data from D1 onto Vout 416. In this way, data comprising D0 and D1 may be presented to GRS transmitter 550 using conventional logic having conventional single-edge synchronous timing, while GRS transmitter 550 time-multiplexes the data for transmission at a double data rate. In other words, two data transfers occur in each period or cycle of the CLK. In a preferred embodiment, D0 is latched when CLK is low to ensure that D0 is stable while being used to control the pre-charge of instance 400(0). Similarly, D1 is latched when CLK is high to ensure D1 is stable while being used to control the pre-charge of instance 400(1).

[0077] In other embodiments, a GRS transmitter comprising more than two instances of GRS data driver 400 is configured to receive a data bit per instance of GRS data driver 400 and to time-multiplex the data at a correspondingly higher data rate. In such embodiments, multiple clock signals may be required to provide appropriate timing for pre-charging and driving data to time-multiplex the data.

[0078] FIG. 5C illustrates a flow chart of a method 560 for generating a ground-referenced single-ended signal, in accordance with one embodiment. Although method 560 is described in conjunction with FIGS. 4A-5B to implement a two to one time-multiplexing ratio of input data to output data, persons of ordinary skill in the art will understand that any system that performs method 560 is within the scope and spirit of embodiments of the present invention.

[0079] Method 560 begins in step 565, where a first data driver, such as instance 400(0) of GRS data driver 400, samples a first bit of data by pre-charging a first capacitor during a first time k. The first capacitor is charged to have a polarity corresponding to a logic level for the first bit of data. In step 570, a second data driver, such as instance 400(1) of GRS data driver 400, samples a second bit of data by pre-charging a second capacitor during a time k+1. The second capacitor is charged to have a polarity corresponding to a logic level for the second bit of data.

[0080] In step 575, the first data driver drives an output signal, such as Vout 416 of FIGS. 4A-4D or Vout 416 of FIG. 5A, to reflect the first bit of data by coupling the first capacitor to the output signal during the time k+1. Here, the first capacitor is coupled between a ground network and the output signal. The polarity of charge on the first capacitor was established in step 565, based on the logic level for the first bit of data. When coupled to the output signal, the first capacitor therefore reflects the logic level for the first bit of data.

[0081] In step 580, the second data driver drives the output signal to reflect the second bit of data by coupling the second capacitor to the output signal during a time k+2. Here, the second capacitor is coupled between a ground network and

the output signal. The polarity of charge on the second capacitor was established in step 570, based on the logic level for the second bit of data. When coupled to the output signal, the second capacitor therefore reflects the logic level for the first bit of data. Method 560 terminates after driving the output signal to reflect the second bit of data.

[0082] In other embodiments, a time-multiplexing ratio of greater than two may be implemented and at least one additional phase-related clock may be provided to orchestrate operation of more than three instances of GRS data driver 400.

#### Multiprocessor System with Ground-Referenced Signaling

[0083] FIG. 6A illustrates a multiprocessor system implemented as a multi-chip module (MCM) 600, in accordance with one embodiment. As shown, MCM 600 comprises an MCM package 190, two or more multi-core processor (MCP) chips 610, and an interconnect 614, configured to facilitate communication between and among MCP chips 610. Each MCP chip 610 may include one or more processor cores. Each MCP chip 610 may also include cache memory for each processor core, as well as cache memory shared by two or more processor cores. For example, each MCP chip 610 may include a first level cache associated with each processor core. Each MCP chip 610 may also include a second level cache shared among one or more processor cores included within MCP chip 610. In certain embodiments, one or more processor cores within an MCP chip 610 are configured to include a vector processor unit (not shown). In one embodiment, MCP chip 610 is configured to include a digital signal processing (DSP) core (not shown).

[0084] In certain embodiments, MCP chip 610(0) is configured to provide high computational performance, while MCP chip 610(1) is configured to provide low power consumption. In such embodiments, MCP chip 610(0) may be fabricated from a high-performance fabrication technology, while MCP chip 610(1) may be fabricated from a low-power fabrication technology. In certain embodiments, MCP chip 610(0) is designed for relatively high performance, while MCP chip 610(1) is designed using the same fabrication technology for relatively low power. In one embodiment, MCP chip 610(0) includes four or more high-performance processor cores, while MCP chip 610(1) includes four or fewer processor cores configured to operate in a low-power mode.

[0085] Each MCP chip 610 is coupled to an interconnect 614 through a corresponding interconnect link 612. As illustrated in greater detail in FIGS. 6B-6D, interconnect 614 may implement different topologies that facilitate communication among MCP chips 610. In one embodiment, each interconnect link 612 comprises one or more GRS transceivers disposed within a corresponding MCP chip 610, and associated electrical traces manufactured within MCM package 190. Each GRS transceiver may include one or more bidirectional or one or more unidirectional data signals, according to implementation-specific requirements.

[0086] Any technically feasible communication protocol may be implemented for transmitting data over interconnect links 612. In one embodiment, the communication protocol specifies, without limitation, a memory read request that includes an access address, a read response (acknowledgement) that includes requested read data, a memory write request that includes an access address and write data, and a



write acknowledgement that indicates a successful write operation. In certain embodiments, the read request and the write request also include an access length specified in bytes, words, or any technically feasible measure of data length. In one embodiment, a given access request comprises a split transaction. In an alternative embodiment, a given access request comprises a blocking transaction. In certain embodiments, the communication protocol specifies a message passing mechanism for transmitting data packets to a destination device. In one embodiment, the communication protocol, implemented for transmitting data through interconnect links **612**, specifies a cache coherence protocol. The cache coherence protocol may provide a broadcast mechanism for maintaining cache coherence, a directory-based mechanism for maintaining cache coherence, or any technically feasible mechanism for maintaining cache coherence among two or more caches or memory subsystems without departing the scope and spirit of embodiments of the present invention. In one embodiment, the cache coherence protocol implements an invalidation mechanism for processing cache writes. Alternatively, the cache coherence protocol implements an update mechanism for processing cache writes. In one embodiment, the cache coherence protocol implements a write-through mechanism for processing certain writes.

[0087] MCM **600** may also include one or more memory subsystems **620**, coupled to MCP chips **610**, either directly or through interconnect **614**. In one embodiment, each memory subsystem **620** comprises a DRAM chip. In another embodiment, each memory subsystem **620** comprises a cache memory chip. The cache memory chip may comprise a second level cache, a third level cache, a cache slice, or any other technically feasible cache memory element. In yet another embodiment, each memory subsystem **620** comprises a stack of memory chips including at least one DRAM chip, or at least one cache memory chip, or a combination thereof. In still other embodiments, each memory subsystem **620** comprises an interface shim chip and at least one DRAM chip, at least one cache memory chip, or at least one DRAM chip and at least one cache memory chip. The interface shim chip may include a memory controller, configured to receive access requests (commands), and process the access requests by generating further access requests that directly target DRAM chips or cache memory chips coupled to the shim chip. In certain embodiments, memory subsystem **620** is configured to communicate through a GRS transceiver comprising one or more data signals to at least one MCP chip **610**. In such embodiments, a given memory subsystem **620** may be coupled directly to the at least one MCP chip **610**, or the memory subsystem may be coupled indirectly through interconnect **614**.

[0088] MCM **600** may also include a system functions chip **618**, coupled to MCP chips **610**. System functions chip **618** may also be coupled to memory subsystems **620**. In one embodiment, system functions chip **618** is configured to implement functionality required by MCM **600**, but not implemented in MCP chips **610**, memory subsystems **620**, or interconnect **614**. For example, system functions chip **618** may implement power management functions, interface functions, system control functions, and watchdog functions, or any combination thereof in conjunction with the operation of MCP chips **610**.

[0089] FIG. 6B illustrates a directly-connected multiprocessor system implemented as MCM **600**, in accordance with one embodiment. As shown, interconnect **614** is configured to

directly connect each MCP chip **610** to each other MCP chip **610**. In the directly-connected topology shown in FIG. 6B, each interconnect link **612** illustrated in FIG. 6A comprises direct-connection links. Specifically, interconnect link **612** (0) comprises link **630**(0), link **630**(3), and link **630**(5); interconnect link **612**(1) comprises link **630**(0), link **630**(1), and link **630**(4); interconnect link **612**(2) comprises link **630**(2), link **630**(3), and link **630**(4); and, interconnect link **612**(3) comprises link **630**(1), link **630**(2), and link **630**(5). In one embodiment, a given link **630** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in one MCP chip **610**, and a second GRS transceiver of the pair of GRS transceivers is included in a different MCP chip **610**. Associated electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link.

[0090] In one embodiment, links **630** comprise independent channels of an interconnect link. In such an embodiment, links **630** implement a communication protocol consistent with a communication protocol for interconnect link **612**. In other embodiments, each link **630** is configured to operate as an independent interconnect link **612**.

[0091] In one embodiment, each MCP chip **610** is coupled to a corresponding memory subsystem **620** through an associated memory link **622**. In certain embodiments, each memory link **622** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in an MCP chip **610**, and a second GRS transceiver of the pair of GRS transceivers is included in a chip comprising a corresponding memory subsystem **620**. Associated electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link. As described previously, memory subsystem **620** may comprise at least one memory chip, such as a DRAM or cache memory chip. The at least one memory chip may be assembled into a stack. In certain embodiments, an MCP chip **610** may be coupled directly to an additional memory subsystem **620** (not shown).

[0092] In one embodiment, each MCP chip **610** is configured to transmit a memory access protocol over a corresponding memory link **622** that specifies, without limitation, a memory read request configured to include an access address, and a memory write request configured to include an access address and write data. In certain embodiments, the read request and the write request also include an access length specified in bytes, words, or any technically feasible measure of data length.

[0093] FIG. 6C illustrates a hub-connected multiprocessor system implemented as an MCM **600**, in accordance with one embodiment. As shown, interconnect **614** is configured to include a hub chip **640** coupled to each MCP chip **610**. In the hub-connected topology shown in FIG. 6C, each interconnect link **612** represents a connection to hub chip **640**. Specifically, interconnect link **612**(0) couples MCP chip **610**(0) to hub chip **640**; interconnect link **612**(1) couples MCP chip **610**(1) to hub chip **640**; interconnect link **612**(2) couples MCP chip **610**(2) to hub chip **640**; interconnect link **612**(3) couples MCP chip **610**(3) to hub chip **640**. In one embodiment, a given interconnect link **612** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in an MCP chip **610**, and a second, corresponding GRS transceiver is included in hub chip **640**. Associated



electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link.

[0094] In one embodiment, each MCP chip **610** is coupled to a memory subsystem **620** through hub chip **640**. Each memory subsystem **620** is coupled to hub chip **640** through a corresponding memory link **622**. In one embodiment, each memory link **622** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in hub chip **640**, and a second GRS transceiver of the pair of GRS transceivers is included in a chip comprising a corresponding memory subsystem **620**. Associated electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link. As described previously, memory subsystem **620** may comprise at least one memory chip, such as a DRAM or cache memory chip. The at least one memory chip may be assembled into a stack.

[0095] In certain embodiments, an MCP chip **610** may be coupled directly to an additional memory subsystem **620** (not shown) through a memory link **622**. In alternative embodiments, each memory subsystem **620** is coupled directly to a corresponding MCP chip **610** through an associated memory link **622**.

[0096] In one embodiment, each MCP chip **610** is configured to transmit a memory access protocol over a corresponding memory link **622** that specifies, without limitation, a memory read request that includes an access address, and a memory write request that includes an access address and write data. In certain embodiments, the read request and the write request also include an access length specified in bytes, words, or any technically feasible measure of data length. In certain embodiments, additional memory subsystems (not shown) are coupled directly to each corresponding MCP chip **610**.

[0097] Hub chip **640** may implement any technically feasible internal communication topology, such as a crossbar, ring, butterfly, Clos, or general mesh network to interconnect links **612** and memory links **622**. Any technically feasible admission control and arbitration mechanism may be implemented for managing and arbitrating ingress to egress traffic. Although MCM **600** is shown in FIG. 6C as comprising four MCP chips **610(0)**-**610(3)**, any number of MCP chips may be included within MCM **600** and coupled to hub chip **640**. Similarly, any number of memory subsystems **620** may be included within MCM **600** and coupled to hub chip **640**.

[0098] FIG. 6D illustrates a network-connected multiprocessor system implemented as an MCM **600**, in accordance with one embodiment. As shown, interconnect **614** comprises two router chips **650** coupled to each other, and to associated MCP chips **610**. In the network-connected topology shown in FIG. 6D, each interconnect link **612** represents a connection to a corresponding router chip **650**. Specifically, interconnect link **612(0)** couples MCP chip **610(0)** to router chip **650(0)**; interconnect link **612(1)** couples MCP chip **610(1)** to router chip **650(0)**; interconnect link **612(2)** couples MCP chip **610(2)** to router chip **650(1)**; and, interconnect link **612(3)** couples MCP chip **610(3)** to router chip **650(1)**. In one embodiment, a given interconnect link **612** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in an MCP chip **610**, and a second, corresponding GRS transceiver is included in a corresponding router chip **650**. Associated electrical traces

manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link.

[0099] In one embodiment, each MCP chip **610** is coupled to a memory subsystem **620** through a collection of two or more router chips **650**. Each memory subsystem **620** is coupled to a router chip **650** through a corresponding memory link **622**. In one embodiment, each memory link **622** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in a router chip **650**, and a second GRS transceiver of the pair of GRS transceivers is included in a chip comprising a corresponding memory subsystem **620**. Associated electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link. As described previously, memory subsystem **620** may comprise at least one memory chip, such as a DRAM or cache memory chip. The at least one memory chip may be assembled into a stack.

[0100] In certain embodiments, an MCP chip **610** may be coupled directly to an additional memory subsystem **620** (not shown) through a memory link **622**. In alternative embodiments, each memory subsystem **620** is coupled directly to a corresponding MCP chip **610** through an associated memory link **622**.

[0101] In one embodiment, each MCP chip **610** is configured to transmit a memory access protocol over a corresponding memory link **622** that specifies, without limitation, a memory read request that includes an access address, and a memory write request that includes an access address and write data. In certain embodiments, the read request and the write request also include an access length specified in bytes, words, or any technically feasible measure of data length.

[0102] During normal operation, router chip **650(0)** may receive a data packet from a source device for delivery to a destination device. The source device and the destination device may each separately comprise an MCP chip **610**, a memory subsystem **620**, a system functions chip **618**, or any other technically feasible destination device. The data packet may comprise a read request, a write request, acknowledgment to a previous, a data message, a command, or any other technically feasible unit of information. Router chip **650(0)** is configured to forward the data packet to the destination device along a forwarding path. The forwarding path may include, without limitation, an interconnect link **612**, a memory link **622**, an inter-router link **652**, or any technically feasible combination thereof. If the source device and the destination device are both directly coupled to router chip **650(0)**, then router chip **650(0)** may forward the data packet directly from the source device to the destination device. If the destination device is instead directly coupled to router chip **650(1)**, then the router chip **650(0)** forwards the data packet through inter-router link **652** to router chip **650(1)**, which then forwards the data packet to the destination device. In one embodiment, inter-router link **652** comprises a pair of GRS transceivers. A first GRS transceiver of the pair of GRS transceivers is included in a router chip **650(0)**, and a second GRS transceiver of the pair of GRS transceivers is included in router chip **650(1)**. Associated electrical traces manufactured within MCM package **190** couple the first GRS transceiver to the second GRS transceiver to complete the link.

[0103] As shown, MCM **600** includes two router chips **650(0)**, **650(1)**, configured to form a two node multi-hop network. However, MCM **600** may include an arbitrary number of router chips **650**, interconnected through a correspond-



ing set of GRS transceivers to form an arbitrary multi-hop network topology such as a mesh, torus, butterfly, or Clos without departing the scope and spirit of embodiments of the present invention.

[0104] A GRS transceiver within the source device includes a GRS transmitter that is configured to transmit serialized data comprising the data packet until the data packet is transmitted in full to a GRS receiver within the destination device or within an interconnection chip preparing to forward the data packet. The GRS transmitter may implement two or more degrees of multiplexing by implementing a corresponding number of GRS data drivers and appropriate clocking circuitry. The GRS transmitter may be configured to perform method 560 to generate individual bits comprising the serialized data for transmission. Exemplary GRS transmitters illustrated in FIGS. 1A-5B implement two to one multiplexing; however, persons skilled in the art will understand that arbitrary degrees of multiplexing may be similarly implemented without departing the scope and spirit of embodiments of the present invention. The destination device may deserialize the serialized bits to construct the access request. If the destination device is configured to serve as a bridge or hub, then the access request may be forwarded to a destination device for processing. In certain embodiments, method 560 of FIG. 5C is performed to generate GRS signals transmitted over one or more interconnect links 612, one or more memory links 622, one or more inter-router link 652, or any combination thereof.

[0105] In one embodiment, interconnect links 612, memory links 622, and inter-router link 652 are implemented as electrical traces within MCM package 190. Each trace may comprise a conductive element affixed to a dielectric substrate, such as an organic substrate layer of MCM package 190. Each electrical trace may be configured to exhibit a controlled electrical impedance.

[0106] In one embodiment, the data packet is generated by the source device for transmission to a destination device for processing. Certain data packets comprise a set of request fields, including, without limitation, an address field, which may uniquely identify the destination device and a specific address within the destination device. The access request is transmitted over a GRS interconnect to the destination device.

[0107] In one embodiment, a non-transitory computer readable medium is configured to represent a detailed design of MCM package 190, including all electrical connections. Such electrical connections include electrical traces designed to support ground-referenced single-ended signals, including, without limitation, interconnect links 612, memory links 622, and inter-router link 652. Each GRS interconnect may include an abstract representation of connectivity, such as connectivity represented within a net list. Individual traces may be represented as code within a net list file. Persons skilled in the art will understand that many net list formats are available, and any technically feasible non-transitory computer readable medium configured to represent system-on-package 600 is within the scope and spirit of the present invention.

[0108] FIG. 7 illustrates an exemplary system 700 in which the various architecture and/or functionality of the various previous embodiments may be implemented. As shown, a system 700 is provided including at least one central processor 701 that is connected to a communication bus 702. The communication bus 702 may be implemented using any suitable protocol, such as PCI (Peripheral Component Intercon-

nect), PCI-Express, AGP (Accelerated Graphics Port), HyperTransport, or any other bus or point-to-point communication protocol(s). The system 700 also includes a main memory 704. Control logic (software) and data are stored in the main memory 704, which may take the form of random access memory (RAM). In one embodiment, central processor 701, graphics processor 706, a portion of bus 702 configured to interconnect the central processor 701 and graphics processor 706, and at least a portion of main memory 704 comprise a system-on-package, such as system-on-package 600 of FIGS. 6A, 6B, and 6C.

[0109] The system 700 also includes input devices 712, a graphics processor 706, and a display 708, i.e. a conventional CRT (cathode ray tube), LCD (liquid crystal display), LED (light emitting diode), plasma display or the like. User input may be received from the input devices 712, e.g., keyboard, mouse, touchpad, microphone, and the like. In one embodiment, the graphics processor 706 may include a plurality of shader modules, a rasterization module, etc. Each of the foregoing modules may even be situated on a single semiconductor platform to form a graphics processing unit (GPU).

[0110] In the present description, a single semiconductor platform may refer to a sole unitary semiconductor-based integrated circuit or chip. It should be noted that the term single semiconductor platform may also refer to multi-chip modules with increased connectivity which simulate on-chip operation, and make substantial improvements over utilizing a conventional central processing unit (CPU) and bus implementation. Of course, the various modules may also be situated separately or in various combinations of semiconductor platforms per the desires of the user.

[0111] The system 700 may also include a secondary storage 710. The secondary storage 710 includes, for example, a hard disk drive and/or a removable storage drive, representing a floppy disk drive, a magnetic tape drive, a compact disk drive, digital versatile disk (DVD) drive, recording device, universal serial bus (USB) flash memory. The removable storage drive reads from and/or writes to a removable storage unit in a well-known manner. Computer programs, or computer control logic algorithms, may be stored in the main memory 704 and/or the secondary storage 710. Such computer programs, when executed, enable the system 700 to perform various functions. The main memory 704, the storage 710, and/or any other storage are possible examples of computer-readable media.

[0112] In one embodiment, the architecture and/or functionality of the various previous figures may be implemented in the context of the central processor 701, the graphics processor 706, an integrated circuit (not shown) that is capable of at least a portion of the capabilities of both the central processor 701 and the graphics processor 706, a chipset (i.e., a group of integrated circuits designed to work and sold as a unit for performing related functions, etc.), and/or any other integrated circuit for that matter.

[0113] Still yet, the architecture and/or functionality of the various previous figures may be implemented in the context of a general computer system, a circuit board system, a game console system dedicated for entertainment purposes, an application-specific system, and/or any other desired system. For example, the system 700 may take the form of a desktop computer, laptop computer, server, workstation, game consoles, embedded system, and/or any other type of logic. Still yet, the system 700 may take the form of various other devices



including, but not limited to a personal digital assistant (PDA) device, a mobile phone device, a television, etc.

[0114] Further, while not shown, the system 700 may be coupled to a network (e.g., a telecommunications network, local area network (LAN), wireless network, wide area network (WAN) such as the Internet, peer-to-peer network, cable network, or the like) for communication purposes.

[0115] In one embodiment, certain signals within bus 702 are implemented as GRS signals, as described above in FIGS. 1A-6D.

[0116] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A system, comprising:
  - a first processor chip configured to include a first ground-referenced single-ended signaling (GRS) interface circuit;
  - a second processor chip configured to include a second GRS interface circuit;
  - a multi-chip module (MCM) package configured to include the first processor chip, the second processor chip, and an interconnect circuit;
  - a first set of electrical traces fabricated within the MCM package and configured to couple the first GRS interface circuit to the interconnect circuit; and
  - a second set of electrical traces fabricated within the MCM package and configured to couple the second GRS interface circuit to the interconnect circuit.
2. The system of claim 1, wherein the first GRS interface circuit comprises:
  - a first GRS driver circuit, configured to:
    - pre-charge a first capacitor to store a first charge during a first pre-charge phase; and
    - drive an output signal relative to a ground network based on the first charge during a first drive phase;
  - a second GRS driver circuit, configured to:
    - pre-charge a second capacitor to store a second charge during a second pre-charge phase; and
    - drive the output signal relative to a ground network based on the second charge during a second drive phase; and
  - a receiver circuit, configured to translate a ground-referenced single-ended input signal to a corresponding logic signal,
 wherein the first set of electrical traces comprise the input signal, the output signal, and the ground network.
3. The system of claim 1, wherein the first processor chip comprises a single processor core and a first level cache.
4. The system of claim 1, wherein the first processor chip comprises two or more processor cores and corresponding first level caches.
5. The system of claim 4, wherein the first processor chip further comprises a vector processor core.
6. The system of claim 4, wherein the first processor chip further comprises a digital signal processor core.
7. The system of claim 1, wherein the first processor chip is configured to operate at relatively high processing throughput

relative to the second processor chip, which is configured to operate at lower throughput and lower power relative to the first processor chip.

8. The system of claim 1, wherein the first processor chip is manufactured from a high-performance fabrication process and the second processor chip is manufactured from a low-power fabrication process.

9. The system of claim 1, further comprising:
 

- a first memory subsystem included within the MCM package and configured to include a third GRS interface circuit;
- a fourth GRS interface circuit included within the first processor chip; and
- a third set of electrical traces fabricated within the MCM package and configured to couple the third GRS interface circuit to the fourth GRS interface circuit.

10. The system of claim 9, wherein the first memory subsystem comprises at least two stacked chips.

11. The system of claim 9, wherein the first memory subsystem comprises a cache memory circuit.

12. The system of claim 9, wherein the memory subsystem comprises:

- a shim chip that includes the third GRS interface circuit and a memory controller circuit coupled to the third GRS interface circuit; and
  - at least one memory chip coupled to the memory controller circuit,
- wherein the memory controller circuit transmits data associated with memory access requests between the third GRS interface circuit and the at least one memory chip.

13. The system of claim 1, wherein the interconnect circuit comprises a first interconnect chip included within the MCM package, and configured to transmit data between the first processor chip and the second processor chip.

14. The system of claim 13, further comprising:
 

- a first memory subsystem included within the MCM package and configured to include a third GRS interface circuit;
- a fourth GRS interface circuit included within the first interconnect chip; and
- a third set of electrical traces fabricated within the MCM package and configured to couple the third GRS interface circuit to the fourth GRS interface circuit.

15. The system of claim 13, further comprising:
 

- a third GRS interface circuit included within the first interconnect chip;
- a second interconnect chip included within the MCM package and configured to include a fourth GRS interface circuit and a fifth GRS interface circuit;
- a third set of electrical traces fabricated within the MCM package and configured to couple the third GRS interface circuit to the fourth GRS interface circuit;
- a third processor chip configured to include a sixth GRS interface circuit; and
- a fourth set of electrical traces fabricated within the MCM package and configured to couple the fifth GRS interface circuit to the sixth GRS interface circuit,

wherein the second interconnect chip is configured to transmit data between the third processor and the first interconnect chip.

16. The system of claim 15, further comprising:
 

- a first memory subsystem included within the MCM package and configured to include a seventh GRS interface circuit;

an eight GRS interface circuit included within the first interconnect chip; and  
 a fifth set of electrical traces fabricated within the MCM package and configured to couple the seventh GRS interface circuit to the eight GRS interface circuit,  
 wherein the first interconnect chip is configured to transmit data between the memory subsystem and the first processor chip, and transmit data between the memory subsystem and the second interconnect chip.

**17.** The system of claim **1**, wherein the interconnect circuit comprises the first set of electrical traces and the second set of electrical traces.

**18.** A non-transitory computer readable medium, comprising:

code representing a first set of electrical traces configured to couple a first processor chip to a interconnection circuit within a multi-chip module (MCM) package; and  
 code representing a second set of electrical traces configured to couple a second processor chip to the interconnection circuit within the MCM package,

wherein the first set of electrical traces and the second set of electrical traces comprise ground-referenced single-ended (GRS) signal lines.

**19.** The non-transitory computer readable medium of claim **18**, further comprising:

code representing a third set of electrical traces configured to couple the first processor chip to a first memory subsystem within a multi-chip module (MCM) package; and

code representing a second set of electrical traces configured to couple the second processor chip to a second memory subsystem within the MCM package.

**20.** The non-transitory computer readable medium of claim **18**, wherein the interconnection circuit comprises a first interconnect chip.

**21.** The non-transitory computer readable medium of claim **18**, wherein the interconnection circuit comprises a first interconnection chip coupled to a second interconnection chip through a third set of electrical traces.

\* \* \* \* \*